

NUMERICAL MATHEMATICS AND SCIENTIFIC COMPUTATION

Books in the series

Monographs marked with an asterisk (*) appeared in the series 'Monographs in Numerical Analysis' which has been folded into, and is continued by, the current series.

- *P. Dierckx: *Curve and surface fittings with splines*
- *J. H. Wilkinson: *The algebraic eigenvalue problem*
- *I. Duff, A. Erisman, & J. Reid: *Direct methods for sparse matrices*
- *M. J. Baines: *Moving finite elements*
- *J. D. Pryce: *Numerical solution of Sturm-Liouville problems*
- K. Burrage: *Parallel and sequential methods for ordinary differential equations*
- Y. Censor & S.A. Zenios: *Parallel optimization: theory, algorithms and applications*
- M. Ainsworth, J. Levesley, W. Light & M. Marletta: *Wavelets, multilevel methods and elliptic PDEs*
- W. Freeden, T. Gervens, & M. Schreiner: *Constructive approximation on the sphere: theory and applications to geomathematics*
- Ch. Schwab: *p- and hp- finite element methods: theory and applications to solid and fluid mechanics*
- J.W. Jerome: *Modelling and computation for applications in mathematics, science, and engineering*
- Alfio Quarteroni & Alberto Valli: *Domain decomposition methods for partial differential equations*
- G. E. Karniadakis & S. J. Sherwin: *Spectral/hp element methods for CFD*
- I. Babuška & T. Strouboulis: *The finite element method and its reliability*
- B. Mohammadi & O. Pironneau: *Applied shape optimization for fluids*
- S. Succi: *The Lattice Boltzmann Equation for fluid dynamics and beyond*
- P. Monk: *Finite element methods for Maxwell's equations*
- A. Bellen & M. Zennaro: *Numerical methods for delay differential equations*
- J. Modersitzki: *Numerical methods for image registration*

Numerical Methods for Image Registration

Jan Modersitzki

Institute of Mathematics, University of Lübeck

OXFORD
UNIVERSITY PRESS

OXFORD UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan South Korea Poland Portugal
Singapore Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© Oxford University Press 2004

The moral rights of the author have been asserted
Database right Oxford University Press (maker)

First published 2004

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above.

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer.

A catalogue record for this book is available from the British Library

Library of Congress Cataloging in Publication Data
(Data available)

ISBN 0 19 852841 8

10 9 8 7 6 5 4 3

Typeset by Newgen Imaging Systems (P) Ltd., Chennai, India
Printed in Great Britain
on acid-free paper by
Biddles Ltd., King's Lynn, Norfolk

ACKNOWLEDGEMENTS

Thanks to *Bernad Visco* for making me wiser.

Thanks to *Oliver Schmitz* for making me see.

Thanks to *Stephen Seligson* for making me speak.

Thanks to *Charles di Paolo* for making me happy.

Also thanks to Martin Böhme, Matthias Bolten, Lutz Dümbgen, Jan Ehrhardt,
Andre Folkers, Jan-Michael Frahm, Stefan Heldmann, Oliver Mahnke, Markus
von Oehsen, Nils Papenberg, Hauke Prenzel, Swen Rist, Dirk Rutsatz, Bodo
Siebert, Stefan Wirtz, and Maike Wolf.

CONTENTS

1	Introduction	1
1.1	The problem	1
1.2	The intention	4
1.3	An overview	5
2	The Human NeuroScanning Project	8
3	The mathematical setting	14
3.1	General remarks	14
3.1.1	Introduction to images	14
3.1.2	From continuous to discrete images	15
3.1.3	From discrete to continuous images	17
3.2	Mathematical notation	19
3.3	The registration problem	21
3.3.1	Restricted transformations	22
3.3.2	The Lagrange and Euler frames	22
Part I	Parametric image registration	25
4	Landmark-based registration	27
4.1	General remarks	27
4.2	Landmark-based parametric registration	28
4.3	Landmark-based smooth registrations	30
4.3.1	Minimal norm solutions for interpolation problems	32
4.3.2	Example: splines and TPSs	38
4.3.3	Smooth minimal norm approximations	41
4.3.4	Smooth transformations	42
4.4	An example of landmark-based registration	44
5	Principal axes-based registration	45
5.1	Geometrical interpretation	46
5.2	Stochastic interpretation	47
5.3	A robust generalization	48
5.4	A simple comparison with $d = 2$	49
5.5	Principal axes under transformation	49
5.6	An example of principal axis-based registration	53
6	Optimal linear registration	55
6.1	Intensity-based registration	56
6.2	An example of intensity-based affine linear registration	58

6.3	Correlation-based registration	59
6.4	Mutual information-based registration	60
6.4.1	Mutual information	60
6.4.2	Gradient of mutual information	66
6.5	An example of mutual information-based affine linear registration	69
6.6	An example of optimal affine linear registration	72
7	Summary of parametric image registration	73
Part II	Non-parametric image registration	75
8	Non-parametric image registration	77
8.1	A general framework	77
8.2	General solution schemes	79
8.3	Computing the forces	80
8.4	General remarks on choosing the discretization	81
9	Elastic registration	83
9.1	Physical motivation	84
9.1.1	Deformation of a body	84
9.1.2	The strain tensor	85
9.1.3	The stress tensor	88
9.1.4	HOKE's law	91
9.2	Eigenfunctions of the Navier-Lamé operator	94
9.2.1	Sliding boundary conditions	96
9.2.2	Bending boundary conditions	97
9.2.3	Periodic boundary conditions	97
9.3	Variational formulation of the elasticity problem	98
9.4	Finite difference approximation of the Navier-Lamé equations	101
9.4.1	The two-dimensional case	101
9.4.2	The three-dimensional case	103
9.5	Factorization of a structured matrix	104
9.5.1	The three-dimensional case	105
9.5.2	The two-dimensional case	109
9.5.3	Inversion of the structured matrices	109
9.5.4	Inversion of the discrete Navier-Lamé equations	111
9.6	Solving the Navier-Lamé equations	113
9.7	Regularizing the updates of the displacements	113
9.8	Elastic registration: $\bullet$ to C	113
9.9	Elastic registration: hands	114
9.10	Elastic registration: brain	115
9.11	Discussion of elastic registration	119

10	Fluid registration	121
10.1	Variational motivation of fluid registration	121
10.2	Physical motivation of fluid registration	121
10.2.1	The basics of fluid mechanics	122
10.2.2	Stokes fluids	124
10.3	Solving the Navier-Lamé equations for the velocity	124
10.4	Implementation details	127
10.4.1	Euler step and regridding	127
10.4.2	Interpolation and regridding	128
10.4.3	Computing forces and velocities	129
10.4.4	Stopping criteria	129
10.5	Fluid registration: $\bullet$ to C	129
10.6	Fluid registration: hands	131
10.7	Fluid registration: hand to disk	133
10.8	Discussion of fluid registration	133
11	Diffusion registration	137
11.1	Continuous and discrete Laplace equations	138
11.2	Factorization of the discrete Laplace operator	139
11.3	Additive operator splitting (AOS)	144
11.3.1	Matrix analysis of AOS	146
11.3.2	Diffusion registration using AOS	148
11.4	Diffusion registration: $\bullet$ to C	150
11.5	Diffusion registration: hands	153
11.6	Thirion's demons registration	155
11.6.1	Demons type	159
11.6.2	Variational interpretation of Thirion's approach	161
11.6.3	Steepest descent interpretation	162
11.7	Discussion of diffusion registration	162
12	Curvature registration	163
12.1	Continuous and discrete bi-harmonic equations	164
12.2	A relation to thin plate splines	166
12.3	Curvature registration: squares	167
12.4	Curvature registration: $\bullet$ to C	168
12.5	Curvature registration: hands	171
12.6	Discussion of curvature registration	173
13	Concluding remarks	175
13.1	Summary of non-parametric image registrations	176
13.2	Timings for non-parametric registrations	177
13.3	A competitive example: hands	180
13.4	How to compare apples with pears	180

13.5 How to choose μ and λ or α and τ	183
13.6 Further accelerations of the schemes	184
13.7 Extensions	185
References	186
Glossary	191
Author Index	195
Index	197

INTRODUCTION

1.1 The problem

In image processing one is often interested not only in analyzing one image but in comparing or combining the information given by different images. For this reason, *image registration* is one of the fundamental tasks within image processing. The task of image registration is to find an *optimal geometric transformation* between *corresponding* image data. In practice, the concrete type of the *geometric transformation* as well as the notions of *optimal* and *corresponding* depend on the specific application.

Image registration is a problem often encountered in many application areas like, for example, astro- and geophysics, computer vision, and medicine. For an overview, see, e.g., Brown (1992), Maintz & Viergever (1998), Maurer & Fitzpatrick (1993), van den Elsen et al (1993), and references therein.

Here, we focus on medical applications. In the last two decades, computerized image registration has played an increasingly important role particularly in medical imaging. Registered images are now used routinely in a multitude of different applications such as the treatment verification of pre- and post-intervention images and time evolution of an injected agent subject to patient motion. Image registration is also useful to take full advantage of the complementary information coming from multimodal imagery, like, for example, computer tomography (CT) and magnetic resonance imaging (MRI).

However, the interpretation of medical images and of the registration result typically requires expert knowledge. For this reason we use the simple test images shown in Fig. 1.1 throughout this book, where even a non-expert has an intuitive understanding of the outcome of a registration procedure. Here, two (modified) X-ray images of a human hand are depicted; see also Amit (1994). Note that a comparison of these two images would be much easier if they could have been aligned.

Another typical example is depicted in Fig. 1.2. In this application, magnetic resonance images of a female breast are taken at different times (images from Bruce Daniel, Lucas Center for Magnetic Resonance Spectroscopy and Imaging, Stanford University). The first image shows an MRI section taken during the so-called wash-in phase of a radiopaque marker and the second image shows the analogous section during the so-called wash-out phase. A comparison of these two images indicates a suspicious region in the upper part of the images. This region can be detected easily since the images have been registered: tissue located at a certain position in the wash-in image is related to tissue in the same position




FIG. 1.1 Two different images of human hands; images from Amit (1994).



FIG. 1.2 Magnetic resonance images of a female breast, LEFT: during the wash-out phase, RIGHT: during the wash-in phase.

in the wash-out phase. Without registration, it would be much harder to decide whether an increase of brightness is related to contrast uptake or to movement of the patient, e.g., due to breathing and/or heart beat. In this application, *rigid registration* (i.e., a registration based on rotations and translations) does not provide a sufficient solution. A non-linear transformation is necessary to correct the local differences in the images.

Histological sectioning processes represent a further important source of registration problems. Here, two-dimensional slices of a three-dimensional object are produced. These two-dimensional views may serve as a basis for further image processing on a very fine scale using a microscope. But often one is interested not only in the microscopic data related to the two-dimensional views but in




FIG. 1.3 Two squares, LEFT: reference image, RIGHT: template image.

the relation to the three-dimensional object. Typically, the sectioning process introduces individual distortions of the two-dimensional views. Thus, a direct reconstruction of the microscopic data is not possible, in general, and there is a demand for high-resolution registration procedures. A detailed description of this registration problem is given in Chapter 2.

The image registration problem can be phrased in only a few words: given a *reference* and a *template* image, find a suitable *transformation* such that the transformed template becomes *similar* to the reference. However, though the problem is easy to state, it is hard to solve. The main reason is that the problem is *ill-posed*. Small changes of the input images can lead to completely different registration results. Moreover, the solution may not be unique. Figure 1.3 illustrates a simplified situation. Suppose we have to register the reference and template images depicted in Fig. 1.3, where, for simplicity reasons, we only allow for rigid transformations, i.e., rotations and translations. We immediately find several solutions: a pure translation, a rotation of 180 degrees, a rotation of 90 degrees followed by a translation, etc. These solutions are equivalent and without additional knowledge it is not possible to decide which one to use.

Another reason for making image registration such a challenging task is the fact that each application has its own demands with respect to the meaning of *similar* and *suitable*. For example, for the registration of X-ray images of bones, one may consider only rigid transformations, whereas for the registration of breast MRI scans one has to consider non-rigid transformations.

In the example of the hand (cf., Fig. 1.1), one could imagine using a deformable image model, based on the deformation properties of muscles, fat, bones, joints, and so forth. However, this model requires additional knowledge, e.g., the location of bones etc., which is in general not available in image registration. To make the problem even worse, image registration is often intended to supply this type of knowledge. A typical application of image

registration is to map a database connected to a reference onto a template image. For example, a very accurate segmentation of a reference image could be used as a basis for a template segmentation after registration.

Finally, some applications allow for time consuming computations while others demand real-time implementation. For example, shape changes of the brain (so-called *brain-shift*) during neurosurgery caused by the intervention need to be corrected numerically during the surgery, while the correction of a histological serial sectioning may run overnight.

Thus, one also has to compromise between complex modeling and accuracy on the one side and computing time and storage requirements on the other side. Therefore, it is not surprising that a huge variety of registration techniques has been proposed in the literature.

1.2 The intention

The intention of this book is to contribute to image registration from a mathematical point of view. A general approach to image registration is presented. This approach is based on a similarity measure and a regularizer. The similarity measure can be viewed as the driving force of the registration whereas the regularizer controls the transformation. This general framework enables the classification of well-known methods as well as the design of novel schemes with particular features. Moreover, it also allows for the design of very fast implementations.

The book is divided into two parts: parametric and non-parametric image registration. Though this partitioning is slightly misleading (because there are some parameters also in non-parametric registration), it provides a good starting point.

The first part summarizes and relates state of the art registration techniques with respect to the new, general approach. The schemes presented are restricted to an expansion of the desired transformation in terms of a finite number of basis functions. Thus, registration results in the computation of optimal parameters for the underlying expansion.

The second and main part demonstrates how non-parametric registration techniques can be designed on the basis of the new approach. It is shown how well-established registration schemes like *elastic* and *fluid registration* can be related. Moreover, it is also shown how to derive new schemes with particular features using this general approach. In addition, a general numerical treatment is presented.

Finally, and most importantly, we present fast, efficient, and stable implementations for the proposed registration techniques. In particular, for *diffusion registration*, an $\mathcal{O}(N)$ implementation is presented, where N denotes the number of pixels or voxels. Using this implementation, diffusion registration is at present to the best of our knowledge the fastest non-parametric registration technique.

1.3 An overview

In the first part of the book, we address the so-called *landmark-based* (cf., Chapter 4), *principal axes-based* (cf., Chapter 5), and *optimal linear registration* (cf., Chapter 6). In landmark-based registration, one locates a number of positions in the reference image and corresponding points in the template image. The transformation is obtained by minimizing the distance (typically the Euclidean or a Mahalanobis distance) between the landmarks in the reference and deformed template image. From a mathematical point of view it is irrelevant whether or not these points are anatomically meaningful (*anatomical landmarks*) and whether these points have been selected before (*fiducial markers*) or after the imaging. The main point is that image features are deduced and the computed transformation relates these features. The registration is based purely on these features further image knowledge is not taken into account. The transformation can be regularized in different ways: it may be taken from a parameterizable space (cf., Section 4.2), e.g., a polynomial or a spline space, or be characterized by an additional smoothness property (cf., Section 4.3). Feature-based registration on a linear space is a delicate matter, because of a possible singularity in the underlying Vandermonde system. Thin plate spline registration provides a remedy. Here, the transformation is not restricted to a linear space but regularized explicitly.

The technique presented in Chapter 5 is based on a similar idea. The difference now is that geometrical features, i.e., center of gravity and the principal axes of a displayed object, are used for registration. The main advantage of this approach is that these features can be deduced automatically. Thus, no user interaction is required.

In Chapter 6, image intensity-based distance measures are introduced. In particular, we consider the *sum of squared differences* (cf., Section 6.1), a *correlation-based* distance measure (cf., Section 6.3), and *mutual information* (cf., Section 6.4). The transformation is restricted to a parameterized space and the registration can be obtained by minimizing the image distance. For presentation reasons, we restrict the transformations to an affine linear space. However, this technique carries over to all sets of parameterizable transformations. Note that the minimization of a distance measure over a parameterized spline space is probably the most commonly used registration technique in today's image registration.

The second part of the book is devoted to non-parametric image registration. The transformation is no longer restricted to a parameterizable set. Instead, a regularizing term is used to circumvent ill-posedness and to privilege more likely solutions. It is desirable to provide different regularizers since different applications require a different treatment in general. Chapter 8 is intended as an introduction to a general approach which comprises the particular methods presented in the following chapters: namely, *elastic* (Chapter 9), *fluid* (Chapter 10), *diffusion* (Chapter 11), and *curvature* registration (Chapter 12).

The registration problem is phrased in a functional setting. The main difference with respect to the parametric case where one is looking for optimal parameters of an expansion of the transformation is that we are now simply seeking a smooth transformation. No parameters are involved in representing the transformation.

This formulation presents a unified approach to image registration based on the regularized minimization of a distance measure. As has already been pointed out, image registration is an ill-posed problem. Thus, regularization is essential and inevitable. The regularizer is needed and used to pick out the most likely solution. Moreover, regularization can be used to supply additional knowledge. Actually, it is the regularizer which distinguishes the methods presented.

In Chapter 9 we comment on elastic registration in detail. The deformation of an elastic body is explained using linear elasticity theory. Note that a linear elasticity model restricts the transformation to small spatial variations. Hence, for some applications, e.g., the registration of images obtained from different human brains, elastic deformations are too restrictive by far. Thus, we also comment on fluid registration (cf., Chapter 10), where the intention is to mimic a flow of fluid in a certain sense. This approach allows for large deformations.

Diffusion registration (cf., Chapter 11) is another registration technique which can easily be derived from the general frame presented in Chapter 8. Here, we borrow the well-known optical flow regularizer for image registration. Moreover, we relate this technique to another popular non-parametric registration technique: Thirion's so-called *demons registration*. The main advantages of this regularizer is that it allows for a fast implementation. Thus, this scheme is attractive for large-scale two-dimensional and three-dimensional registration problems.

Finally, we illustrate how to use this general framework to design regularizers with particular features. In Chapter 12 we discuss a novel non-parametric registration technique which relies on a curvature-based penalizing term. This approach not only provides smooth solutions but also allows for automatic rigid alignment. Thus, in contrast to other popular registration schemes, the pre-registration step becomes redundant.

A further important aspect of image registration is the acceptance of a method in a practical environment. For this reason, robust and efficient schemes together with high-quality results become key issues. For parametric registration standard minimization techniques, like, for example, Gauss-Newton or Levenberg-Marquardt techniques, are used commonly and, with some additional effort, can even be applied to the minimization of mutual information.

Unfortunately, the computation of a numerical solution for a non-parametric registration is not straightforward. Thus, we also present algorithms for all four non-parametric registration schemes. The corresponding Euler-Lagrange equations serve as a starting point for all implementations. After an appropriate discretization, this approach leads to high-dimensional linear systems of equations which have to be solved numerically. Though this dimension can get

quite large, e.g., $N = 52\,428\,800$ for the registration of the histological serial sectioning presented in Section 9.10, the underlying systems are highly structured. The structure enables the design of efficient solution schemes. The complexity of the devised schemes is $\mathcal{O}(N \log N)$ for elastic, fluid, and curvature registration and $\mathcal{O}(N)$ for diffusion registration, where N denotes the number of image pixels or voxels or whatever. These fast solution schemes form the backbone of our implementations.

Though the proposed techniques are very general and can be used for various applications, we focus on one particular application: the reconstruction of a histological serial sectioning of a human brain (cf., Chapter 2). However, additional examples are presented to illustrate the advantages and disadvantages of the various techniques. The hand images (cf., Fig. 1.1) serve as the example in this respect.

2

THE HUMAN NEUROSCANNING PROJECT

The Human NeuroScanning Project (HNSP) was initiated by O. Schmitt; cf., Schmitt (2001). The goal of this project is a three-dimensional reconstruction down to particular neurons of a whole human brain based on microscopic modalities. These data should then be used as the basic structure for the integration of functional data based on stochastic mapping and later on for modeling and simulation studies of a virtual brain; see Schmitt (2001) for details.

In order to locate the spatial position of the neurons, the post mortem brain from a 55 year old male human voluntary donor was prepared in several steps. The brain was fixed in a neutral buffered formaldehyde solution for three months. An MRI scan of the brain was produced after fixation. Dehydration and embedding of the brain in paraffin required three further months.

This preparatory work was followed by sectioning the brain in $20\text{ }\mu\text{m}$ thick slices (about 5000 for this brain) using a sliding microtome; cf., Fig. 2.1. Before each slicing step a high-resolution episcopic image (1352×1795 pixels, three colors) was taken; cf., Fig. 2.2.

Figure 2.3 displays a tissue slice after sectioning. The tissue slice was then stretched in warm water at 55°C . Thereafter, it was transferred onto a microscopic slide and dried overnight. After drying, the sections were stained in galloxyanin chrome alum and mounted under cover-glasses.

A specialized light microscope (LMAS) with an extraordinarily large object range of $250 \times 250\text{ mm}^2$ is used to visualize all cells of the large tissue sections; cf., Fig. 2.2. Different neuronal entities were analyzed on different structural scales, i.e., from macroscopic details down to the cellular level; see Schmitt (2001) for the image processing.

In order to relate the microscopic data to a macroscopic view of the slide and to recover the geometrical deformation of the tissue introduced by the various sectioning steps, transparent flat bed scans (FBS) of the slides were produced; cf., Fig. 2.3. Using a resolution of 2032 parts per inch (ppi) in an 8 Bit gray-scale mode the digitized images range between 5000×2000 and 11000×7000 pixels (about 196 MBytes storage for the largest scan).

The sectioning, stretching, and drying processes are essential to produce high-quality slides for further microscopic analysis; cf., Schmitt (2001). However, these methodological steps produce non-linear (i.e., not necessarily linear) deformations of the tissue. A deformed tissue section is visible in Fig. 2.3. These deformations prohibit a direct three-dimensional reconstruction of the data. Image registration techniques are required to by-pass these problems.




FIG. 2.1 The sectioning process. The episcopic camera is located at the top, the sliding microtome and the paraffin embedded brain are located at the bottom.

Figure 2.4 displays the arbitrarily chosen section 3799 (reference) and section 3800 (template) as well as the difference between the non-registered, the linearly registered, and the non-linearly registered sections. As is apparent from this figure, non-linear registration techniques may provide effective tools for the reconstruction of geometrical deformations.

In this phase of the HNSP the images were re-sampled to 1024×1024 pixels in order to limit the memory requirements and the computation times while providing adequate resolution. A straightforward re-sampling based on bi-linear interpolation was used.




FIG. 2.2 TOP: Episcopic image of the paraffin embedded brain before sectioning section 100; BOTTOM: slide under a particular light microscope.






FIG. 2.3 TOP: microtome with tissue section on top of the blade, BOTTOM LEFT: flat bed scanner, BOTTOM RIGHT: transparent flat bed scan of the tissue section after fixation.

Elastic matching (to be explained in Chapter 9) can be used also for multimodal matching of histological sections with non-deformed MRI scans or episcopic images. The episcopic images might be derived from image processing before sectioning the embedded brain. The uncompressed amount of flat bed scanned data was approximately 700 GBytes + 40 GBytes episcopic data for one human brain.




FIG. 2.4 Sections 3799 (TOP LEFT) and 3800 (TOP RIGHT) as well as the difference images non- (MIDDLE LEFT), linearly (MIDDLE RIGHT), and non-linearly registered.