

Helsinki University of Technology
Department of Engineering Physics and Mathematics
Low Temperature Laboratory, Brain Research Unit
Espoo 2007

QUALITY IMPROVEMENTS FOR MULTI-MODAL NEUROIMAGING

Mika Seppä



TEKNILLINEN KORKEAKOULU
TEKNISKA HÖGSKOLAN
HELSINKI UNIVERSITY OF TECHNOLOGY
TECHNISCHE UNIVERSITÄT HELSINKI
UNIVERSITE DE TECHNOLOGIE D'HELSINKI

Helsinki University of Technology
Department of Engineering Physics and Mathematics
Espoo 2007

QUALITY IMPROVEMENTS FOR MULTI-MODAL NEUROIMAGING

Mika Seppä

Dissertation for the degree of Doctor of Science in Technology to be presented with due permission of the Department of Engineering Physics and Mathematics, for public examination and debate in Auditorium F1 at Helsinki University of Technology (Espoo, Finland) on the 17th of September, 2007, at 12 o'clock noon.

Helsinki University of Technology
Low Temperature Laboratory
Brain Research Unit

Teknillinen korkeakoulu
Kylmälaboratorio
Aivotutkimusyksikkö

Distribution:

Low Temperature Laboratory
Helsinki University of Technology

P.O. Box 2200

FIN-02015 HUT, FINLAND

Tel. +358-9-451 2962

Fax +358-9-451 2969

This thesis is downloadable at <http://lib.hut.fi/Diss/>

Copyright 2007 Mika Seppä

ISBN 978-951-22-8927-1

ISBN 978-951-22-8928-8 (PDF)

Picaset Oy

Helsinki 2007



ABSTRACT OF DOCTORAL DISSERTATION		HELSINKI UNIVERSITY OF TECHNOLOGY P. O. BOX 1000, FI-02015 TKK http://www.tkk.fi	
Author Mika Seppä			
Name of the dissertation Quality Improvements for Multi-Modal Neuroimaging			
Manuscript submitted 12.06.2007		Manuscript revised 21.08.2007	
Date of the defence 17.09.2007			
<input type="checkbox"/> Monograph		<input checked="" type="checkbox"/> Article dissertation (summary + original articles)	
Department		Department of Engineering Physics and Mathematics	
Laboratory		Systems Analysis Laboratory	
Field of research		Applied Mathematics	
Opponent		Prof. Rainer Goebel	
Supervisor		Prof. Raimo P. Hämmäläinen	
Instructors		Prof. Matti Hämmäläinen & Prof. Riitta Hari	
Abstract			
<p>This work, concentrating on improving the quality of multi-modal medical image fusion, was carried out in the Brain Research Unit of the Low Temperature Laboratory at Helsinki University of Technology (TKK) in close collaboration with the Advanced Magnetic Imaging (AMI) Centre of TKK.</p> <p>Modern medical imaging devices produce large amounts of highly detailed information about the anatomy and function of various body parts. Different imaging modalities are typically sensitive to different properties of the underlying tissue and therefore produce complementary information. In multi-modal neuroimaging, data from different modalities is fused together to facilitate better analysis of the structure and activity of the brain or of other parts of the nervous system.</p> <p>Multi-modal image fusion in human neuroimaging has many uses both in clinical settings and in research. Since the different imaging modalities reveal different properties of the nervous system, joint visualization helps to combine all this information for interpretation. Typically, data are combined to visualize the anatomical locations of the functional activations. In addition to joint visualization, multi-modal neuroimaging can also incorporate information from one modality to the analysis of data in another modality and thereby lead to more accurate results.</p> <p>In the core of multi-modal imaging lies image registration that brings together the information from two or more imaging modalities. Since the spatial alignments and resolutions of the registered images typically differ, resampling is required to bring the data into a common coordinate frame for visualization or for further analysis. This thesis work concentrates on the three key stages of the multi-modal image fusion: registration, resampling, and visualization. The introduced enhancements for mutual-information registration allow for sub-sample accuracy even in the worst-case scenarios and the novel two-stage resampling algorithm produces smaller resampling errors than any of the currently used methods. Furthermore, the proposed enhancement to combine triangle meshes with volume rendering techniques provides fast high-quality visualization. In addition to these technical improvements, an application of diffusion tensor imaging to delineate the course of peripheral nerves is presented.</p>			
Keywords medical imaging, image registration, resampling, visualization			
ISBN (printed) 978-951-22-8927-1		ISSN (printed)	
ISBN (pdf) 978-951-22-8928-8		ISSN (pdf)	
Language English		Number of pages 66 p. + app. 59 p.	
Publisher Low Temperature Laboratory, Helsinki University of Technology			
Print distribution Low Temperature Laboratory, Helsinki University of Technology			
<input checked="" type="checkbox"/> The dissertation can be read at http://lib.tkk.fi/Diss/2007/isbn9789512289288/			



VÄITÖSKIRJAN TIIVISTELMÄ		TEKNILLINEN KORKEAKOULU PL 1000, 02015 TKK http://www.tkk.fi	
Tekijä Mika Seppä			
Väitöskirjan nimi Laadullisia parannuksia multimodaaliseen neurokuvantamiseen			
Käsikirjoituksen päivämäärä 12.06.2007		Korjatun käsikirjoituksen päivämäärä 21.08.2007	
Väitöstilaisuuden ajankohta 17.09.2007			
<input type="checkbox"/> Monografia		<input checked="" type="checkbox"/> Yhdistelmäväitöskirja (yhteenveto + erillisartikkelit)	
Osasto		Teknillisen fysiikan ja matematiikan osasto	
Laboratorio		Systeemianalyysin laboratorio	
Tutkimusala		Sovellettu matematiikka	
Vastaväittäjä		prof. Rainer Goebel	
Työn valvoja		prof. Raimo P. Hämäläinen	
Työn ohjaajat		prof. Matti Hämäläinen & prof. Riitta Hari	
Tiivistelmä			
<p>Tämä väitöskirjatyö, jossa esitetään uusia tapoja parantaa eri kuvantamismenetelmien tuottamien kuvien yhdistämistä, on tehty Teknillisen korkeakoulun Kylmälaboratorion aivotutkimusyksikössä läheisessä yhteistyössä TKK:n AMI-keskuksen (Advanced Magnetic Imaging Centre) kanssa.</p> <p>Nykyaikaiset lääketieteelliset kuvantamismenetelmät tuottavat suuret määrät tarkkaa tietoa kehon eri osien anatomiasta sekä toiminnasta ja eri kuvantamismenetelmillä saatu tieto täydentää yleensä toisiaan. Multimodaalisessa neurokuvantamisessa eri menetelmillä saadun tiedon yhdistäminen helpottaa aivojen tai muun hermoston rakenteen ja toiminnan tutkimista. Kuvantamismenetelmien yhdistämisestä on paljon hyötyä sekä kliinisissä sovelluksissa että perustutkimuksessa. Tyypillisessä käytössä aivojen aktivaatiota ja rakennetta erikseen kuvaavien menetelmien tulokset visualisoidaan niin, että mitattu aivotoiminta on helppo paikallistaa suhteessa aivojen anatomiaan. Eri menetelmien tuottaman tiedon yhdistäminen mahdollistaa usein myös tarkemmat tulokset, kun yhden kuvausmenetelmän tietoa voidaan käyttää toisen menetelmän tulosten mallinnuksessa ja analyysissä.</p> <p>Eri menetelmien tuottamien kuvien tarkalla kohdennuksella on keskeinen rooli kuvatietojen yhdistämisessä. Koska kuvien tarkkuus sekä kolmiulotteinen asemointi poikkeavat yleensä toisistaan, kuvat joudutaan uudelleennäytteistämään ennenkuin ne voidaan tuoda yhteiseen koordinaatistoon visualisointia tai analyysia varten.</p> <p>Tämä väitöskirjatyö keskittyy multimodaalisen kuvantamisen kolmeen keskeiseen vaiheeseen: kohdennukseen, uudelleennäytteistykseen sekä visualisointiin. Työssä esitetyt parannukset kuvien kohdennukseen mahdollistavat alle näytekoon kohdennustarkkuuden huonoimmassakin tapauksessa ja uudelleennäytteistykseen kehitetty uusi kaksivaiheinen menetelmä tuottaa pienempiä virheitä kuin mikään aiempi menetelmä. Visualisoinnissa työ yhdistää kolmioverkkoihin perustuvan pintaesityksen tilavuuskuvien suoraan visualisointiin käytettyihin säteenheittomenetelmiin ja mahdollistaa erittäin nopean korkeatasoisen visualisoinnin. Teknisten parannusten lisäksi esitetään sovellus, jossa diffuusiotensorikuvauksen avulla erotetaan ääreishermit niitä ympäröivästä kudoksesta.</p>			
Asiasanat lääketieteellinen kuvantaminen, kuvien kohdennus, uudelleennäytteistys, visualisointi			
ISBN (painettu) 978-951-22-8927-1		ISSN (painettu)	
ISBN (pdf) 978-951-22-8928-8		ISSN (pdf)	
Kieli englanti		Sivumäärä 66 s. + liit. 59 s.	
Julkaisija Kylmälaboratorio, Teknillinen korkeakoulu			
Painetun väitöskirjan jakelu Kylmälaboratorio, Teknillinen korkeakoulu			
<input checked="" type="checkbox"/> Luettavissa verkossa osoitteessa http://lib.tkk.fi/Diss/2007/isbn9789512289288/			

Academic dissertation

Quality Improvements for Multi-Modal Neuroimaging

Author: Mika Seppä
Brain Research Unit, Low Temperature Laboratory
Helsinki University of Technology
Finland

Supervising professor: Prof. Raimo P. Hämmäläinen
Systems Analysis Laboratory
Helsinki University of Technology
Finland

Supervisors: Prof. Matti Hämmäläinen
MGH/MIT/HMS Athinoula A. Martinos Center for
Biomedical Imaging
Charlestown, USA

Prof. Riitta Hari
Brain Research Unit, Low Temperature Laboratory
Helsinki University of Technology
Finland

Preliminary examiners: Prof. Karen Egiazarian
Institute of Signal Processing
Tampere University of Technology
Finland

Doc. Jyrki Lötjönen
VTT Technical Research Centre of Finland
Tampere, Finland

Official opponent: Prof. Rainer Goebel
Department of Neurocognition
Maastricht University
Netherlands

Table of contents

List of publications	xiii
Abbreviations	xiv
Preface and acknowledgments	xv
1 Introduction	1
1.1 Contributions of the Thesis	2
1.2 Outline of the Thesis	3
2 Medical imaging modalities	5
2.1 Magnetic resonance imaging	5
2.1.1 Nuclear magnetic resonance	5
2.1.2 Signal formation	7
2.1.3 Image formation	8
2.2 Diffusion tensor imaging	8
2.2.1 Diffusion	9
2.2.2 Diffusion weighted images	9
2.2.3 Diffusion tensors	10
2.3 Functional magnetic resonance imaging	10
2.4 Magnetoencephalography and electroencephalography	12
2.5 Computed tomography	13
3 Image registration	15
3.1 Transformations	15
3.1.1 Linear transformations	15
3.1.2 Non-linear transformations	16
3.2 Match metrics	17
3.2.1 Landmark and structure-based metrics	17
3.2.2 Metrics based on image intensity	18
3.3 Optimizers	19
3.3.1 Algorithms	19
3.3.2 Multi-resolution techniques	19
3.4 Mutual-information registration	20
3.4.1 Entropy	20
3.4.2 Mutual information	21

3.4.3	Mutual-information in image registration	21
3.4.4	Joint-probability estimation	22
3.4.5	Artifacts	23
3.4.6	Improvements in accuracy (P1)	23
4	Resampling	25
4.1	Sampling theory	25
4.2	Convolution interpolation	26
4.3	Generalized interpolation	27
4.4	Specialized resampling methods	28
4.4.1	Fourier-space methods	28
4.4.2	Two-stage resampling (P2)	30
5	Visualization	35
5.1	Scalar volumes	35
5.1.1	Cross-sections	35
5.1.2	Surface-based rendering	36
5.1.3	Volume rendering	37
5.1.4	Texture-mapped triangle meshes (P3)	39
5.2	Tensor volumes	43
5.2.1	Direct tensor visualization	43
5.2.2	Tractography	45
5.2.3	Tractography of distal peripheral nerves (P4)	47
6	Discussion and Conclusions	49
	Bibliography	53

List of publications

The dissertation consists of the present summary article and the following publications:

- P1 Seppä, M. (2007). On continuous sampling and sub-sample accuracy in mutual-information registration. *TKK Report*, TKK-KYL-017.
- P2 Seppä, M. (2007). High-quality two-stage resampling for 3-D volumes in medical imaging. *Medical Image Analysis*, 11:346–360.
- P3 Seppä, M. and Hämäläinen, M. (2005). Visualizing human brain surface from T_1 -weighted MR images using texture-mapped triangle meshes. *NeuroImage*, 26:1–12.
- P4 Hiltunen, J., Suortti, T., Arvela, S., Seppä, M., Joensuu, R., and Hari, R. (2005). Diffusion tensor imaging and tractography of distal peripheral nerves at 3 T. *Clinical Neurophysiology*, 116:2315–2323.

Contributions of the author

I was the principal author in publications P1–P3 and I developed, implemented, and tested the new methods presented. I wrote these three manuscripts with active participation of the second author in publication P3.

In publication P4, I implemented new software to compute diffusion tensors and tractography from diffusion-weighted MR images. This software employed the visualization technique presented in publication P3. The main data analysis in publication P4 was performed with free computer programs available on-line. My newly implemented software employed tensor interpolation and a different tractography algorithm, and thereby allowed the verification of the results. I also actively participated in the writing of the manuscript.

Abbreviations

ADC	Apparent diffusion coefficient
BEM	Boundary-element model
BOLD	Blood oxygenation level dependent
BSP	Binary space partitioning
CPU	Central processing unit
CT	Computed tomography
DEC	Directionally-encoded color (map)
DFT	Discrete Fourier transform
DOF	Degrees of freedom
DTI	Diffusion tensor imaging
DVR	Direct volume rendering
DWI	Diffusion-weighted image
EEG	Electroencephalography
FA	Fractional anisotropy
fMRI	Functional magnetic resonance imaging
GI	Generalized interpolation
GPU	Graphics processing unit
ICP	Iterative closest-point (algorithm)
IVR	Indirect volume rendering
LI	Linear interpolation
MEG	Magnetoencephalography
MI	Mutual information
MOMS	Maximal-order-minimal-support
MR	Magnetic resonance
MRI	Magnetic resonance image/imaging
NMR	Nuclear magnetic resonance
OMOMS	Optimal-MOMS
PD	Proton density
PVI	Partial volume interpolation
RA	Relative anisotropy
RF	Radio frequency
RIU	Ratio-image uniformity
SDFT	Shifted-DFT (algorithm)
SNR	Signal-to-noise ratio
SQUID	Superconducting quantum interference device
ROI	Region of interest
TE	Echo time
TR	Repetition time
VIR	Variance of intensity ratios
VR	Volume ratio

Preface and acknowledgments

This thesis work was carried out in the Brain Research Unit of the Low Temperature Laboratory at the Helsinki University of Technology (TKK). The medical imaging data were collected at the Advanced Magnetic Imaging Centre of TKK and the Department of Radiology at the Helsinki University Central Hospital. This work was financially supported by the Academy of Finland, the Foundation of Vilho, Yrjö, and Kalle Väisälä (Finnish Academy of Sciences and Letters), the Helsinki University of Technology, and the Hospital District of Helsinki and Uusimaa (HUSLAB). I am also thankful for the support of the Finnish Graduate School of Neuroscience (FGSN).

It has been a great privilege and joy to work in the Low Temperature Laboratory which provides excellent facilities and good spirit for its many researchers to stay at the cutting edge of science. These circumstances are largely due to the efforts of the laboratory's first director, late Academician Olli V. Lounasmaa, and the current director Prof. Mikko Paalanen.

I am equally grateful to the primus motor of the Brain Research Unit, Prof. Riitta Hari, for creating the top-class possibilities for our interdisciplinary group to work on the forefront of the brain research. Her aptitude to master all the sides from administration to science and from neurophysiology to technical details has never ceased to amaze me.

I thank my supervisors, Prof. Riitta Hari and Prof. Matti Hämäläinen, for all the guidance, support, and positive attitude I have received during these years, both related and unrelated to this thesis. You gave me the inspiration and the possibility to work on this field that I have so much enjoyed.

I am grateful to my supervising professor, Prof. Raimo P. Hämäläinen from the Systems Analysis Laboratory of TKK, for inspiring lectures about applied mathematics and for his patience when I, after years of silence, suddenly appeared with tight self-imposed deadlines for the thesis.

My sincere thanks go to Prof. Karen Egiazarian and Doc. Jyrki Lötjönen for their valuable time and effort to prereview my thesis. I also want to thank my FGSN follow-up group, Doc. Jyrki Lötjönen and Prof. Ulla Ruotsalainen, for their guidance and for the opportunity to reflect my work with experts outside this university.

It has been a pleasure to work with Phil.Lic. Jaana Hiltunen, who I could always trust in matters concerning diffusion tensor imaging and magnetic resonance imaging in general. I am also thankful to my other co-authors: Mr. Sakari Arvela, Dr. Raimo Joensuu, and Ms. Taru Suortti. In addition, I want to express my deepest gratitude to Mr. Lauri Parkkonen and Dr. Kimmo Uutela for their valuable comments on my manuscripts.

To finalize this thesis, I thank Ms. Linda Henriksson, Phil.Lic. Jaana Hiltunen, and Mr. Lauri Parkkonen for kindly providing me with pictures and I am in debt to Dr. Cathy Nangini for language corrections despite the tight schedule. To be able to focus on the thesis, I am especially grateful to Mr. Jan Kujala whose invaluable efforts on daily basis have made it possible.

I want to thank for and acknowledge the efforts of all the senior scientists of our

group, past and present, for shaping the Brain Research Unit to its success. I have especially enjoyed my talks with Prof. Riitta Salmelin and Doc. Simo Vanni.

During my years in the Low Temperature Laboratory, I have had the privilege to share the company of many colleagues and have great time both at work and at free time with them. Although they are too numerous to list here and thereby take the risk of accidentally overlooking somebody, I am truly grateful for their role in my work and as friends.

I want to thank also the secretaries and other support personnel for keeping the laboratory running and helping with all possible practical matters, including those related to the finalization of this thesis. Special thanks are due to all the members, past and present, of the computer support team for keeping the computer systems running smoothly. In addition, I feel especially privileged for the fruitful collaboration with Elekta Neuromag Oy and its multi-talented personnel.

Outside the working environment, I wish to thank my many friends like Esa, Kalle, Kim, Kimmo, Pete, and Ville with whom I have shared enjoyable time doing sports, playing games, or just relaxing otherwise.

I am deeply grateful to my parents for their continuous support and unconditional love at all times. To my dear wife Isabel, I owe everything I deem precious in life. She and our lovely son Luis are the real sunshines of my days.

Espoo, August 2007

Mika Seppä

Chapter 1

Introduction

Modern medical imaging devices produce large amounts of highly detailed information about the structure and function of various body parts. Different medical imaging modalities are sensitive to different properties of the underlying tissue and therefore produce complementary information. In multi-modal neuroimaging, data from different modalities are fused together to facilitate better analysis of the structure and function of the brain or of other parts of the nervous system.

Medical imaging modalities can be divided into anatomical and functional, according to the data they produce. Anatomical imaging provides information about the structure and other time-invariant properties of the tissue. On the other hand, functional imaging produces information about the dynamic behavior of the imaged body part, typically in a time scale ranging from a few milliseconds to a few seconds. The functional imaging of the human brain reflects neural activations, either directly or indirectly.

Multi-modal neuroimaging has many applications in both clinical and research studies. Joint visualization of data from different imaging modalities is needed to combine all this information for interpretation. For example, functional and structural information is typically combined to visualize the anatomical locations of the functional data. In addition to joint visualization, multi-modal neuroimaging can bring information available from one modality to the analysis of data in another modality. For example, boundary element models (BEMs), created from anatomical images, allow more accurate forward modeling of magnetoencephalographic and electroencephalographic signals and thus lead to more precise estimates of the underlying sources.

At the core of multi-modal neuroimaging lies the data registration (also known as co-registration or spatial normalization) that brings together the information from two or more different data-sets. In case of 2-D or 3-D data, the specific terms “image registration” or “volume registration” are often used. Since the spatial alignments and resolutions of the registered data-sets typically differ, resampling is required to bring the data into a common coordinate frame for visualization or for further analysis.

Despite the diversity of applications in multi-modal imaging, the following three key stages are practically always present: registration, resampling, and visualization. Even in the absence of an explicit resampling stage, the data are typically implicitly resampled during visualization or already in the registration stage.

1.1 Contributions of the Thesis

This work aims at improving the quality of multi-modal neuroimaging. It presents improvements to the three key aspects of image fusion: registration (P1), resampling (P2), and visualization (P3). The last publication (P4) presents an application of multi-modal imaging.

In the analysis of multi-modal data, high-quality image processing algorithms are essential to guarantee that the correlations found are actually real and not produced by the registration and resampling errors or by other artifacts of the methods. In particular, the otherwise very versatile mutual-information (MI) registration metrics are known to contain in certain situations sampling-related artifacts that reduce the registration accuracy. These artifacts are currently one of the key concerns and limit the applicability of these metrics in *e.g.* motion-correction applications.

For the joint visualization of the analyzed multi-modal data, the tendency is to combine more and more information for simultaneous viewing. To allow interactive manipulation of such complex scenes, it is crucial to render each geometrical model or other displayed data in minimal time while maintaining high visual quality.

Publication P1 of this Thesis concentrates on the sub-sample accuracy of the MI registration. A new concept of continuous sampling is developed and it leads to a higher-order sampling method that reduces the sampling-related artifacts and improves the registration accuracy. The simulation experiments reveal and measure the registration errors inherent in the MI metric itself with sub-sample accuracy. The results indicate a worst-case average error of 0.2–0.3 samples for the currently used methods and of only 0.04 samples for the suggested higher-order method.

Publication P2 introduces a novel two-stage resampling method that combines Fourier-space up-sampling with traditional interpolation. Efficient implementation details are discussed and a new algorithm for Fourier-space up-sampling is presented. The experiments show that the proposed two-stage resampling method produces 2–10 dB better signal-to-noise ratio (*i.e.* signal to interpolation error ratio) than the previously used methods.

Publication P3 introduces a novel technique for visualizing T_1 -weighted magnetic resonance images (MRIs) of the human brain. A previously used method in volume rendering is combined with texture-mapped triangle meshes. The work concentrates on minimizing errors when the 3-D geometry is opened onto a flat 2-D texture image. A novel idea of measuring the distortion with respect to the data resolution rather than to the geometry is introduced. The proposed visualization technique allows the use of larger triangles, thereby providing high-quality visualization with excellent speed.

Publication P4 presents the first diffusion tensor imaging and tractography results of human distal peripheral nerves. The results show that the maximum fractional anisotropy (FA) of the diffusion at the nerves covaries with the nerve size. The nerves were consequently delineated by tractography using two pieces of software: one freely available on-line and other implemented by our laboratory. Our software employed tensor interpolation and different tractography algorithm, and thereby allowed the reliability of the results to be assessed. The visualization combined the tracts with anatomy and employed the method of publication P3.

This Thesis introduces enhancements to the key methods used in medical image processing. Therefore, these novel algorithms are not limited to any specific use but can be employed in a large variety of applications ranging from the intra-subject registration to inter-subject registration, group averages, and brain atlases.

1.2 Outline of the Thesis

This summary provides background information for multi-modal neuroimaging in a concise form, briefly presenting the current state-of-the-art of the developed methods. First, Chapter 2 provides an overview of the medical imaging modalities relevant to this work. Then, the three key aspects of the multi-modal imaging are described in detail: the registration (Chapter 3), resampling (Chapter 4), and visualization (Chapter 5). Finally, the results are discussed and the conclusions are drawn in Chapter 6.

For simplicity, medical data-sets in this thesis are called “images” regardless of whether they are 2-D images, 3-D volumes, or something else. The samples in these images are called pixels. The specific terms volume, voxel, etc. are used when it is necessary to differentiate from 2-D images.

Chapter 2

Medical imaging modalities

This Chapter describes the anatomical and functional imaging modalities directly relevant to this Thesis. Here, magnetic resonance imaging (MRI), diffusion tensor imaging (DTI), and functional magnetic resonance imaging (fMRI) are considered as separate modalities, although they employ the same imaging device, *i.e.* an MRI scanner. However, MRI, DTI, and fMRI data reflect very different properties of tissues. Especially for image registration (Chapter 3), images from these methods and even MR images with different weightings (see Fig. 2.1) are considered to originate from different modalities.

MRI and DTI are central to this Thesis and are therefore explored in greater detail whereas fMRI, magnetoencephalography (MEG), electroencephalography (EEG), and computed tomography (CT) are described more briefly. Nowadays, fMRI, MEG, and EEG are the main noninvasive functional imaging modalities used in brain research. The Brain Research Unit at the Low Temperature Laboratory, where this work has been conducted, is known as one of the pioneers in MEG development. The results of functional imaging are typically combined with the anatomical information provided by MRI, although CT provides better resolution with lower geometrical distortions. However, the soft-tissue contrast is poor in CT and the inherent radiation dose is acceptable only for clinical studies.

2.1 Magnetic resonance imaging

Magnetic resonance imaging (Lauterbur, 1973; Mansfield, 1977) employs nuclear magnetic resonance (NMR) (Bloch, 1946; Purcell et al., 1946) for non-invasive 3-D imaging. The NMR phenomenon is based on the quantum-mechanical properties of the protons and neutrons in an atom's nucleus. All nuclei containing odd numbers of protons or neutrons have a net spin angular momentum and a related net magnetic moment. In medical imaging, hydrogen ^1H (single proton) is the most commonly measured nucleus due to its natural abundance in the tissues. In the following discussion of NMR, we assume that the target nucleus is that of ^1H .

2.1.1 Nuclear magnetic resonance

The spin of the proton can take either a value of $1/2$ or $-1/2$. When a proton is placed in a magnetic field with field strength B_0 , the spin-related magnetic moment is either parallel (low-energy spin state) or anti-parallel (high-energy spin state) to the field and precesses about it. The angular frequency of this precession is $\omega_0 = \gamma B_0$ with the respective Larmor frequency $f = \omega_0/2\pi = \gamma B_0/2\pi$ (Freeman, 1998). Here, γ is

the nucleus-specific gyromagnetic ratio; for hydrogen protons $\gamma/2\pi = 42.58$ MHz/T (Huettel et al., 2004).

The difference between the proton's two energy levels is $\Delta E = \gamma B_0 h/2\pi$, where h is Planck's constant. At a thermal equilibrium state, the lower energy level is slightly more populated (about 20 parts per million at 37° C and 3 T field) according to the Boltzmann equation $N_{\text{high}}/N_{\text{low}} = \exp(-\Delta E/kT)$ (Freeman, 1998). Here, k is the Boltzmann constant, T is the absolute temperature, and N_{high} and N_{low} are the number of protons in the high and low energy states, respectively.

In a group of protons, this population difference in the energy levels produces a net magnetization M_0 . As the protons precess at random phases, the net magnetization vector is parallel to the magnetic field B_0 . The MRI manipulates this net magnetization vector by applying radio-frequency (RF) electromagnetic pulses at the Larmor frequency. At this specific frequency, protons in the lower energy state can absorb a photon with energy ΔE and enter the higher energy state. Soon after the excitation, the protons re-emit their excess energy at the Larmor frequency and return to the lower state.

The net magnetization vector can be rotated into the plane orthogonal to the magnetic field B_0 by using an RF excitation pulse with correct length and orthogonal to B_0 . Such a pulse is called the 90° pulse as it flips the net magnetization vector by 90°. The magnetic moments of the individual protons are now in phase and give rise to the precessing transverse net magnetization that produces the measured NMR signal. Over time, the synchrony is gradually lost and the transverse magnetization disappears. This phenomenon, known as the transverse relaxation, is caused by two effects.

Spin–spin interactions, the disturbances the spins cause to one another, are the intrinsic cause for the transverse relaxation. The signal loss by this mechanism is called T_2 decay and is characterized by the spin–spin relaxation time T_2 . The transverse magnetization M_{xy} disappears according to

$$M_{xy}(t) = M_0 e^{-t/T_2} \quad (2.1)$$

where t is the time after the 90° pulse (Huettel et al., 2004). An extrinsic cause for the transverse relaxation is the magnetic field inhomogeneity experienced by the protons. Variation in the local field strength affects the precession frequency and leads to the loss of coherence. The signal loss due to the combined effects of spin–spin interactions and the field inhomogeneities is known as T_2^* decay and is characterized by the time constant T_2^* . The T_2^* decay is faster than the T_2 decay ($T_2^* < T_2$) due to the added effect of field inhomogeneity.

The local magnetic-field variations cause some of the protons to precess faster and some slower. The slow protons will lag behind, whereas the fast protons will be ahead. In a spin-echo sequence, a 180° refocusing RF pulse inverts the phase. Now, the faster protons are behind in phase and will soon catch the slower ones again. As the phase coherence returns for a moment, the transverse magnetization re-emerges and causes an echo in the measured RF signal. The echo signal is employed for the MRI and in addition to spin-echo, other common pulse sequences employ gradient-echo imaging where the echo is produced by magnetic field gradients.

Immediately following the original 90° pulse, the longitudinal net magnetization is zero and the two energy levels are equally populated. The time constant that describes how fast the longitudinal component returns back to the thermal equilibrium is called the spin–lattice or longitudinal relaxation time (T_1). The recovery of the longitudinal

Table 2.1: Approximate values for the relaxation times at 3 T field strength (Wansapura et al., 1999).

	T_1 (ms)	T_2 (ms)	T_2^* (ms)
Gray matter	1300	110	45
White matter	800	80	45




Figure 2.1: T_1 -weighted (T_1), T_2 -weighted (T_2), and proton-density weighted (PD) MR images. Imaged at 3 T at the AMI Centre of TKK.

component is given by equation (Huettel et al., 2004)

$$M_z(t) = M_0(1 - e^{-t/T_1}) \quad (2.2)$$

and is always slower than the loss of the transverse component: $T_1 > T_2 > T_2^*$.

2.1.2 Signal formation

The NMR phenomenon is applicable for medical imaging because each tissue type has its specific relaxation time (see Table 2.1 for examples). By varying the RF pulse sequences, the measured signal can be made sensitive to T_1 , T_2 , or T_2^* effects. The timing of the pulse sequences is governed by two important parameters.

The first parameter, repetition time (TR), is the time interval between successive excitation pulses. If TR is small relative to the relaxation time T_1 , only a part of the original longitudinal magnetization M_0 has recovered (Eq. 2.2) before the new excitation. Therefore, the transverse magnetization component M_{xy} after the second excitation is reduced to

$$M_{xy}(t) = M_0(1 - e^{-TR/T_1})e^{-t/T_2}. \quad (2.3)$$

The second parameter, echo time (TE), is the time interval between the excitation and the data acquisition occurring at the signal echo. The transverse magnetization is

$$M_{xy}(TE) = M_0(1 - e^{-TR/T_1})e^{-TE/T_2} \quad (2.4)$$

at the time of the acquisition and relates directly to the measured MR signal.

By controlling the parameters TR and TE, it is possible to maximize tissue contrast, *i.e.* the difference in signal strength between the tissue types. Fig. 2.1 shows examples of T_1 -weighted (intermediate TR, short TE), T_2 -weighted (long TR, intermediate TE), and proton-density (long TR, short TE) MRIs. Proton-density (PD) imaging minimizes the effects of the exponential terms in 2.4 and measures the signal strength related to M_0 directly, *i.e.* the density of the protons producing the signal.

2.1.3 Image formation

Magnetic-field gradients are central for MR image formation (Lauterbur, 1973; Mansfield, 1977). These gradients cause spatial variations in the magnetic field and thus the magnetic-field dependent Larmor frequency is different at different points of the tissue. When such a field gradient is applied during the narrow-band RF excitation pulse, only protons in a thin slice of the imaged volume are capable of absorbing the energy. This procedure is called slice selection as the subsequently measured MR signal will only come from this single excited slice.

Two other gradients, orthogonal to each other and to the slice selection gradient, are employed for spatial encoding of the received signal. These gradients correspond to the row and column directions of the image on the excited slice. One of these gradients is applied during the signal acquisition and it varies the field strength along the image row direction. Therefore, the protons at each column position emit their energy at a different Larmor frequency.

Another magnetic field gradient is applied for a short time just prior to signal acquisition. This gradient is along the column direction so that the Larmor frequency for each image row is different. Therefore, the protons in a certain image row precess faster than protons in another row. After this gradient is turned off, all protons return to the original precession frequency but maintain their phase difference acquired during the presence of the field gradient.

These two gradients create a frequency–phase encoding scheme, called k -space, for the acquired signal. The k -space representation of the image is typically collected one row at a time, *i.e.* the frequency-encoded signal is recorded for a certain phase encoding and then repeated with another phase encoding. The final MR image can be reconstructed from this collected k -space image by a 2-D Fourier transform.

The explanation above describes the basics of the MR imaging and many different variations of the MRI pulse sequences exist. Sequences can employ different flip angles other than 90° , use multiple echoes to collect more than one row of the k -space image for each excitation, sample the k -space along special trajectories, and even in an in-homogeneous fashion. A 3-D volume is typically imaged one slice at a time using the technique above. However, 3-D volumes can also be collected using 3-D k -space pulse sequences, where the slice selection gradient is discarded and a second phase-encoding gradient is introduced. The primary advantage is the higher signal-to-noise ratio (SNR) as the volume contains more excited spins contributing to the MR signal. Such pulse sequences are used especially for high-resolution anatomical images (Huettel et al., 2004).

2.2 Diffusion tensor imaging

Diffusion tensor imaging (DTI) measures the random movements of water molecules (Basser et al., 1994b; Le Bihan, 1995) and can reveal organized micro-structure that limits these movements in certain directions. DTI has been mainly applied in the study of fiber tracts and white matter structures of the human brain (Pierpaoli et al., 1996), although other applications have also emerged. The basics of the DTI are explained below, largely based on a recent review by Le Bihan (2003).

2.2.1 Diffusion

Molecular diffusion is caused by the random thermal motion of molecules. During a given time in a free medium, this random walk (Brownian motion) leads to molecular displacements that obey a 3-D Gaussian distribution. These displacements are statistically described by a diffusion coefficient D which depends only on the size (mass) of the molecules, the temperature, and the nature (viscosity) of the medium (Le Bihan, 2003).

The diffusion along one dimension is characterized by the equation $E\{\Delta x^2\} = 2DT_d$, where T_d is the diffusion time and $E\{\Delta x^2\}$ is the variance of the molecular displacements Δx . For example, the diffusion coefficient for free water molecules in 37° C water is approx. $3 \times 10^{-9} \text{ m}^2\text{s}^{-1}$ giving a diffusion distance (standard deviation) of 17 μm during 50 ms (Le Bihan, 2003).

During their diffusion in a non-homogeneous medium, the molecules interact with obstacles present and thereby probe the medium on a microscopic scale. In tissues, cell membranes, fibers, and macromolecules impede the water-molecule movements and thereby reduce their diffusion distances. Depending on the local microscopic structures, the diffusion might be limited more in one direction than in another, thereby causing anisotropy in the displacement distribution.

Measures of the average diffusion distances and of the diffusion anisotropy provide invaluable clues to the organization of the fine-structure features in the tissue. As the water diffusion distances are in the scale of 10 μm during typically used diffusion times, the water molecules probe the surroundings well below the usual millimetric image resolution.

2.2.2 Diffusion weighted images

Diffusion-weighted images (DWIs) (Merboldt et al., 1985; Taylor and Bushell, 1985; Le Bihan, 1995) are acquired by special MRI pulse sequences that sensitize the NMR signal to spin diffusion (Stejskal and Tanner, 1965). Two strong magnetic field gradient pulses control the diffusion weighting. The first pulse induces a phase shift along the applied gradient direction whereas the second pulse induces an equal but opposite phase shift. Therefore, the phase shift is canceled for all spins that have not moved in the gradient direction. On the other hand, those spins that have diffused along the gradient direction will experience different phase shifts and will not be completely re-focused at the signal echo. Thus, diffusion along the applied gradient direction causes signal loss.

To measure this signal loss and remove the signal dependency to PD, T_1 , and T_2 effects, at least two measurements for the diffusion are required. Typically, one image is acquired without diffusion weighting and another with the weighting applied. Then, the apparent diffusion coefficient (ADC) D can be calculated from equation (Stejskal and Tanner, 1965)

$$S = S_0 e^{-bD} \quad (2.5)$$

where S_0 is the signal without diffusion weighting, S is with the weighting, and b is the diffusion weighting factor (Le Bihan et al., 1986). This factor is

$$b = \gamma^2 \delta^2 \left(\Delta - \frac{\delta}{3} \right) |\mathbf{g}|^2 \quad (2.6)$$

where δ is the duration of the diffusion gradient pulses, Δ is the time between these pulses, and $|\mathbf{g}|$ is the strength of the diffusion gradients. The measured ADC value D depends on the diffusion weighting factor b and on the direction of the sensitizing gradient \mathbf{g} . Top row of Fig. 2.2 shows examples of diffusion weighted images.

2.2.3 Diffusion tensors

The strength of the diffusion can be different in different spatial directions, and the formalism of the diffusion tensor (Basser et al., 1994b,a) was developed to characterize this anisotropy. Using the diffusion tensor \mathbf{D} (second rank tensor), described as a symmetric 3×3 matrix

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{bmatrix}, \quad (2.7)$$

Eq. 2.5 above can be expressed as

$$S = S_0 e^{-b\hat{\mathbf{g}}^T \mathbf{D} \hat{\mathbf{g}}} \quad (2.8)$$

where $\hat{\mathbf{g}} = \mathbf{g}/|\mathbf{g}|$ is the normalized gradient vector. The symmetric tensor \mathbf{D} has six independent coefficients, and thus at least six different gradient directions in addition to the baseline data S_0 are required to estimate it. More directions can be used to suppress noise and to provide more robust tensor estimate.

For a total of n applied gradient directions, the following set of linear equations can be used for estimating the tensor coefficients (Masutani et al., 2003, with corrections).

$$\begin{bmatrix} x_1^2 & y_1^2 & z_1^2 & 2x_1y_1 & 2y_1z_1 & 2z_1x_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_i^2 & y_i^2 & z_i^2 & 2x_iy_i & 2y_iz_i & 2z_ix_i \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_n^2 & y_n^2 & z_n^2 & 2x_ny_n & 2y_nz_n & 2z_nx_n \end{bmatrix} \begin{bmatrix} D_{xx} \\ D_{yy} \\ D_{zz} \\ D_{xy} \\ D_{yz} \\ D_{xz} \end{bmatrix} = -\frac{1}{b} \begin{bmatrix} \ln \frac{S_1}{S_0} \\ \vdots \\ \ln \frac{S_i}{S_0} \\ \vdots \\ \ln \frac{S_n}{S_0} \end{bmatrix} \quad (2.9)$$

where $\hat{\mathbf{g}}_i = [x_i \ y_i \ z_i]^T$ is the i th gradient direction and S_i is the respective signal intensity. If $n > 6$, these equations can be solved in the least-squares sense by the singular value decomposition method (Press et al., 1988). Examples of estimated tensor coefficients are shown as images in Fig. 2.2.

2.3 Functional magnetic resonance imaging

Functional magnetic resonance imaging measures neural activity indirectly by exploiting the interrelations between physiological function, energy metabolism, and local blood supply. Various techniques exist to image different aspects of the hemodynamic response. For example, perfusion-based fMRI estimates the blood-supply changes by quantitatively measuring the blood flow. Here, we concentrate on the most widely used fMRI technique that employs the blood oxygenation level dependent (BOLD) MR signal (Ogawa et al., 1992).

Measuring neural activation using BOLD fMRI is possible because oxygenated and deoxygenated hemoglobin have different magnetic properties (Thulborn et al.,




Figure 2.2: Examples of diffusion imaging. The top row shows DWIs, one without weighting (S_0) and two with different directions of diffusion sensitizing gradient (direction shown by the white lines). The bottom part visualizes the computed six tensor coefficients. The intensity in these images is related to the absolute value of the respective coefficient and the off-diagonal elements are scaled by 5 to make them better visible. The schematic (lower left) shows fiber bundle locations (black lines) in the anterior and posterior parts of the corpus callosum. Note the signal drop in the DWIs when the bundles are aligned with the gradient direction. The fiber-direction related intensities in the tensor coefficients (e.g. D_{xx} and D_{xy}) can also be seen. White region in the illustration denotes the ventricles that have strong isotropic diffusion. *Imaged at 3 T at the AMI Centre of TKK.*

1982) and the local concentration of the deoxyhemoglobin changes during brain activation. The paramagnetic deoxyhemoglobin has a significant magnetic moment, unlike the diamagnetic oxyhemoglobin, and thus causes changes in the local magnetic field. These local field changes lead to faster spin dephasing in the neighboring protons, which the T_2^* -decay sensitive MR pulse sequences can detect as a signal drop.

The complex relationships between neural activity, blood flow, and blood oxygenation are not completely understood (Di Salle et al., 1999; Turner and Ordidge, 2000; Logothetis, 2002; Huettel et al., 2004; Lauritzen, 2005; Raichle and Mintun, 2006). The increased neural and metabolic activities cause an increase in the consumption of glucose and oxygen, thereby resulting in larger amounts of deoxygenated hemoglobin. However, the response measured by the BOLD-fMRI actually shows a signal increase relative to a resting state. This response suggests a decrease in the concentration of the deoxyhemoglobin in the tissues, likely resulting from the increase in the local cerebral blood flow supplying excessive amounts of oxygenated hemoglobin.

The signal change caused by the BOLD response is only a few percentages of the normal signal amplitude. Therefore, statistical analysis is typically required to detect subtle consistent changes from the image noise. In MRI, the SNR is proportional to the voxel volume and therefore the spatial imaging resolution also affects the signal strength. Typical spatial resolution for human fMRI is 1–5 mm, but animal experiments with high-field MRI devices can obtain even 0.1 mm resolutions. For localizing brain activations, the question still remains how the detected blood flow changes precisely relate to the actual neural activity.

The temporal resolution of the fMRI is in the order of a few seconds and is affected by many factors such as MRI hardware, imaging sequence, spatial resolution (number of slices), experimental setup, and the hemodynamic response (Huettel et al., 2004). However, despite the several second delay of the hemodynamic response, some specific measurements, such as differences in onset times, can be made with temporal accuracy of even tens of milliseconds (Menon et al., 1998; Menon and Kim, 1999; Formisano and Goebel, 2003).

2.4 Magnetoencephalography and electroencephalography

Magnetoencephalography (Cohen, 1972; Hämäläinen et al., 1993) and electroencephalography provide recordings that are directly related to electric currents in the neurons. MEG measures the resulting magnetic field with arrays of magnetometers or gradiometers placed outside the head whereas EEG uses electrodes to record the produced electrical potential distribution on the scalp.

The magnetic fields produced by neural activations and measured outside the head are very weak, typically on the order of 10^{-9} times the static geomagnetic field of the earth and 10^{-7} times the magnetic field disturbances caused by electrical equipment in a typical laboratory. The MEG measurements are usually carried out in a special magnetically shielded room using extremely sensitive superconducting quantum interference devices (SQUIDs) embedded in liquid helium.

In brain research, the main modeling task in both MEG and EEG is to estimate the source current distributions producing the recorded fields. This inverse problem is ill posed and further constraints or models for the sources are required. Furthermore, regularization is required to avoid error magnification in the inverse source estimates.

Typically, spherically symmetric models or boundary-element models (BEMs) are employed for the volume conductor and the source currents are modeled as constellations of point sources (equivalent current dipoles) or as continuous current distributions with location and orientation constraints derived from MRI. The required BEMs are usually constructed by segmenting and tessellating the respective surfaces from the individual anatomic images.

Both MEG and EEG have excellent temporal resolution and are able to follow neural activation on a millisecond scale. The spatial accuracy of the MEG is a few millimeters in favorable conditions (Hämäläinen et al., 1993), even with the simple spherically symmetric conductor models. The poorly conducting skull confines most of the volume currents inside the intra-cranial volume and, consequently, the forward modeling is considerably more straightforward than in EEG where the electrical potential on the scalp is heavily affected by such differently conducting layers and by their inhomogeneities. As a result, the spatial resolution and localization accuracy of EEG is inferior to MEG.

2.5 Computed tomography

Computed tomography (Hounsfield, 1980), originally known as computed axial tomography, collects a series of X-ray projection images taken around a single axis of rotation to compute a 3-D volumetric representation. The CT scans can provide anatomical information with good sub-millimeter spatial resolution and with minimal geometrical distortions.

The CT images show the absorption of X-ray beams which is related to the density of the tissue. Bone, soft tissue, and air cavities of the head are clearly visible and easy to separate by means of image segmentation. However, different soft-tissue classes, such as white and gray matter of the brain, cannot be well distinguished. Various contrast agents can be employed in CT scans to enhance the visibility of certain tissues or pathologies by imaging the build-up of the agent.

The use of the CT for brain research is very limited due to the inherent radiation dose and it is mainly applied in clinical settings for acute neurology and for pre-surgical or radiation-therapy planning.

Chapter 3

Image registration

A crucial step in multi-modal imaging is image registration which creates a correspondence (a mapping) between two different images. The result of the registration is a coordinate transformation that maps coordinates from one image to the coordinates of the respective location in the other image.

The process of image registration consists of three key elements: a specification of geometrical transformation, a match metric, and an optimizer. The geometrical transformation, with its parameters, specifies what kind of transformations are considered for the best registration. The match metric measures the goodness of the applied registration and allows to compare different candidates of transformation parameters. The optimizer tries to efficiently find the best registration from the parameter space of the selected geometrical transformation.

The following sub-sections describe these elements. Due to their importance in multi-modal image registration, mutual-information (MI) metrics have their own sub-section at the end, separate from the other match metrics. For a more in-depth view of medical image registration, see surveys by Maintz and Viergever (1998), Hill et al. (2001), and Pluim et al. (2003).

3.1 Transformations

The suitability of a given geometrical transformation depends on the registration application and on the medical images used. Different imaging modalities may have different geometrical distortions in their images and successful registration should allow for correction of these distortions. Furthermore, the purpose for which the registration is used might dictate which transformations are sensible.

3.1.1 Linear transformations

Linear (also known as affine, linear affine) transformation is usually suitable for intra-subject (within subject) registration of rigid body parts. In brain research, the head and brain are typically considered rigid because skull effectively limits soft-tissue movement and deformation. Linear transformation has been employed in a variety of cases, for example by Collignon et al. (1995), Hill et al. (1998), Maes et al. (1997), Studholme et al. (1996a, 1997, 1999), Thévenaz and Unser (2000), Viola and Wells (1995, 1997), Wells et al. (1996), and Zhu and Cochoff (2002). Under linear transformation, straight and parallel lines remain straight and parallel, but the angles between non-parallel lines are not necessarily preserved.

The simplest of linear transformations is the rigid-body transformation with 6 degrees of freedom (DOFs). It allows for translation (3 DOFs) and rotation (3 DOFs)

only and is sufficient for registering images from two correctly calibrated imaging modalities when image distortions are not present. Registrations using a rigid-body transform with (1 DOF) global scaling allow for uniform scale correction between the imaging modalities. Both of these special cases of linear transformation preserve the angles.

Non-uniform scale correction requires at least 3 DOFs for scaling, in addition to the rigid-body transform. However, this minimum of 9 DOFs limits the scale correction to a fixed set of coordinate axes, typically those in one of the registered images. Non-uniform scale correction with a free set of axes requires the full 12 DOFs affine transformation, which allows for translation, rotation, scale, and skew.

3.1.2 Non-linear transformations

In contrast to intra-subject registration, inter-subject (between subject) registration tries to match images from different subjects. Since the individual anatomy varies, higher-order transformations are required to allow for local deformations in addition to global matching. Non-linear registration is also useful in intra-subject registration to correct for non-linear distortions in images and to match images of soft-tissue body parts. In many of the non-linear methods, linear registration is performed first to bring the images in rough correspondence and the non-linear transformation then refines the registration result.

Several different non-linear (curved) transformations have been employed in medical imaging, and their usability typically depends on the requirements of the registration. For registrations that match feature-based homologous point landmarks, thin-plate splines (Bookstein, 1989) can be used for interpolating the transformation between these landmarks. This interpolation matches the landmarks exactly and the method has been extended into approximating transformations (Rohr et al., 1996, 2001; Rohr, 1998) that allow for errors in the landmark locations.

Above, the geometrical transform was directly linked to the match metric but usually these two are not coupled and the transform can be selected separately from the metric. Woods et al. (1998) and Horsfield (1999) use linear combinations of polynomial terms and Ashburner and Friston (1999) employ linear combinations of smooth basis functions to define the transformation. Thin-plate splines have been also applied with feature-independent control points (Kim et al., 1997; Meyer et al., 1997) when the match is defined by an image intensity based metric.

The non-linear geometrical transformations often employ control points in a regular grid or otherwise distributed throughout the images. These control points can be individually adjusted to define local deformations and the transformation is then propagated from the control points by linear interpolation (Kjems et al., 1999), by B-splines (Rueckert et al., 1999; Studholme et al., 2000, 2001; Rohlfing and Maurer, 2001; Schnabel et al., 2001), by Gaussian kernels (Gaens et al., 1998), or by other kernels (Lötjönen et al., 1999; Lötjönen and Mäkelä, 2001; Rohde et al., 2003).

Even if the transform is defined by control points, a common way to represent the final transformation is a vector field providing the local displacement vector for each pixel. To guarantee the smoothness of the transformation, this vector field can be explicitly smoothed (Thirion, 1998; Kjems et al., 1999; Guimond et al., 2001; Lau et al., 2001) or regularization terms can be used in the optimization (Kjems et al., 1999; Hellier et al., 1999; Hellier and Barillot, 2000; Hermosillo and Faugeras, 2001; Lötjönen and Mäkelä, 2001). Some authors base their transformations on physical models of

elastic materials and fluids (Christensen et al., 1994; Bro-Nielsen and Gramkow, 1996; Schormann et al., 1996; Christensen et al., 1996, 1997; Hata et al., 1998; Wang and Staib, 1998).

Special transformations for rigid bodies in deformable tissue have also been created (Little et al., 1997; Tanner et al., 2000) and they are suitable for modeling bones that can move with respect to each other. Such transformations are especially useful for imaging the spinal column, for example.

3.2 Match metrics

A match metric allows comparison of different candidates of transformation parameters. It represents the goodness of a particular geometrical transformation by a single scalar value which is consequently optimized. The metric can measure the registration error or distance from an optimal match, in which case the value is minimized. Alternatively, the metric can measure the similarity of the images, in which case the maximum is sought.

The match metrics can be divided into two classes: landmark or structure-based metrics and image intensity based metrics. The first class typically requires extensive preprocessing and possibly user intervention to extract the landmarks or structures for matching. On the other hand, metrics based on image intensities often use the acquired images directly with only minimal or no preprocessing.

The requirements for the match metric depend on the application and, in particular, on the two images being registered. Inter-modal registration is usually more difficult than intra-modal registration, due to the differences in images. One imaging modality might capture details that are not present in the images of the other modality, for example.

3.2.1 Landmark and structure-based metrics

Landmarks used in a registration metric can be anatomical, geometrical, or external. Anatomical landmarks are usually prominent, accurately locatable points in the visible anatomy that can be easily identified by the user. Geometrical landmarks are the extrema of some geometrical property, *e.g.* the locus of maximum curvature, and they are usually identified automatically. External landmarks employ markers attached to the skin surface or to a stereotaxic frame and consequently identified from the images by the user.

The benefit of landmark-based registration is that it can be applied to a variety of different medical data as long as the respective landmarks can be identified. Furthermore, the optimization is relatively fast because the images are reduced to a set of homologous landmarks. The landmark-based match metrics typically minimize the average distance between the corresponding landmarks and use most often the least-squares error measure (*e.g.* Arun et al., 1987). The geometrical transform using thin-plate splines (Bookstein, 1989), mentioned above, is a special non-linear transform matching the landmarks exactly and defining a smooth transform between the points.

Other structures, such as tissue boundaries and surfaces, are typically more distinct than landmarks in medical images. Various segmentation methods exist to delineate such high-contrast structures. The boundary between tissue and air is clearly visible in

most imaging modalities, which makes the skin surface a very widely used structure for matching.

Several different methods have been successfully applied in structure-based registration. Pelizzari et al. (1989) developed the “head-and-hat” algorithm for matching two surfaces representing the scalp. Here, the match metric is the squared distance between the points on the “hat” surface and the respective closest points on the head surface, as measured in the direction to the centroid of the head.

Borgefors (1988) employs a similar metric with distance transforms. A segmented binary image is transformed into a distance image, where each pixel contains the distance from that location to the segmented structure. Such a transform allows for fast computation of distances when a point-set is matched to the segmented structure.

Another very versatile registration method is the iterative closest point (ICP) algorithm (Besl and McKay, 1992) that can be used with a variety of different geometrical data-sets, including point sets, lines, curves, and surfaces. The algorithm matches a point set to geometrical data by first finding the closest locations on the geometrical structure to given points. A transform is then computed that minimizes the squared distance between these point pairs. As the transform changes, the closest points on the geometrical structure also change, and the procedure is repeated until it converges to a solution.

The drawback of the landmark and structure-based metrics is that their accuracy is limited by the accuracy of landmark and structure identification or segmentation. Some of the method also require user intervention in the form of landmark localization or image segmentation.

3.2.2 Metrics based on image intensity

Metrics based on image intensity employ directly the information in the pixel intensities. For intra-modality registration, the two registered images should have very similar intensities at the corresponding pixels when the images are correctly aligned. If the images differ only by Gaussian noise, the optimal match metric is the least-squares intensity difference (Friston et al., 1995; Hajnal et al., 1995b,a).

Usually, at least a global intensity scaling difference is present between the images, even for intra-modality registration. A least-squares metric with intensity rescaling (Alpert et al., 1996) has been proposed to correct for this difference. The ratio image uniformity (RIU) (Woods et al., 1992), frequently also called the variance of intensity ratios (VIR), is another metric to tackle the same problem. A ratio of image intensities is computed and the registration aims to minimize the variance of this ratio.

Cross-correlation metrics (Junck et al., 1990) have been widely employed directly to image intensities or to feature images extracted in a pre-processing step. If sufficiently similar features can be extracted from different modalities, the cross-correlation, as well as other intra-modality match metrics, can be applied also to inter-modality registration. Typically, high-contrast changes, such as intensity ridges (van den Elsen et al., 1995), are used.

Another approach to employ intra-modality metrics with inter-modality registration is to apply intensity re-mapping (Andersson et al., 1995). A modality-specific intensity mapping function is estimated that maps intensities from one modality to the corresponding intensities in the other modality. Typically, cross-correlation is used as the match metric between the re-mapped image and the other image.

Currently, the most versatile multi-modal match metrics employ techniques from information theory. In particular, metrics based on the concept of mutual information have gained a lot of attention. Due to their importance, these metrics will be described in detail in Section 3.4, below.

3.3 Optimizers

The third key element of image registration is the optimizer which finds the parameters of the geometrical transform that optimize the selected match metric. In this sense, image registration is like any other mathematical optimization problem where a function is either minimized or maximized with respect to a set of parameters. A variety of well-known optimization algorithms (see, e.g., Press et al., 1988) have been employed. Below, these algorithms as well as the most widely used speed-up technique are explored.

3.3.1 Algorithms

If the gradient of the match metric is impossible or too time consuming to compute, algorithms without gradient information can be employed. In image registration, the most popular methods of this class are the Powell's direction-set method and the downhill-simplex method (Maintz and Viergever, 1998; Hill et al., 2001; Pluim et al., 2003). Simulated annealing and genetic algorithms have been also employed. With the gradient information present, the most often used algorithms employ gradient-descent, conjugate-gradient, quasi-Newton, and Levenberg-Marquardt methods.

With the MI-based metrics, the gradient depends on the selected MI measure and also on the exact implementation. Wells et al. (1996) compute a stochastic approximation for the gradient employing Parzen windowing for the MI measure. Maes et al. (1999) show the exact gradient in the case of the partial-volume interpolation (PVI) (Collignon et al., 1995) method and also compare different optimization strategies. They conclude that the downhill-simplex, conjugate-gradient, and Levenberg-Marquardt methods provide the best performance for the MI registration in a multi-resolution optimization context. With the MI measures, the gradient of the match metric is relatively time consuming to evaluate and therefore the downhill-simplex method performs equally well with the gradient-based algorithms (Maes et al., 1999).

3.3.2 Multi-resolution techniques

The most widely used technique to speed-up optimization is to employ the multi-resolution scheme. The technique can be applied to the registered images, to the deformation field (deformation grid), or to both. When applied to the deformation field or grid, the technique is sometimes called multi-grid optimization.

In both cases, the idea is to down-sample either the images or the deformation field into a lower-resolution version. Several resolution levels can be employed, and the optimization starts with the most coarse level. When a sufficient match is found at a certain resolution level, the optimization continues at the next higher resolution.

In a lower resolution scale, the evaluation of the match metric is much faster and therefore the multi-resolution techniques greatly enhance the speed of convergence.

They also allow the global structures to be registered first and thereby typically increase the robustness of the applications.

3.4 Mutual-information registration

The information-theory concept of mutual information was first brought to the registration of medical images by Collignon et al. (1995) and Viola and Wells (1995). The MI metrics are based directly on image intensities and typically require no preprocessing. Furthermore, they make only minimal assumptions about the correspondence of the intensities between the images. Below, the definition and use of MI metrics are described. A detailed view with the current state of the art can be found in an excellent survey by Pluim et al. (2003).

3.4.1 Entropy

In information theory, the entropy is a measure of randomness in a probability distribution. The higher the uncertainty of the outcome of an event, the higher is the entropy. Entropy measures the “information content” of an event and it also relates to the minimum coding length of information for transmission. The most widely used entropy measure is the Shannon entropy (Shannon, 1948)

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x) = -E\{\log p(X)\} = E\{\log \frac{1}{p(X)}\} \quad (3.1)$$

where X is a discrete random variable, \mathcal{X} is the set of all possible outcomes, $p(x)$ is the probability of an outcome x , $E\{g(X)\} = \sum_{x \in \mathcal{X}} p(x)g(x)$ is the expected value of the random variable $g(X)$, and $a \log a = 0$ if $a = 0$.

The entropy $H(X)$ yields positive values, or 0 if the outcome is certain (*i.e.* $p(x) = 1$ for some x). The value of the entropy is at maximum for an even distribution, *i.e.* if $p(x) = 1/N$ for all x where N is the number of outcomes in \mathcal{X} . In this case, all outcomes are equally likely and the uncertainty of an event is at maximum.

For a pair of random variables X and Y with a joint-probability distribution $p(x, y)$, the joint (Shannon) entropy is

$$H(X, Y) = - \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x, y) \log p(x, y) = -E\{\log p(X, Y)\} \quad (3.2)$$

where \mathcal{Y} is the set of all possible outcomes for Y . Similarly, the conditional entropy $H(Y|X)$ is based on the conditional probabilities $p(y|x)$:

$$H(Y|X) = - \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x, y) \log p(y|x) = -E_{p(x, y)}\{\log p(Y|X)\}. \quad (3.3)$$

The probability $p(x, y)$ in the sub-index of the expectation E denotes the applied probability distribution for clarity. The chain rule for the entropies states that

$$H(X, Y) = H(X) + H(Y|X). \quad (3.4)$$

3.4.2 Mutual information

Mutual information is a measure of the dependence of two random variables X and Y (Cover and Thomas, 1991). It is based on the Kullback-Leibler distance between the joint distribution $p(x,y)$ and the product distribution $p(x)p(y)$ representing total independence. The MI metric $I(X,Y)$ measures the amount of information that one random variable contains about another, and it is defined as

$$I(X,Y) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x,y) \log \frac{p(x,y)}{p(x)p(y)} \quad (3.5)$$

$$= H(X) - H(X|Y) = H(Y) - H(Y|X) \quad (3.6)$$

$$= H(X) + H(Y) - H(X,Y) \quad (3.7)$$

with the following properties (Pluim et al., 2003):

- $I(X,Y) = I(Y,X)$.
Mutual information is symmetric.
- $I(X,X) = H(X)$.
The information a random variable X contains about itself (self-information) is equal to the information (entropy) in the variable X .
- $I(X,Y) \leq \min\{H(X), H(Y)\}$.
The information a random variable contains about another variable can never be greater than the information in either of the variables.
- $I(X,Y) \geq 0$.
The information cannot be negative, *i.e.* the uncertainty in a random variable X can never be increased by learning about Y .
- $I(X,Y) = 0$, if and only if X and Y are independent.
When X and Y are independent, X does not contain any information about Y .

3.4.3 Mutual-information in image registration

For applying MI in image registration, the intensities of two images A and B represent the random variables. The joint probability distribution $p(a,b)$ is estimated, and it represents the probability of having intensity a in A and intensity b in B at the same location. The mutual information $I(A,B)$ measures the change in uncertainty about an intensity in image A when the intensity in the corresponding location of B is given, and *vice versa*. The measure $I(A,B)$ reaches maximum when one image predicts best the values of the other image and this maximum is considered as the correct registration between the two images. For evaluating the MI measure $I(A,B)$, the required marginal distributions $p(a)$ and $p(b)$ are calculated by summation from $p(a,b)$.

The joint probability $p(a,b)$ is computed for the overlapping area of the two images. Therefore, the MI measure $I(A,B)$ is dependent on the extent of overlap between the images. A less sensitive normalized MI measure

$$Y(A,B) = \frac{H(A) + H(B)}{H(A,B)} \quad (3.8)$$

proposed by Studholme et al. (1999) has been widely employed to minimize this effect. Other MI-based measures mentioned in the literature include $\rho(A, B) = H(A, B) - I(A, B)$ and the entropy correlation coefficient $ECC(A, B) = 2I(A, B)/(H(A) + H(B)) = 2 - 2/Y(A, B)$ (Maes et al., 1997).

Pure MI measures are based only on the probabilities of co-occurring intensities without any spatial information about neighboring pixels. Both connected region labeling (Studholme et al., 1996b) and image gradients (Pluim et al., 2000a; Lötjönen and Mäkelä, 2001) have been used for integrating spatial information into the MI metric.

3.4.4 Joint-probability estimation

From the implementational point of view, the key aspect of any MI metric is the estimation of the joint probability distribution $p(a, b)$, and different methods have been developed. One frequently used method is Parzen windowing (Parzen, 1962) that allows the estimation of a probability distribution $p(x)$ from a set of samples S . The estimated probability is defined as

$$p(x) \approx P(x, S) = \frac{1}{n} \sum_{s \in S} W(x - s) \quad (3.9)$$

where n is the number of samples in S and $W(x)$ is a weighting function. Most implementations employ Gaussian weighting (Wells et al., 1996; Viola and Wells, 1997; Hermosillo and Faugeras, 2001; Mangin et al., 2001; Sarrut and Clippe, 2001), although double exponential function (Kim et al., 2000) and splines (Thévenaz and Unser, 2000) have also been used.

However, to estimate $p(a, b)$, the majority of the implementations (see survey by Pluim et al. 2003) use a straight forward alternative employing joint histogram. The joint histogram $h(a, b)$ represents the number of times when intensity a in image A coincides with intensity b in image B . It is typically computed by sampling through all the pixels of one image and collecting the values at respective places in the other image. Division of the histogram bins by the sum of all entries normalizes the values and yields the joint-probability distribution $p(a, b)$.

In image registration, the geometrical transformation does not typically map the centers of pixels (grid-points) of one image to the grid-points of the other image. Therefore, the joint histogram $h(a, b)$ cannot be computed just by picking the intensities from both images. This problem can be overcome by interpolation or fractional incrementation.

Applications of the first class employ some interpolation method to calculate the intensity values at non-grid locations. Typically, linear interpolation (LI) is chosen for its speed, but also higher order interpolation methods can be used. Consequently, the histogram bin corresponding to the interpolated intensity is incremented by one. On the other hand, fractional incrementation methods do not compute a new intensity value from the surrounding pixels but increment the histogram bins corresponding to these surrounding pixels by fractional amounts. This idea was first introduced by Collignon et al. (1995) in a form of the PVI and it was later generalized by Chen and Varshney (2003) to higher-order kernels.




Figure 3.1: Artifacts in the MI metric $I(A,B)$ as a function of translation (left) and rotation (right) with different sampling methods. The vertical lines show the correct registration at -0.2 pixels for translation and 0.1 degrees for rotation. *Adapted from publication P1.*

3.4.5 Artifacts

In certain situations, the two most used joint-histogram estimation methods, LI and PVI, produce artifacts in the MI metric (Pluim et al., 2000b; Ji et al., 2003; Tsao, 2003). These artifacts manifest especially when the axes of the images are parallel and the pixel sizes are equal. In such circumstances, these methods tend to “grid-lock” due to the artifacts and therefore the registration accuracy is limited to pixel size. Examples of such artifacts in the MI measure $I(A,B)$ are shown in Fig. 3.1. The BS_4 denotes the support-4 B-spline method explained below in Section 3.4.6.

The origin of these artifacts is different for the interpolation and the fractional incrementation methods. Interpolation tends to blur images and to affect small details and noise (Pluim et al., 2000b; Tsao, 2003). It can also produce new intensity values that do not occur in the original image, for example in images with only a few distinct intensity levels. Both of these effects spread the joint histogram and thereby affect the entropies. The amount of spread varies with the pixel-grid alignment and thereby creates the artifact patterns visible in Fig. 3.1.

The fractional-incrementation methods do not produce new intensity values but they transfer the histogram weights gradually from one original sample to another as the location changes. The PVI method uses linear weights and the artifacts are attributed to the non-linearity of the entropy (Pluim et al., 2000b). For example, if the weight from one bin is distributed equally into two bins, the total entropy increases ($H(p) < 2H(p/2)$ if $p > 0$). As above, this effect depends on the pixel-grid alignment.

Although the MI measures are theoretically symmetric, the implementational details cause asymmetry to the metric. Typically, both the interpolation and fractional-incrementation methods sample through all the pixel centers of image A when the histogram $h(a,b)$ is estimated. A slightly different histogram would result if the occurred in the pixel centers of image B instead.

3.4.6 Improvements in accuracy (P1)

Publication P1 introduces a novel theoretical concept of continuous fractional incrementation. In a normal fractional-incrementation method, the histogram $h(a,b)$ is sampled at every pixel of image A overlapping image B . The respective location in B is

computed and the surrounding pixels determine which histogram entries receive the incrementation. The amount of incrementation, *i.e.* the fractional weight, is determined by a kernel $\varphi_{\text{fi}}(x)$ chosen by the application. The kernel function for fractional incrementation has to fulfill the following two conditions (Chen and Varshney, 2003) for all $x \in \mathbb{R}$:

$$\varphi_{\text{fi}}(x) \geq 0 \quad (3.10)$$

$$\sum_{n=-\infty}^{\infty} \varphi_{\text{fi}}(n+x) = 1 \quad (3.11)$$

For example, the PVI method uses the same kernel function as the linear interpolation:

$$\varphi_{\text{PVI}}(x) = \varphi_{\text{LI}}(x) = \begin{cases} 1+x, & \text{if } -1 \leq x < 0 \\ 1-x, & \text{if } 0 \leq x \leq 1 \\ 0, & \text{otherwise.} \end{cases} \quad (3.12)$$

In contrast to sampling at every pixel center of A , the newly proposed method samples the images for histogram $h(a, b)$ everywhere, in a continuous fashion. Publication P1 demonstrates that this theoretical sampling concept can be reduced into the normal fractional-incrementation procedure in the special circumstances where the most severe artifacts appear. Under these circumstances, the continuous sampling with kernel $\varphi_{\text{fi}}(x)$ is equivalent to employing the normal fractional-incrementation procedure with the auto-correlated kernel $\varphi_{\text{fi}}(x) * \varphi_{\text{fi}}(x)$. Publication P1 further shows that this auto-correlated kernel fulfills the conditions 3.10 and 3.11 if the original kernel $\varphi_{\text{fi}}(x)$ also does. Therefore, the kernel produced by auto-correlation can be used in every situation, and not only in the special circumstances in which it was developed.

The novel concept of continuous sampling provides new insight into the kernels already in use. For example, the well-known PVI kernel is the auto-correlation result of the nearest-neighbor kernel

$$\varphi_{\text{NN}}(x) = \begin{cases} 1, & \text{if } |x| < \frac{1}{2} \\ \frac{1}{2}, & \text{if } |x| = \frac{1}{2} \\ 0, & \text{otherwise.} \end{cases} \quad (3.13)$$

Therefore, applying the PVI method can be considered as sampling with $\varphi_{\text{NN}}(x)$ in a continuous fashion¹. Similarly, the higher-order kernels suggested by Chen and Varshney (2003) can be connected to the lower-order kernels. For example, the auto-correlation of the PVI kernel produces support-4 third-order B-spline kernel (BS₄ in Fig. 3.1).

The experiments in Publication P1 study the registration errors of the linear interpolation, the PVI, and the BS₄ methods with sub-sample accuracy. The effect of image quality, in the form of noise and anti-aliasing, as well as method implementation parameters are explored. The results show that the higher-order BS₄ method provides slightly better accuracy in normal cases and significantly better registration accuracy under the worst case scenario. In circumstances when the sampling-related artifacts are typically present, the BS₄ method has the average and maximum errors of 0.04 and 0.2 pixels, respectively, whereas both the LI and the PVI methods produce average and maximum errors of 0.2–0.3 and 0.5 samples, respectively. To achieve the better accuracy, the BS₄ method requires a computation time 2–4 times that of the LI and PVI methods.

¹under the special circumstances

Chapter 4

Resampling

Image registration produces a coordinate transform (a mapping) between registered images. For joint visualization or for analysis, one of the images is typically transformed (mapped) into the coordinate space of the other image. Alternatively, both of the images can be transformed into some other, common coordinate frame.

This process of transforming an image from one pixel-grid to another is called resampling and is a very common process that can happen explicitly or implicitly in many stages. For instance, if an image is not shown on a computer screen pixel-by-pixel, some form of resampling has taken place.

Naturally, high-quality resampling algorithms are needed to minimize artifacts and the loss of information. Even if the artifacts are not readily visible, they can affect the image analysis. The effects of interpolation in mutual-information registration form a good example of such artifacts, see Page 23.

In resampling, a pixel-grid point (*i.e.* the center of a pixel) in one image does not generally transform onto a grid-point in the other image. Therefore, an interpolation method is required to estimate the pixel value at a particular location between the grid points. This Chapter reviews methods commonly used in medical image processing.

First, the relevant aspects of the sampling theory are briefly discussed. The theory is explored in the one-dimensional case, and for clarity we refer to the image and the pixels as signal and samples, respectively. Next, the traditional convolution interpolation and the generalized interpolation methods are described. Finally, two specialized Fourier space methods and the two-stage method developed in Publication P2 are explored.

4.1 Sampling theory

The Nyquist-Shannon sampling theorem (Shannon, 1949) states that exact reconstruction of continuous-time signal $g(t)$ ($t \in \mathbb{R}$) from its samples $x(n) = g(nT)$ ($n \in \mathbb{Z}$) is possible if the signal is band-limited and the sampling frequency f_s is greater than twice the signal bandwidth. Here, T is the sampling interval ($f_s = 1/T$) and the Nyquist frequency (“critical frequency”) of the sampling system is $f_s/2$. If all signal frequencies are below the Nyquist frequency, the original signal can be completely recovered by

$$g(t) = \sum_{n \in \mathbb{Z}} x(n) \cdot \text{sinc}(f_s t - n) \quad (4.1)$$

where $\text{sinc}(t) = \sin(\pi t)/(\pi t)$ is the normalized sinc function, and $\text{sinc}(0) = 1$.

The sampling and reconstruction above can be represented in an alternative way by using the Dirac delta function (unit impulse function) $\delta(t)$ which is a special function

that has the value of infinity at $t = 0$, the value zero elsewhere, and it integrates to unity: $\int_{-\infty}^{\infty} \delta(t) dt = 1$ (Bracewell, 1986). Using the Dirac comb (impulse train)

$$\Delta_T(t) = T \sum_{n \in \mathbb{Z}} \delta(t - nT) \quad (4.2)$$

the sampling can be represented as multiplication with $\Delta_T(t)$:

$$g_s(t) = g(t) \cdot \Delta_T(t) = T \sum_{n \in \mathbb{Z}} g(t) \delta(t - nT) \quad (4.3)$$

$$= T \sum_{n \in \mathbb{Z}} g(nT) \delta(t - nT) = T \sum_{n \in \mathbb{Z}} x(n) \delta(t - nT) \quad (4.4)$$

where the sampled signal $g_s(t)$ retains the information in $g(t)$ at the discrete time points $t = nT$ and is zero elsewhere.

Now, the reconstruction of $g(t)$ can be expressed as a (continuous) convolution between the sampled signal $g_s(t)$ and the function $h_{f_s}(t) = f_s \text{sinc}(f_s t)$:

$$g(t) = g_s(t) \circ h_{f_s}(t) \quad (4.5)$$

$$= \left\{ T \sum_{n \in \mathbb{Z}} x(n) \delta(t - nT) \right\} \circ f_s \text{sinc}(f_s t) \quad (4.6)$$

$$= \sum_{n \in \mathbb{Z}} x(n) \{ \delta(t - nT) \circ \text{sinc}(f_s t) \} \quad (4.7)$$

$$= \sum_{n \in \mathbb{Z}} x(n) \cdot \text{sinc}(f_s(t - nT)) = \sum_{n \in \mathbb{Z}} x(n) \cdot \text{sinc}(f_s t - n). \quad (4.8)$$

Convolution with the signal domain function $h_{f_s}(t)$ is equivalent to ideal low-pass filtering in the frequency domain. The rectangular pulse

$$\text{rect}_{f_s}(f) = \begin{cases} 1 & \text{if } |f| \leq f_s/2 \\ 0 & \text{otherwise} \end{cases} \quad (4.9)$$

in the frequency domain can be inverse Fourier transformed as

$$\mathcal{F}^{-1} \{ \text{rect}_{f_s}(f) \} = \int_{-f_s/2}^{f_s/2} e^{i2\pi ft} df = f_s \text{sinc}(f_s t) = h_{f_s}(t). \quad (4.10)$$

Comparison of the Shannon's Eq. 4.1 and the convolution Eq. 4.5 reveals that the convolution with $h_{f_s}(t) = f_s \text{sinc}(f_s t)$ is equivalent to Shannon's method. Furthermore, Eq. 4.10 shows that the signal-domain convolution with $h_{f_s}(t)$ is equal to ideal low-pass filtering in the frequency domain. In a summary, all three equations and methods reconstruct the original continuous-time signal from the samples.

4.2 Convolution interpolation

In the following sections, unit sampling frequency ($f_s = 1$) is assumed for clarity and the theory is explored in terms of sample spacing. The convolution interpolation can be then expressed as

$$g_{\text{int}}(t) = g_s(t) \circ \varphi_{\text{int}}(t) \quad (4.11)$$

where $g_{\text{int}}(t)$ is the interpolated continuous signal and $\varphi_{\text{int}}(t)$ is the interpolating convolution kernel. To reproduce the original samples $g_{\text{int}}(k) = g_s(k)$ for all $k \in \mathbb{Z}$, the




Figure 4.1: Interpolating convolution kernels: nearest neighbor (solid line), linear (dashed line), and sinc (dotted line). Circles mark the values that interpolating kernels must obtain for integer offsets.

kernel has to fulfill $\varphi_{\text{int}}(0) = 1$ and $\varphi_{\text{int}}(k) = 0$ for all integer $k \neq 0$ (circles in Fig. 4.1). Similarly as above, the continuous convolution can be converted into sum

$$g_{\text{int}}(t) = \sum_{n \in \mathbb{Z}} x(n) \cdot \varphi_{\text{int}}(t - n). \quad (4.12)$$

For implementing the convolution interpolation algorithm, the summation in Eq. 4.12 is employed. The support of kernel $\varphi_{\text{int}}(t)$, *i.e.* the range of t where $\varphi_{\text{int}}(t)$ assumes non-zero values, affects the computational complexity of the selected kernel. It determines how many samples $x(n)$ and how many multiplications are required for each interpolated outcome. Therefore, the choice of the interpolating kernel $\varphi_{\text{int}}(t)$ is a trade-off between the computational effort and the interpolation quality.

As is obvious from the sampling theory (Eq. 4.5), kernel $\varphi_{\text{int}}(t) = \text{sinc}(t)$ (see Fig. 4.1) is the optimal interpolator that reproduces exactly the original continuous signal $g(t) = g_{\text{int}}(t)$ if all signal frequencies were below the Nyquist limit. However, this kernel is impractical as the sinc function extends to infinity and either truncated or windowed sinc kernels have to be employed. Furthermore, the sinc function decays very slowly and therefore relatively large support kernels are needed. These aspects make the sinc-based convolution-interpolation methods very slow.

The fastest convolution interpolation kernels are the nearest neighbor (Eq. 3.13) and the linear interpolation (Eq. 3.12) kernels (see Fig. 4.1). Especially the support-2 linear interpolation is widely employed due to its speed. Other common finite-support methods include cubic (Keys, 1981) and polynomial (Meijering et al., 1999) interpolation. Surveys by Lehmann et al. (1999, 2001) and by Meijering et al. (2001) provide a comprehensive comparison of different interpolation kernels.

4.3 Generalized interpolation

Similarly as above, the generalized interpolation (GI) introduced by Blu et al. (Blu et al., 1999, 2001; Thévenaz et al., 2000) is also written as a summation

$$g_{\text{int}}(t) = \sum_{n \in \mathbb{Z}} y(n) \cdot \varphi(t - n). \quad (4.13)$$

Here, the original samples $x(n)$ of Eq. 4.12 are replaced by prefiltered samples $y(n)$, and the interpolating kernel $\varphi_{\text{int}}(t)$ is replaced by a general kernel $\varphi(t)$. The prefiltering

creating $y(n)$ from $x(n)$ is linked to $\varphi(t)$ by the interpolation condition $g_{\text{int}}(n) = g_s(n) \forall n \in \mathbb{Z}$, i.e. the function $g_{\text{int}}(t)$ must reproduce the original samples.

The support of the kernel is the most crucial factor determining the computational efficiency of convolution interpolation (Eq. 4.12) and GI (Eq. 4.13), especially with multi-dimensional data. If the support for the kernel is S and the data set has D dimensions, the total of S^D multiplications are required for each interpolated output value. For example, support-2 and support-4 kernels require $2^3 = 8$ and $4^3 = 64$ multiplications, respectively, in the 3-D case.

For convolution interpolation, larger support kernels allow for better interpolation quality since more samples are taken into account. The generalized interpolation tries to circumvent this effect by prefiltering the data and thereby spreading the information in the samples. Therefore, a GI kernel has information available from a larger area using the same support size. For typical GI methods, the prefiltering can be implemented efficiently as a succession of forward and backward recursive filters (Blu et al., 1999) and the extra time required is minimal. Furthermore, if repeated interpolation is performed on the same image, the data can be kept in the prefiltered form for efficiency.

Due to the required prefiltering, spline interpolation (Hou and Andrews, 1978; Unser et al., 1993a,b; Unser, 1999) and the shifted linear interpolation (Blu et al., 2004) also fall into this GI category. Blu et al. (2001) introduced a class of maximal-order-minimal-support (MOMS) functions made of linear combinations of B-splines and its derivatives. These functions were shown to have minimal support for a given interpolation accuracy. The optimal versions were labeled as OMOMS and they are currently one of the most accurate interpolation methods.

4.4 Specialized resampling methods

Some specialized methods exist for certain image resampling tasks. For example, the rotation matrix in 2-D image rotation can be decomposed into two or three one-dimensional image shears (Catmull and Smith, 1980; Tsuchida et al., 1987; Unser et al., 1995). In shape-based interpolation (Goshtasby et al., 1992; Grevera and Udupa, 1996), the interpolation process itself is influenced by the underlying images. Typically, the interpolation kernels are separable although non-separable two-dimensional cubic kernels have been studied as well (Reichenbach and Geng, 2003).

Below, two different Fourier-space methods and a two-stage resampling method are explored in more detail. The Fourier-space methods efficiently employ full-width sinc interpolation through discrete Fourier transforms (DFTs) and they are especially suited for data up-sampling. The two-stage resampling method combines these up-sampling methods with traditional interpolation methods for general resampling.

4.4.1 Fourier-space methods

Fourier space methods usually employ the fact that a signal-space convolution corresponds to multiplication with a transfer function in the Fourier domain. Large-support convolution kernel can therefore be efficiently implemented as a (complex) multiplication of the transformed signal and the transfer function. Below, two Fourier space up-sampling methods and their limitations are discussed after a brief summary of the DFT.

Discrete Fourier transform

The discrete Fourier transform $G(u)$ of an N -sample array $g(x)$ and its inverse are defined as

$$G(u) = \mathcal{F}\{g(x)\} = \sum_{x=0}^{N-1} g(x)e^{-i2\pi ux/N} \quad (4.14)$$

$$g(x) = \mathcal{F}^{-1}\{G(u)\} = \frac{1}{N} \sum_{u=0}^{N-1} G(u)e^{i2\pi ux/N} \quad (4.15)$$

where $i = \sqrt{-1}$, and $\mathcal{F}\{\dots\}$ and $\mathcal{F}^{-1}\{\dots\}$ denote the Fourier and the inverse Fourier transforms, respectively. The Fourier transform and the inverse Fourier transform have several well-known properties (Gonzalez and Woods, 1992) such as separability, translation, periodicity, and conjugate symmetry.

Shifted-DFT (SDFT) algorithm

The specific task of up-sampling a signal with a discrete factor can be performed by sub-sample shifts (Du et al., 1994; Yaroslavsky, 1997). The original signal and the sub-sample shifted signals can be composed together to form a higher-resolution version. For example, interleaved original signal and signals shifted by $1/3$ and $2/3$ samples form a up-sampled signal with magnification factor of 3.

The sub-sample shift can be performed by an appropriate phase-shift in the Fourier domain. The signal is Fourier transformed, multiplied by a transfer function $H_s(u)$ causing the shift, and finally inverse transformed. It is crucial to note that the well-known Fourier-space phase-shift equation $H_d(u) = e^{i2\pi ud/N}$ is valid only for discrete shifts of d samples. For sub-sample shifts of s samples, the Fourier-space transfer function is (Yaroslavsky, 1997)

$$H_s(u) = \begin{cases} e^{i2\pi us/N} & \text{if } 0 \leq u < N/2 \\ \cos(\pi s) & \text{if } u = N/2 \\ e^{-i2\pi(N-u)s/N} & \text{if } N/2 < u \leq N-1 \end{cases} \quad (4.16)$$

In case N is odd, the middle term $\cos(\pi s)$ is left out in Eq. 4.16.

Zero-filled up-sampling

The Fourier-space zero-filling methods employ ideal low-pass filtering in the frequency domain (Eq. 4.9). The Fourier-transformed sample array is extended to contain higher frequencies that are consequently filled with zeros. The inverse Fourier transform yields the up-sampled signal.

With this method, the original signal length N and the final length $L > N$ can be any size. However, due to the available Fast Fourier Transform (FFT) algorithms, both of these values are often restricted to powers of two. In the usual case where N is even, the sample $G(N/2)$ at the Fourier transformed array contains aliased frequencies and therefore needs to be handled specially when converting into the extended array (Yaroslavsky, 1997; Seppä, 2007).

The factor for up-sampling is typically 2 or higher, meaning that the zero-filled array contains mostly zeros. With the normal FFT algorithms, the inverse transform leads to many unnecessary multiplications and additions with zeros. To remediate




Figure 4.2: Illustration of the up-sampled stage (on the right) of the two-stage resampling method. White frame indicates the position of the new image to be computed. *Adapted from Publication P2.*

this drawback, Markel (1971) developed a pruned decimation-in-frequency type FFT algorithm. Similarly, Skinner (1976) introduced a slightly more efficient version for the decimation-in-time type FFT algorithm. These algorithms were originally developed for computation of high-resolution spectra and modified FFT algorithms for zero-filled up-sampling has been also introduced (Nagai, 1986; Smit et al., 1990; Seppä, 2007).

Limitations

Fourier-space methods bear some limitations that arise from the FFT algorithms used. Typically, radix-2 FFT algorithms are employed and therefore the signal array sizes need to be powers of two. Naturally, this obstacle can be circumvented by (zero) padding the array to the next suitable size. The up-sampling factor needs to be an integer for the SDFT method and a power of two for the zero-filled up-sampling methods with radix-2 algorithms. Fractional factors are also possible by discarding part of the up-sampled data. For example, factor $4/3$ can be realized by up-sampling first by factor 4 and then selecting every third sample.

With the Fourier-space methods above, the up-sampled image grid is always aligned with the original image grid and free placement (rotation, scaling, etc.) of the resulting grid is not possible. Furthermore, the Fourier transforms assume implicitly that the images are periodic, *i.e.* the data repeats end-to-end. Therefore, the up-sampled pixels at one edge of an image depend on the pixels at the opposite edge.

Fourier-space methods also compute the up-sampling simultaneously for the whole image. Therefore, they are not necessarily efficient if the result is needed only for a relatively small part of the image.

4.4.2 Two-stage resampling (P2)

The two-stage resampling method introduced in Publication P2 combines the Fourier-space methods above with the traditional interpolation methods. The first stage employs Fourier-space methods for high-quality up-sampling of the image. Next, the second stage computes the resampled image from this up-sampled intermediate version (see Fig. 4.2). Either normal convolution interpolation or the generalized interpolation can be employed in the second stage. The GI prefiltering is conveniently performed by transfer-function multiplication during the first-stage Fourier-space processing with virtually no extra computational cost. Publication P2 also represents an improved algorithm for zero-filled up-sampling that provides a slight 1–8% speed increase compared to the previously used algorithms.

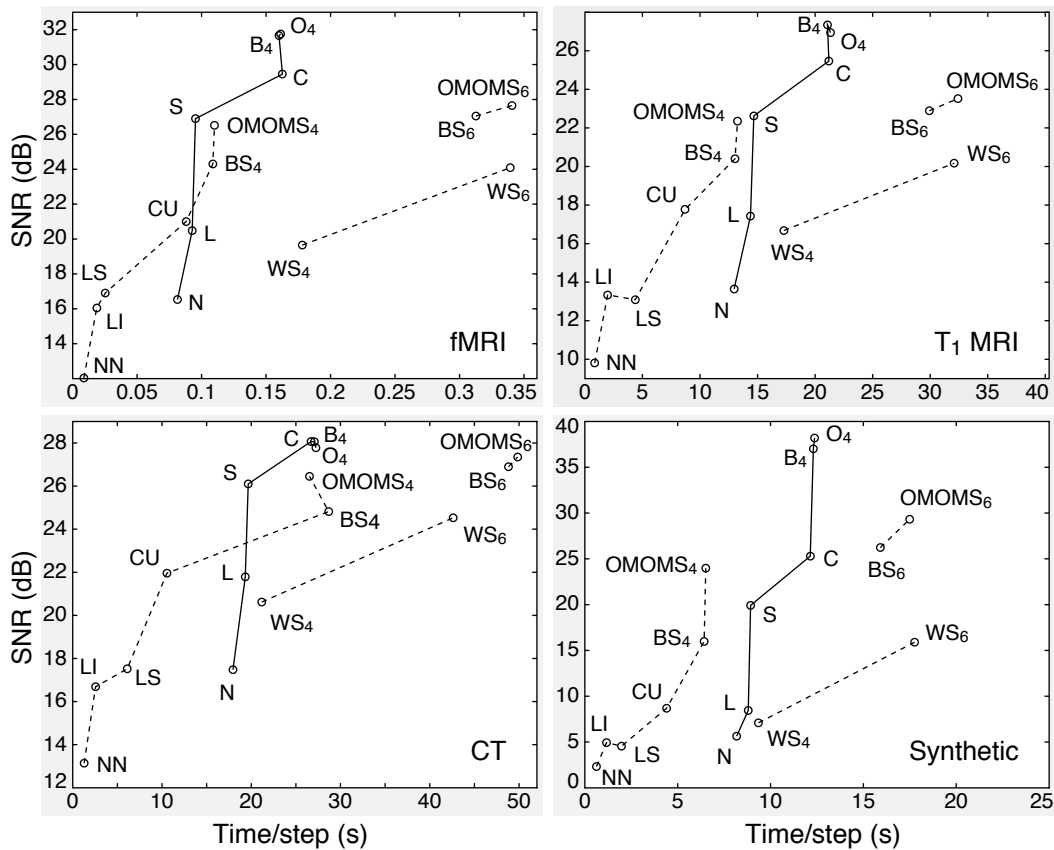


Figure 4.3: Quality and speed of different resampling methods with 3-D data. For details of the abbreviations of the methods, see text. *Adapted from Publication P2.*

The combination of these two stages has many benefits. Because high-quality full-width sinc interpolation is used in the up-sampling, the second stage can employ relatively small-support interpolation kernels. The size of the support is always a trade-off between computational efficiency and interpolation quality, and the high-quality intermediate stage now leads to small interpolation errors even with small-support kernels. Especially in 3-D medical imaging, the size of the support is the most crucial factor for fast computation. The combination of the two stages also circumvents the limitations of the first stage. Specifically, the second stage allows free positioning and sizing of the resampling grid which is not possible with Fourier-space methods alone.

The experiments in Publication P2 show that the two-stage resampling method produces very high-quality results. With a suitable up-sampling factor and second-stage kernel, the resampling errors are smaller than for any typically used resampling method. The results show that the up-sampling factor 2 is optimal in normal cases in combination with support-4 kernels in the second stage. Special cases requiring repetitive resampling of the same image (*e.g.* image registration) can benefit from higher up-sampling factors in combination with faster support-2 kernels. In such circumstances, the first stage is computed just once and only the second stage is repeated.

Figures 4.3 and 4.4 illustrate the speed and quality of various resampling methods. The quality is measured by the signal-to-noise ratio (SNR) after 15 successive resampling steps. In each panel, the solid line combines the two-stage methods and they all use up-sampling with factor 2. The letter identifies the second-stage resampling method and is either nearest neighbor (N), linear (L), shifted-linear (S), cubic (C), support-4 B-spline (B₄), or support-4 OMOMS (O₄). Dashed lines connect the

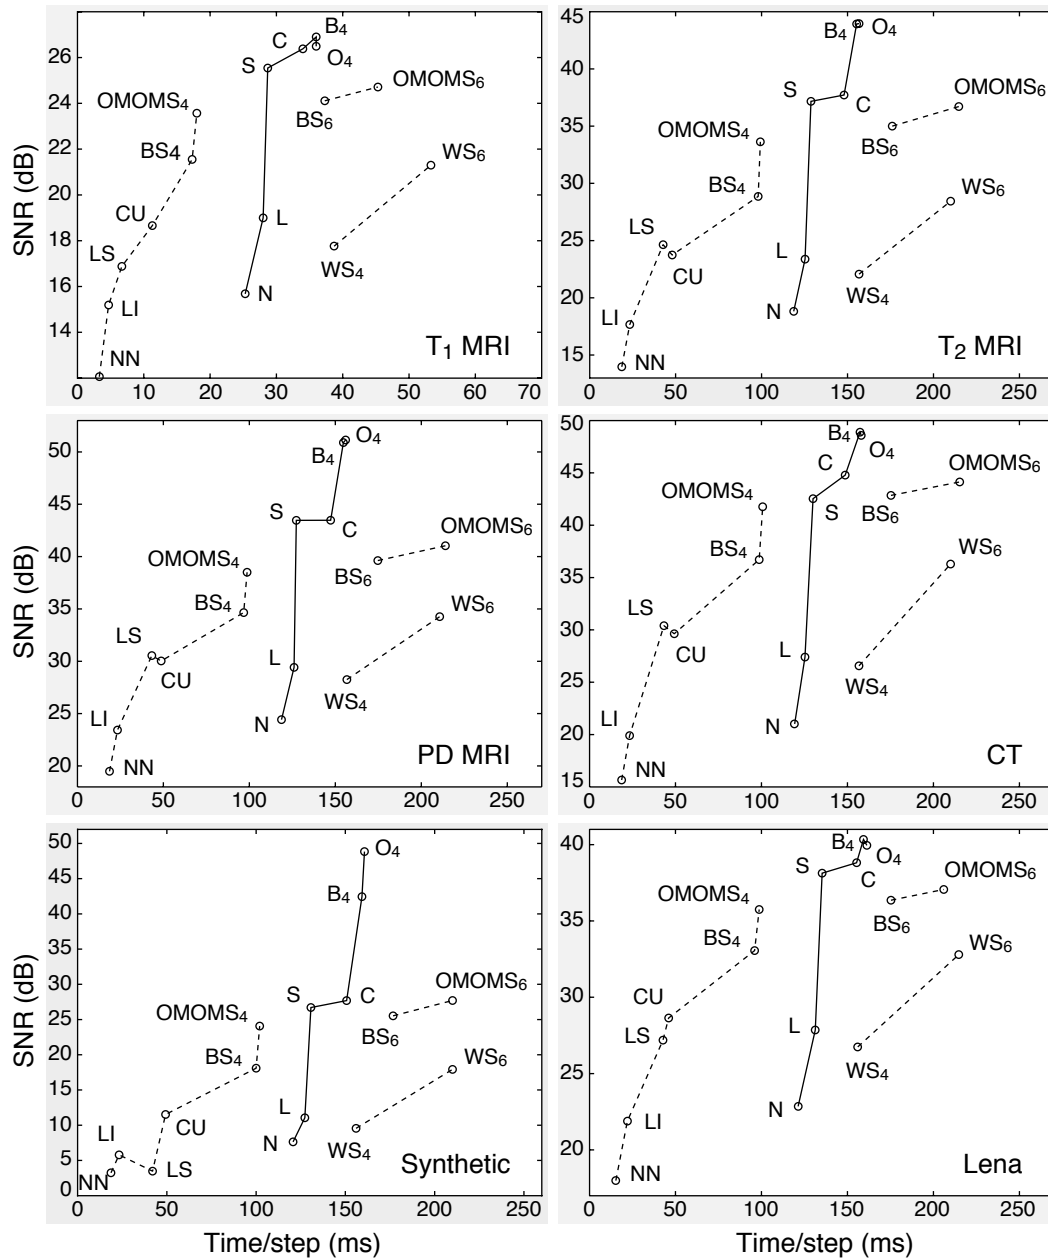


Figure 4.4: Quality and speed of different resampling methods with 2-D data. For details of the abbreviations of the methods, see text. *Adapted from Publication P2.*

previously known methods (in 3 groups) which employ nearest neighbor (NN), linear (LI), shifted-linear (LS), Keys cubic (CU), B-spline (BS_x), OMOMS ($OMOMS_x$), or windowed sinc (WS_x) interpolation. The windowed sinc methods employ Hamming windowing and the subscript x in the labels indicate the size of the support.

The 2-D image or 3-D volume used in the experiments is identified in the lower-right corner of each panel. The experiments employed T_1 -weighted MRIs (T_1 MRI), T_2 -weighted MRIs (T_2 MRI), proton-density weighted MRIs (PD MRI), computed tomography images (CT), functional MRIs (fMRI), synthetic images (Synthetic), and a gray-scale photograph (Lena).

As is shown in Publication P2 and evident in Figs. 4.3 and 4.4, the two-stage resampling method is capable of producing higher-quality results than the previously known methods. With support-4 cubic (C), B-spline (B_4), or OMOMS (O_4) kernels, the two-stage resampling is faster and produces better results than normal support-6 methods. The method is best suited for applications requiring minimal resampling errors. Such applications typically perform statistical analysis (*e.g.* fMRI data analysis) or otherwise further process the resampled data.

Chapter 5

Visualization

Visualization, as the final step of medical imaging, should be easy and intuitive to understand, leaving no room for misinterpretations. If interactive real-time manipulation of the view is possible, the identification of 3-D spatial relationships becomes easier. Therefore, the speed of a particular visualization method also affects the usability of the method.

Typically, analysis results of the functional images or other processed information is overlaid on the anatomical images. Different tissue properties, borders, and structures of the anatomical images allow identification of the physical locations of the overlaid data. The visibility of such anatomy can be further enhanced by extracting various structures by means of image segmentation and by visualizing them as individual objects.

Nowadays, the typically employed medical imaging methods produce 3-D volumetric data, especially in brain imaging. Therefore, the following sections concentrate on exploring volume visualization and the techniques are divided into scalar and tensor volumes, according to the type of information present in each voxel.

5.1 Scalar volumes

Scalar volumes contain one value for each voxel and the interpretation for the value depends on the imaging method. These volumes are not necessarily even the direct outcome of the imaging but might be the results of further processing and data analysis. A single imaged volume might even lead to several different scalar volumes, each expressing different analysis results. In the simplest form, the imaged values might directly represent densities, concentrations of particular substances, or some other physical characteristics of the tissue. On the other hand, the analysis results could express probabilities or confidences of a particular event (*e.g.* of neural activation), estimated current densities, or some more complex physical properties such as diffusion anisotropies. Below, the most common methods to visualize scalar volumes are explored.

5.1.1 Cross-sections

Two-dimensional planar cross-sections (Fig. 5.1) are perhaps the most common way to visualize 3-D volume data. A 2-D plane cuts through the volume and the values on this plane are visualized. In the simplest approach, the 2-D plane is constrained to match a pair of axes of the 3-D volume data. No interpolation is necessary if each 2-D pixel on the cut plane coincides with a 3-D voxel, and the data can be directly copied into




Figure 5.1: Two cross-section images of a 3-D T_1 -weighted MRI volume. One original sagittal slice is shown on the left (no interpolation required) with a line representing the location of the interpolated oblique cross-section on the right.




Figure 5.2: Examples of surface-based rendering using triangle meshes. The cortical surface (left) is inflated (right) to reveal areas hidden inside sulci. *Visualization employs inflation techniques of Fischl et al. (1999). Images by courtesy of Lauri Parkkonen.*

the cross-section image. However, such constraints severely limit the possible views and 3-D interpolation can be employed to allow for oblique cross-sections. Tri-linear interpolation is commonly used for speed, although higher-order methods can be also employed for better quality results.

5.1.2 Surface-based rendering

Surface-based rendering methods, also known as indirect volume rendering (IVR) methods, require preprocessing to extract objects of interest from the 3-D volumes. Such objects can be iso-surfaces obtained by simple thresholding or they might require more complex segmentation. After extraction, the surfaces are represented by surface primitives, such as polygons or patches. In medical imaging, triangle meshes are by far the most commonly employed surface representations.

The surface extraction reduces the number of data primitives significantly. A 3-D volume contains typically millions of voxels that are consequently represented by some thousands or hundred thousands of surface primitives. The visualization of these reduced number of primitives is usually much faster than the original 3-D volume data. Current display hardware and graphics processing units (GPUs) are especially

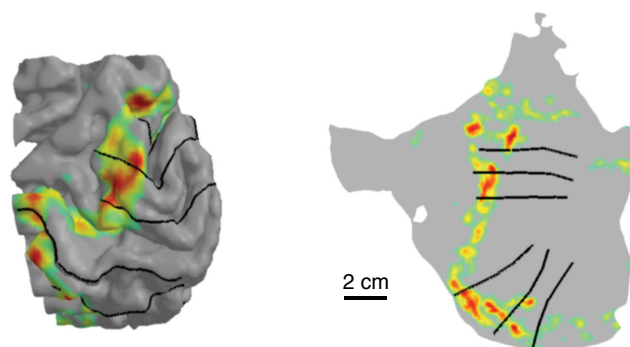


Figure 5.3: Example of the flat maps where part of the right visual cortex (left) is unfolded onto a plane (right) to reveal the fMRI activation patterns. *Visualization is computed using Brain à la Carte software employing retinotopic mapping techniques of Warnking et al. (2002). Image by courtesy of Linda Henriksson.*

suited for fast rendering of triangle meshes which makes them very attractive surface representations.

As an additional benefit, triangle meshes and other surface representations allow further processing and analysis of the surfaces. For example, local surface curvatures or distances measured along the surfaces can provide useful information. Highly convoluted surfaces, such as the human cortex, can be inflated (Fig. 5.2) and even unfolded onto a plane (flat maps, Fig. 5.3) to reveal surface areas otherwise hidden from view (Fischl et al., 1999). Due to the intrinsic curvature of the unfolded surface, some degree of geometrical distortion is inevitably introduced and various methods have been developed for minimizing these distortions.

Unfolding a surface onto a 2-D plane is equivalent to parametrizing the surface with a 2-D coordinate system. Low-distortion surface parametrization is widely studied in computer graphics and, in addition to surface unfolding, it has many applications such as multi-resolution analysis and re-meshing (Eck et al., 1995; Lee et al., 1998; Alliez et al., 2003), surface fitting (Floater, 1997), digital geometry processing (Guskov et al., 1999), texture synthesis (Turk, 2001), and mesh compression (Gu et al., 2002). To minimize the distortions, many popular methods use a set of forces that aim to preserve local distances, areas, and angles (Maillot et al., 1993; Carman et al., 1995; Eck et al., 1995; Floater, 1997). Another possibility is to minimize the changes in vertex distances (Schwartz et al., 1989; Wolfson and Schwartz, 1989; Fischl et al., 1999) that are induced by the unfolding. The use of conformal maps have been also investigated (Hurdal et al., 1999; Haker et al., 2000; Lévy et al., 2002; Gu and Yau, 2003) in applications where minimization of the angle distortion is important.

5.1.3 Volume rendering

Volume rendering, also known as direct volume rendering (DVR), is performed using the volumetric 3-D data as such. Typically, no preprocessing is required although additional steps, such as segmentation and tissue classification, can be used for removing certain parts or for otherwise enhancing the data (Fig. 5.4). To reveal the structures in the volume, voxels have different opacities which can depend directly on the voxel values or be separately assigned by a preprocessing step. Some methods limit the voxels to be either totally opaque or transparent, while others allow different levels of transparency.




Figure 5.4: Visualizations of the brain surface from T_1 -weighted MR images. Same segmented volume is rendered using ray-casting (left) and hardware-assisted 3-D texture mapping (right) techniques. *Images created with software written by the author.*

The four most common (Meissner et al., 2000) volume rendering techniques are ray-casting (Tuy and Tuy, 1984; Levoy, 1988), shear-warp factorization (Lacroute and Levoy, 1994), splatting (Westover, 1990), and hardware-assisted 3-D texture mapping (Cabral et al., 1994). Ray-casting has been widely studied and it has collected the largest number of publications over the years (Meissner et al., 2000). Virtual rays are traced from each pixel in the view plane (*i.e.* camera) into the 3-D volume and their interaction with the data is computed on a voxel-by-voxel basis. Different space subdivisions, such as octrees (Levoy, 1990) and binary-space-partitioning (BSP) trees, have been employed to enhance the speed of the ray-casting methods.

The shear-warp factorization is a further enhancement to ray-casting and it is recognized as the fastest software rendering method (Meissner et al., 2000). A clever encoding scheme for the volume and image data is employed that shears the 3-D volume for ray-casting. The resulting rays are always perpendicular to the volume slices which simplifies the calculations. A final warping step transforms the volume-parallel image into the final image.

In contrast to ray-casting, the splatting algorithm is a forward mapping technique that processes through the 3-D volume and maps the voxels to the rendered image in back to front order. Voxels are represented by overlapping basis functions, typically Gaussian kernels, with amplitudes scaled by the voxel values. These basis functions are consequently projected onto the screen and efficient rasterization is achieved with precomputed footprint look-up tables.

Lately, the increase in the capabilities and on-board memories of modern consumer-level PC graphics adapters have made 3-D texture-mapping a very attractive and widely available (Rezk-Salama et al., 2000) volume-rendering technique. The method off-loads the majority of the rendering effort from the central processing unit (CPU) to a dedicated GPU specialized in such computations. A further benefit is the easy incorporation of other solid objects and standard graphics processing methods, such as clipping planes (Fig. 5.5).

Nowadays, the large GPU memories allow multi-channel volumes to be employed in visualization. For example, intensity gradients can be computed from the original volume and the magnitudes of these gradients can be encoded as the second channel.




Figure 5.5: Example of a combined volume and surface-based rendering. T_1 -weighted MRI volume is visualized by hardware-assisted 3-D texture mapping employing two cut planes. Bone structures, visualized as solid objects with standard triangle meshes, are cut open with a set of three cut planes. *Image created with software written by the author.*

The borders between different tissue types are represented as intensity changes and are easily identified from this second channel. Such a method does not require a separate segmentation step as the voxel opacities can be interactively and efficiently changed with 2-D color-table look-ups during the visualization. The tissue boundaries (Fig. 5.6) are especially easy to visualize by setting voxels with a minimal gradient value to transparent.

5.1.4 Texture-mapped triangle meshes (P3)

Publication P3 presents improvements to the surface-based rendering methods employing triangle meshes. The suggested method is especially suited for visualizing the brain surface from T_1 -weighted MRI data although it can be also applied to other surfaces. The method combines the benefits of volume rendering and surface-based rendering and produces high-quality images with excellent speed.

The introduced method uses well-known 2-D texture mapping techniques (Blinn and Newell, 1976; Heckbert, 1986) to add surface details on the rendered triangle meshes. Such techniques have not been previously employed in medical imaging as proper texture images are not readily present. One novel idea in publication P3 is to lend volume rendering techniques, namely depth integration (Bomans et al., 1990), for the texture image generation. The method thereby provides accurate surface texture images from the same 3-D volume that is used for extracting the triangle mesh.

Publication P3 pays careful attention to minimizing the distortions when the sur-




Figure 5.6: Example of tissue boundaries, skin and bone, visualized from a CT volume using hardware-assisted 3-D texture mapping. Image gradient magnitudes are augmented as a second channel to allow easy identification of the surfaces. *The picture was created with the Simian (Kniss, 2006) software using the techniques by Kniss et al. (2001, 2002).*




Figure 5.7: Brain surface visualizations using ray-casting (a), hardware-assisted 3-D texture mapping (b), solid-colored triangle mesh (c), and 2-D texture-mapped triangle mesh (d). *Adapted from publication P3. Images created with software written by the author.*




Figure 5.8: Example combining two 2-D texture-mapped surfaces (skin and brain), one non-textured surface (skull), and volume rendering (in gray). Cut planes are employed to reveal inside structures. *Image created with software written by the author.*

face mesh is flattened on the 2-D texture images and to minimizing possible sampling artifacts present in the textures. The novel idea of measuring the distortions in terms of data resolution rather than mesh geometry is presented. Furthermore, a new error measure, especially suitable for minimizing distortions with respect to data resolution, is introduced.

The proposed method benefits from the texture-mapped surfaces in many ways. The texture images are capable of presenting surface details more accurately than would be practical by increasing the accuracy of surface tessellation. Typically, the triangle size can be increased, producing far less triangles for the surface model. Furthermore, segmentation and triangulation into sulci is no longer necessary for visualization of the cortical mantle, which makes the surface extraction significantly easier. The above effects typically reduce the number of triangles from around 100,000 for human cortex to a mere 10,000–20,000 with textures applied. Such a reduction in triangle count naturally speeds-up the rendering process.

Fig. 5.7 illustrates the visual quality of the proposed method in comparison with volume rendering and solid-colored triangle meshes. The experiments in publication P3 show that such texture-mapped meshes can be rendered many times faster (approx. 7-fold) than a solid-colored mesh with corresponding surface details. Comparing with the volume rendering methods, the difference in speed is around 100-fold in favor of the textured meshes. The method is not limited to brain surfaces alone and can easily incorporate multiple textured objects and employ other standard rendering techniques, such as cut planes and stencil buffers. Fig. 5.8 employs the method for brain and skin surfaces with multiple cut planes and volume rendering. The speed of the proposed method allows real-time interactive manipulation of such complex scenes. Fig. 5.9 shows the 2-D texture images computed for the surfaces of Fig. 5.8.




Figure 5.9: 2-D texture images created for the surface models of Fig. 5.8.




Figure 5.10: Prolate ($\lambda_1 \gg \lambda_2 \geq \lambda_3$), oblate ($\lambda_1 \approx \lambda_2 \gg \lambda_3$), and isotropic ($\lambda_1 \approx \lambda_2 \approx \lambda_3$) tensors.

5.2 Tensor volumes

In medical imaging, tensor volumes can be produced by diffusion tensor imaging (see Section 2.2) and also as a result of complex image analysis. For example, the displacement vector fields in image registration can be analyzed over multiple subjects and the distribution between individual displacements and the average displacement can be summarized as a covariance tensor (Kindlmann et al., 2004).

Such tensor volumes contain a symmetric 3×3 tensor (matrix) in each voxel and many different attributes, such as fractional anisotropy (FA), relative anisotropy (RA), and volume ratio (VR), can be computed from them (Basser and Pierpaoli, 1996; Pierpaoli and Basser, 1996; Uluğ and van Zijl, 1999). These attributes produce scalar volumes that can be visualized with the methods described above, and the subsections below concentrate on methods visualizing the tensor data.

5.2.1 Direct tensor visualization

The tensors can be decomposed into 3 eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ with corresponding eigenvectors \mathbf{e}_1 , \mathbf{e}_2 , and \mathbf{e}_3 . The eigenvectors specify three orthogonal axes such that the diffusion or variance, for example, is the largest along \mathbf{e}_1 and smallest along \mathbf{e}_3 . The strength of the effect (*e.g.* diffusion or variance) along each axis is specified by the respective eigenvalue. Fig. 5.10 shows the three main shapes of the tensors as ellipsoids.

The tensor data can be visualized using the eigenvalue decomposition. In the simplest form, only the direction of the strongest effect, *i.e.* the principal direction \mathbf{e}_1 , is shown. In Fig. 5.11 (a,b), the principal directions are visualized using directionally encoded color (DEC) maps (Douek et al., 1991; Pajevic and Pierpaoli, 1999) that represent the 3-D direction with colors. For example, in Fig. 5.11 (a,b), red color is used for encoding left–right direction, green for anterior–posterior direction, and blue for superior–inferior direction. The principal directions can also be visualized directly as vectors (line segments) showing the local tensor orientation (Fig. 5.11 c,e,f).

Although the visualizations of the principal axis do not reveal anything about the other two components, they are very informative for prolate tensors. Further data, such as FA or other derived attributes, can be used for modifying the brightness and transparency of the DEC map pixels or the color, visibility (transparency), and size of the vector representation. However, if $\lambda_1 \approx \lambda_2 \gg \lambda_3$ (oblate tensor), the principal direction becomes arbitrary due to noise and could just as well be any direction in the




Figure 5.11: Examples of tensor visualization using DEC maps (a,b), vectors for principal direction (c), tractography (d), and combination of vectors and tractography (e,f). Panels (c) and (d) use simulated depth-of-field effect to add sense of depth and to remove visual clutter. Panels (c), (e), and (f) show vectors for tensors with $FA > 0.4$, and panels (e) and (f) further constrain them into 1 cm thick volume over the transaxial cut plane. A sagittal cut plane in panels (e) and (f) separates the visualizations in the two hemispheres. Panels (a) and (b) were created with *DtiStudio* software (Jiang et al., 2006) and provided by courtesy of Jaana Hiltunen. Panels (c)–(f) were created with software written by the author. Data acquired at 3 T at the AMI Centre of TKK.




Figure 5.12: Example of tensors at the anterior part of the corpus callosum visualized as ellipsoids. White rectangle on the FA image (right) shows the location of the magnified portion on the left. Visualization created using *dTV* (Masutani, 2005a) and *Volume-One* (Masutani, 2005b) software employing methods by Masutani et al. (2003). Data acquired at 3 T at the AMI Centre of TKK. Images by courtesy of Jaana Hiltunen.

plane spanned by \mathbf{e}_1 and \mathbf{e}_2 . In these and isotropic cases, the visualization of the full tensor is necessary.

The entire tensor can be visualized using glyphs that illustrate the direction and the tensor size in 3-D. The three eigenvectors specify a local coordinate frame for the 3-D glyph object and the eigenvalues determine the shape and size along these axes. Typically, ellipsoids (Basser et al., 1994b,a; Pierpaoli et al., 1996) are employed as the glyphs (Fig. 5.10 and 5.12). However, the exact 3-D shape of an ellipsoid is sometimes difficult to view in 2-D pictures, especially if the view cannot be interactively rotated. Other glyph objects and shapes, such as super-quadric glyphs (Kindlmann et al., 2004), have been suggested to overcome this problem.

5.2.2 Tractography

A tensor volume can also be visualized indirectly by proper modeling of the effects producing those tensors. For example, diffusion anisotropy in the white matter of the brain is known to roughly be caused by bundles of axonal fibers running in parallel (Le Bihan and van Zijl, 2002). Such fiber tracts can be delineated and extracted from the tensor volume by tractography (Fig. 5.11 d,e,f).

DTI tractography has been intensively studied in recent years (Conturo et al., 1999; Jones et al., 1999; Mori et al., 1999; Westin et al., 1999, 2002; Basser et al., 2000; Poupon et al., 2001; Tench et al., 2002; Parker et al., 2002, 2003; Masutani et al., 2003) and many different tract reconstruction techniques exist. Tract tracing is most often started from a user-selected region of interest (ROI), although the whole tensor volume can also serve as the ROI. Tracing stops when selected termination criteria are fulfilled. Typically, the value of the FA is employed and the tracing is terminated when the FA drops below a predetermined limit, for example 0.2. Another usual termination criterion is the high local tract curvature, *i.e.* large change in the tract angle.

In addition to a single ROI, multiple ROIs can also be employed. One of the ROIs serves as the launch site for the fibers and the others serve as filters. Extraneous fibers

are removed by keeping only those that connect from the launch site to one of the filtering ROIs. Such elimination removes clutter from the visualization and allows the study of possible connections between two specified brain regions.

For tracing the fiber tracts, the principal direction \mathbf{e}_1 of the tensor is assumed to indicate the average fiber bundle direction in a voxel. Some of the methods (*e.g.* Mori et al., 1999; Conturo et al., 1999) employ only this principal direction for the tract propagation. Mori et al. (1999) compute the principal direction for each voxel and an entering tract follows that direction through the whole voxel. A new direction is assigned only at the border of the next voxel. On the other hand, tensor interpolation (*e.g.* Conturo et al., 1999) produces smoother paths and offers better tracing accuracy, especially if the resolution is low with respect to tract curvature. The tract is followed in short steps, for example 0.5 mm, and a new direction is obtained from the interpolated tensor at each step.

Even if an imaged voxel contains only a single fiber bundle, its principal direction will be disturbed by imaging noise. The tensorlines techniques (Weinstein et al., 1999; Westin et al., 1999, 2002) utilize the full information of the diffusion tensor for more accurate tracing. The current tract direction \mathbf{v}_i is modulated by the tensor \mathbf{D} to produce a new direction $\mathbf{v}_{i+1} = \mathbf{D}\mathbf{v}_i$. In effect, the vector \mathbf{v}_i is divided into components along the eigenvectors of \mathbf{D} and these components are modulated by the respective eigenvalues. Therefore, if $\lambda_1 \gg \lambda_2 \geq \lambda_3$ (prolate tensor), the vector \mathbf{v}_{i+1} will be almost parallel to the principal direction \mathbf{e}_1 as the components along \mathbf{e}_2 and \mathbf{e}_3 are suppressed. On the other hand, if $\lambda_1 \approx \lambda_2 \gg \lambda_3$ (oblate tensor), the direction in the plane spanned by \mathbf{e}_1 and \mathbf{e}_2 will be almost unchanged and only the orthogonal component along \mathbf{e}_3 is suppressed.

The tensorlines technique is more robust than the techniques using only the principal direction. Even in the prolate case, the principal direction of a tensor is easily perturbed due to image noise. Furthermore, fiber crossings cause problems as the diffusion tensor is typically oblate in those areas and the principal direction is not well defined. In such cases, the tensorlines technique allows extrapolation of the old direction.

The line propagation methods above produce a tract from a given seed point, but do not provide any information about the sensitivity to perturbations (*e.g.* noise or seed location). Naturally, the tracing can be repeated with small variations of the seed point location to assess the spread caused by such effects. Probabilistic methods (Koch et al., 2002; Behrens et al., 2003; Parker and Alexander, 2003, 2005) utilize probability density functions to describe the local uncertainty of the tract orientation at each spatial location and can therefore provide probability estimates for anatomical connections. Another approach employs regional energy minimization (minimal tract bending) to select the most probable trajectory from several candidates (Poupon et al., 2000).

Nevertheless, partial volume effects (*i.e.* multiple tissue types or fiber bundles in a voxel), image noise, and fiber branching still pose challenges to all tractography techniques. However, as concluded by Mori and van Zijl (2002), even the simple tractography methodologies are able to visualize major white matter connections in animals and humans, and thus DTI and tractography provide valuable *in vivo* information about brain connectivity on an individual level.

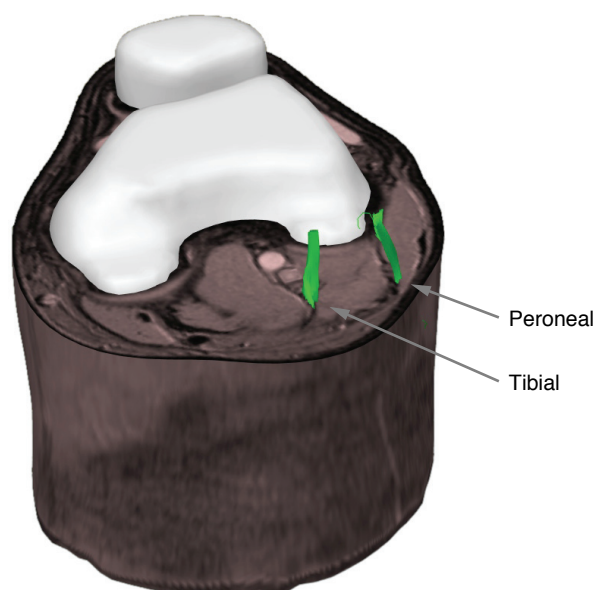


Figure 5.13: Visualization of tibial and peroneal nerves of right knee region, adapted from publication P4. Data acquired at 3 T at the AMI Centre of TKK. Image created with software written by the author.

5.2.3 Tractography of distal peripheral nerves (P4)

Publication P4 presents the first DTI and tractography results of human distal peripheral nerves. Diffusion in the median, ulnar, and radial nerves of the upper limb and the tibial and peroneal nerves of the lower limb was quantified using the apparent diffusion coefficient and FA. These nerves were subsequently delineated and visualized with tractography (see Fig. 5.13).

Our results showed that peripheral nerves are extremely difficult to delineate from standard T_1 - and T_2 -weighted anatomical MRIs. The nerves and the surrounding tissues have very similar intensities making thresholding and 3-D region growing approaches unsuccessful even in high-resolution volumes. However, the diffusion tensors contain information about the nerve direction (*i.e.* diffusion direction) which is lacking from the anatomical scalar images. Tractography can be considered as a special form of 3-D region growing allowing the delineation to follow the diffusion direction only. Consequently, DTI and tractography are better at distinguishing the peripheral nerves from the surrounding tissues.

As shown in publication P4, tractography delineated successfully the majority of the studied peripheral nerves. A total of 21 nerves were traced: 3 lower limb nerves in 3 subjects and 3 upper limb nerves in 4 subjects. Only in one of the subjects, the radial nerve running very superficially at the wrist was not detected. Anatomical, FA, and DEC images were used to manually position the ROIs for tract tracing. FA thresholds and tract lengths were used to distinguish nerves from other nearby tissues. Muscle fibers and ligaments in the joint areas sometimes showed high anisotropy but produced only short tract bundles that were easy to separate from the long nerve tracts.

The maximum FA values in the nerves tended to covary with the nerve size, being smaller for smaller nerves. This change in FA could reflect the more compact tissue structure in the larger nerves. On the other hand, partial volume effects have a larger influence on small nerves, thereby also lowering the measured FA value. In any case,

the maximum FA value seems to be an indicator for the peripheral nerve size.

The main analysis in Publication P4 used dTV (Masutani, 2005a) and Volume-One (Masutani, 2005b) software available on-line. The tractography algorithm of this software employs the principal direction technique (Masutani et al., 2003) and is therefore more susceptible to noise. To increase the reliability and weight of the results, tractography for one subject was also performed with our own software (Fig. 5.13) employing tensor interpolation and the tensorlines tracing algorithm (Westin et al., 1999, 2002). This software used the original diffusion-weighted images and computed the tensors and the eigenvalue decompositions independently from the other software, thereby also validating these processing stages. The tracts were visualized in combination with the techniques developed in Publication P3 and validated the results of the other software.

Chapter 6

Discussion and Conclusions

In this Thesis, three novel enhancements of the three key stages of multi-modal image fusion were presented. These enhancements allow more accurate image registration, better quality image resampling, and fast high-quality visualization. In addition, an application of the diffusion tensor imaging to peripheral nerve tractography successfully delineated these nerves.

The proposed improvement to the mutual-information-based image registration (P1) allows sub-sample registration accuracy even under worst-case conditions. With the currently used sampling schemes, the MI registration is known to grid-lock and produce registration bias when the matched images have equal pixel size and when the image axes align. Under such conditions, a maximum bias of 0.5 pixels and an average bias of 0.2–0.3 pixels produced by the current methods is reduced by the proposed method to a maximum bias of 0.2 pixels and an average bias of 0.04 pixels. The computation time increases with the more accurate method 2–4 fold.

The sub-sample registration accuracy is important for medical data-sets with relatively large pixel sizes, especially when the registered images are used for further analysis; good examples are the functional magnetic resonance images and the diffusion tensor images. Both fMRI and DTI typically employ resolutions of a few millimeters and comprise large sets of serially imaged volumes registered together for statistical analysis (fMRI) or for tensor estimation (DTI). The identical resolution in the consecutive volumes and the possible alignment of the image axes during registration constitute the worst-case conditions for the MI metric. This metric would otherwise be ideal for the fMRI and DTI data since it makes only minimal assumptions about the correspondence between pixel intensities. Especially in diffusion-weighted images, pixel intensities vary greatly depending on the applied gradient direction. In addition, activation differences in fMRI volumes are known to mislead some registration metrics. Thus, the proposed improvements to the worst-case MI registration accuracy are very important.

The proposed improvement to the image resampling (P2) combines Fourier-space up-sampling with conventional convolution-based interpolation. The Fourier-space stage employs efficiently the theoretically-perfect full-width sinc interpolation for image up-sampling. The convolution stage allows the limitations of the first stage to be circumvented and this combination leads to a very high-quality resampling method with better SNR than the other currently employed methods (*e.g.* OMOMS and B-splines). The proposed method is especially suitable for fMRI and DTI volumes.

The minimization of the resampling errors is necessary when the resampled images are used for further analysis. For example, fMRI volumes are registered to remove subject motion and are subsequently resampled before statistical analysis. High-quality registration and resampling guarantee that the analyzed correlations are actually due

to brain activation and not produced by image registration or resampling errors. The proposed two-stage resampling method is also especially fit for image registration applications which resample one of the original images while iteratively searching for the best match. In such applications, the first up-sampling stage can employ higher magnification and the second stage can use fast small-support interpolation. For the repeated resampling during the iterative search, only the fast second stage needs to be re-computed, thereby providing high-quality results with high speed.

The introduced enhancement for the triangle-mesh rendering (P3) combines the benefits of the surface-based visualization and the volume visualization. The applied standard 2-D texture mapping, optimized with respect to data resolution, provides accurate surface details and allows the use of bigger triangles, thereby reducing their count. Consequently, the suggested method provides 5–100 fold increase in the speed to produce equal or even better quality results than the previously used methods.

The high-quality high-speed visualization is important to allow interactive manipulation of the rendered scene. Certain special techniques, such as depth-of-field simulation and concave cuts, require multiple rendering passes for a single scene. Typically, the scene contains several complex models, and thus a single model has to be rendered as quickly as possible. Fast interactive manipulation of the scene, even a simple rotation of the view, allows better interpretation of the 3-D spatial relationships.

The application of DTI to peripheral nerves (P4) illustrates that the diffusion imaging, typically used for the white matter of the brain, can be successfully applied to delineating distal peripheral nerves. The maximum of the fractional anisotropy measure was also shown to covary with the nerve size. The DTI could provide useful information for the diagnosis and follow-up of nerve lesions, entrapments, and regeneration. Furthermore, the peripheral nerves could serve as 'living phantoms' for testing and validating different tractography methods as the course of these nerves can be readily verified from known anatomy.

The fields of multi-modal neuroimaging and medical imaging in general are rapidly developing and full of potential for useful applications. High-quality image processing steps, such as those presented in this Thesis, are crucial for accurate image fusion, analysis, and visualization.

Topics of future research

Although the proposed improvements to the MI-based image registration significantly reduced the sampling-related artifacts, the remaining maximum worst-case bias of 0.2 pixels warrants further research. Naturally, this maximum bias can be reduced by increasing the kernel size, but that will happen at the expense of computation time. The present kernels are derived from the nearest-neighbor and the linear kernels, and thus fall into the B-spline family of kernels. Therefore, some other type of kernel in support sizes 2–4 could perhaps further improve the worst-case accuracy.

The proposed two-stage resampling method was tested with the first-stage magnification factors up to 8 using the 2-D data. Curiously, factor 2 was already capable of producing the best results, with no further benefit with the increased magnification. With higher factors, all versions employing different second-stage interpolators approached an upper limit, already set by the best versions with magnification factor 2. Since a change from the floating-point calculations to double-precision calculations had no effect, this limit was not caused by the numerical round-off errors and seemed to be image-specific. Therefore, it will be interesting to see if any future resampling

method with any support size could clearly exceed this limit.

To minimize the possibility of visible texture discontinuities, the proposed visualization method was designed to minimize the number of parts that the mesh was divided into. However, these discontinuities turned out to be invisible to the eye with the typical texture generation parameters. The current design limits the method to meshes that are topologically spheres and it also causes larger distortions with surfaces having high local curvature. Since discontinuities are not an issue, a dynamic method could be developed that divides the surface into any required number of parts to circumvent the present problems.

The current trend in computer science is to increase the number of cores in a CPU. Although the speed of a single core is also gradually increasing, the real performance boost comes from parallel processing with multi-thread programming. The vast majority of the current image processing algorithms and software are designed for serial processing and the next challenge in the field is to utilize this parallel computation power in the best possible way.

Bibliography

- Alliez, P., Cohen-Steiner, D., Devillers, O., Lévy, B., and Desbrun, M. (2003). Anisotropic polygonal remeshing. *ACM Transactions on Graphics*, 22: 485–493.
- Alpert, N. M., Berdichevsky, D., Levin, Z., Morris, E. D., and Fischman, A. J. (1996). Improved methods for image registration. *NeuroImage*, 3: 10–18.
- Andersson, J. L. R., Sundin, A., and Valind, S. (1995). A method for coregistration of PET and MR brain images. *The Journal of Nuclear Medicine*, 36: 1307–1315.
- Arun, K. S., Huang, T. S., and Blostein, S. D. (1987). Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9: 698–700.
- Ashburner, J. and Friston, K. J. (1999). Nonlinear spatial normalization using basis functions. *Human Brain Mapping*, 7: 254–266.
- Basser, P. J., Mattiello, J., and Le Bihan, D. (1994a). Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance, Series B*, 103: 247–254.
- Basser, P. J., Mattiello, J., and Le Bihan, D. (1994b). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66: 259–267.
- Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., and Aldroubi, A. (2000). In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44: 625–632.
- Basser, P. J. and Pierpaoli, C. (1996). Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance, Series B*, 111: 209–219.
- Behrens, T. E. J., Woolrich, M. W., Jenkinson, M., Johansen-Berg, H., Nunes, R. G., Clare, S., Matthews, P. M., Brady, J. M., and Smith, S. M. (2003). Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50: 1077–1088.
- Besl, P. J. and McKay, N. D. (1992). A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14: 239–256.
- Blinn, J. F. and Newell, M. E. (1976). Texture and reflection in computer generated images. *Communications of the ACM*, 19: 542–547.
- Bloch, F. (1946). Nuclear induction. *Physical Review*, 70: 460–474.
- Blu, T., Thévenaz, P., and Unser, M. (1999). Generalized interpolation: Higher quality at no additional cost. In *Proceedings of the 1999 IEEE International Conference on Image Processing (ICIP'99)*, Vol. III, pp. 667–671, Kobe, Japan.

- Blu, T., Thévenaz, P., and Unser, M. (2001). MOMS: Maximal-order interpolation of minimal support. *IEEE Transactions on Image Processing*, 10: 1069–1080.
- Blu, T., Thévenaz, P., and Unser, M. (2004). Linear interpolation revitalized. *IEEE Transactions on Image Processing*, 13: 710–719.
- Bomans, M., Höhne, K.-H., Tiede, U., and Riemer, M. (1990). 3-D segmentation of MR images of the head for 3-D display. *IEEE Transactions on Medical Imaging*, 9: 177–183.
- 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.
- Borgefors, G. (1988). Hierarchical chamfer matching: A parametric edge matching algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10: 849–865.
- Bracewell, R. N. (1986). *The Fourier Transform and Its Applications*. McGraw-Hill, New York.
- Bro-Nielsen, M. and Gramkow, C. (1996). Fast fluid registration of medical images. In Höhne, K. H. and Kikinis, R., editors, *Visualization in Biomedical Computing*, Vol. 1131 of *Lecture Notes in Computer Science*, pp. 265–276. Springer-Verlag, Berlin, Germany.
- Cabral, B., Cam, N., and Foran, J. (1994). Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In Kaufman, A. and Krüger, W., editors, *Proceedings of Volume Visualization '94*, pp. 91–98, New York, USA. ACM Press.
- Carman, G. J., Drury, H. A., and Van Essen, D. C. (1995). Computational methods for reconstructing and unfolding the cerebral cortex. *Cerebral Cortex*, 5: 506–517.
- Catmull, E. and Smith, A. R. (1980). 3-D transformations of images in scanline order. In *Proceedings of SIGGRAPH '80*, pp. 279–285, New York, USA. ACM Press.
- Chen, H. and Varshney, P. K. (2003). Mutual information-based CT–MR brain image registration using generalized partial volume joint histogram estimation. *IEEE Transactions on Medical Imaging*, 22: 1111–1119.
- Christensen, G. E., Joshi, S. C., and Miller, M. I. (1997). Volumetric transformation of brain anatomy. *IEEE Transactions on Medical Imaging*, 16: 864–877.
- Christensen, G. E., Rabbitt, R. D., and Miller, M. I. (1994). 3d brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39: 609–618.
- Christensen, G. E., Rabbitt, R. D., and Miller, M. I. (1996). Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5: 1435–1447.
- Cohen, D. (1972). Magnetoencephalography: Detection of the brain's electrical activity with a superconducting magnetometer. *Science*, 175: 664–666.

- Collignon, A., Maes, F., Delaere, D., Vandermeulen, D., Suetens, P., and Marchal, G. (1995). Automated multi-modality image registration based on information theory. In Bizais, Y., Barillot, C., and Di Paola, R., editors, *Information Processing in Medical Imaging*, pp. 263–274, Dordrecht, The Netherlands. Kluwer.
- Conturo, T. E., Lori, N. F., Cull, T. S., Akbudak, E., Snyder, A. Z., Shimony, J. S., McKinstry, R. C., Burton, H., and Raichle, M. E. (1999). Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences of the United States of America*, 96: 10422–10427.
- Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley-Interscience, New York, NY, USA.
- Di Salle, F., Formisano, E., Linden, D. E. J., Goebel, R., Bonavita, S., Pepino, A., Smaltino, F., and Tedeschi, G. (1999). Exploring brain function with magnetic resonance imaging. *European Journal of Radiology*, 30: 84–94.
- Douek, P., Turner, R., Pekar, J., Patronas, N., and Le Bihan, D. (1991). MR color mapping of myelin fiber orientation. *Journal of Computer Assisted Tomography*, 15: 923–929.
- Du, Y. P., Parker, D. L., Davis, W. L., and Cao, G. (1994). Reduction of partial-volume artifacts with zero-filled interpolation in three-dimensional MR angiography. *Journal of Magnetic Resonance Imaging*, 4: 733–741.
- Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., and Stuetzle, W. (1995). Multiresolution analysis of arbitrary meshes. In *Proceedings of SIGGRAPH '95*, pp. 173–182, New York, USA. ACM Press.
- Fischl, B., Sereno, M. I., and Dale, A. M. (1999). Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, 9: 195–207.
- Floater, M. S. (1997). Parametrization and smooth approximation of surface triangulations. *Computer Aided Geometric Design*, 14: 231–250.
- Formisano, E. and Goebel, R. (2003). Tracking cognitive processes with functional MRI mental chronometry. *Current Opinion in Neurobiology*, 13: 174–181.
- Freeman, R. (1998). *Spin Choreography, Basic Steps in High Resolution NMR*. Oxford University Press, Oxford.
- Friston, K. J., Ashburner, J., Frith, C. D., Poline, J.-B., Heather, J. D., and Frackowiak, R. S. J. (1995). Spatial registration and normalization of images. *Human Brain Mapping*, 3: 165–189.
- Gaens, T., Maes, F., Vandermeulen, D., and Suetens, P. (1998). Non-rigid multimodal image registration using mutual information. In Wells, W. M., Colchester, A., and Delp, S., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI'98*, Vol. 1496 of *Lecture Notes in Computer Science*, pp. 1099–1106. Springer-Verlag, Berlin, Germany.
- Gonzalez, R. C. and Woods, R. E. (1992). *Digital Image Processing*. Addison Wesley, Reading, Massachusetts.

- Goshtasby, A., Turner, D. A., and Ackerman, L. V. (1992). Matching of tomographic slices for interpolation. *IEEE Transactions on Medical Imaging*, 11: 507–516.
- Grevera, G. J. and Udupa, J. K. (1996). Shape-based interpolation of multidimensional grey-level images. *IEEE Transactions on Medical Imaging*, 15: 881–892.
- Gu, X., Gortler, S. J., and Hoppe, H. (2002). Geometry images. In *Proceedings of SIGGRAPH '02*, pp. 355–361, New York, USA. ACM Press.
- Gu, X. and Yau, S.-T. (2003). Global conformal surface parametrization. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, pp. 127–137, Aire-la-Ville, Switzerland. Eurographics Association.
- Guimond, A., Roche, A., Ayache, N., and Meunier, J. (2001). Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections. *IEEE Transactions on Medical Imaging*, 20: 58–69.
- Guskov, I., Sweldens, W., and Schröder, P. (1999). Multiresolution signal processing for meshes. In *Proceedings of SIGGRAPH '99*, pp. 325–334, New York, USA. ACM Press.
- Hajnal, J. V., Saeed, N., Oatridge, A., Williams, E. J., Young, I. R., and Bydder, G. M. (1995a). Detection of subtle brain changes using subvoxel registration and subtraction of serial MR images. *Journal of Computer Assisted Tomography*, 19: 677–691.
- Hajnal, J. V., Saeed, N., Soar, E. J., Oatridge, A., , Young, I. R., and Bydder, G. M. (1995b). A registration and interpolation procedure for subvoxel matching of serially acquired MR images. *Journal of Computer Assisted Tomography*, 19: 289–296.
- Haker, S., Angenent, S., Tannenbaum, A., Kikinis, R., Sapiro, G., and Halle, M. (2000). Conformal surface parametrization for texture mapping. *IEEE Transactions on Visualization and Computer Graphics*, 6: 181–189.
- Hata, N., Dohi, T., Warfield, S., Wells III, W., Kikinis, R., and Jolesz, F. A. (1998). Multimodality deformable registration of pre- and intraoperative images for MRI-guided brain surgery. In Wells, W. M., Colchester, A., and Delp, S., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI'98*, Vol. 1496 of *Lecture Notes in Computer Science*, pp. 1067–1074. Springer-Verlag, Berlin, Germany.
- Heckbert, P. S. (1986). Survey of texture mapping. *IEEE Computer Graphics and Applications*, 6: 56–67.
- Hellier, P. and Barillot, C. (2000). Multimodal non-rigid warping for correction of distortions in functional MRI. In Delp, S. L., DiGoia, A. M., and Jaramaz, B., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000*, Vol. 1935 of *Lecture Notes in Computer Science*, pp. 512–520. Springer-Verlag, Berlin, Germany.
- Hellier, P., Barillot, C., Mémin, E., and Pérez, P. (1999). Medical image registration with robust multigrid techniques. In Taylor, C. and Colchester, A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI'99*, Vol. 1679 of *Lecture Notes in Computer Science*, pp. 680–688. Springer-Verlag, Berlin, Germany.

- Hermosillo, G. and Faugeras, O. (2001). Dense image matching with global and local statistical criteria: A variational approach. In Jacobs, A. and Baldwin, T., editors, *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2001. CVPR 2001*, Vol. 1, pp. I-73–I-78, Los Alamitos, California, USA. IEEE Computer Society.
- Hill, D. L. G., Batchelor, P. G., Holden, M., and Hawkes, D. J. (2001). Medical image registration. *Physics in Medicine and Biology*, 46: R1–R45.
- Hill, D. L. G., Maurer, Jr., C. R., Studholme, C., Fitzpatrick, J. M., and Hawkes, D. J. (1998). Correcting scaling errors in tomographic images using a nine degree of freedom registration algorithm. *Journal of Computer Assisted Tomography*, 22: 317–323.
- Horsfield, M. A. (1999). Mapping eddy current induced fields for the correction of diffusion-weighted echo planar images. *Magnetic Resonance Imaging*, 17: 1335–1345.
- Hou, H. S. and Andrews, H. C. (1978). Cubic splines for image interpolation and digital filtering. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26: 508–517.
- Hounsfield, G. N. (1980). Computed medical imaging, Nobel lecture, December 8, 1979. *Journal of Computer Assisted Tomography*, 4: 665–674.
- Huettel, S. A., Song, A. W., and McCarthy, G. (2004). *Functional Magnetic Resonance Imaging*. Sinauer Associates, Inc, Sunderland, Massachusetts USA.
- Hurdal, M. K., Bowers, P. L., Stephenson, K., Summers, D. L., Rehm, K., Schaper, K., and Rottenberg, D. A. (1999). Quasi-conformally flat mapping the human cerebellum. In *Proceedings of the Second International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 279–286, London, UK. Springer-Verlag.
- Hämäläinen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., and Lounasmaa, O. V. (1993). Magnetoencephalography – theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics*, 65: 413–497.
- Ji, J. X., Pan, H., and Liang, Z.-P. (2003). Further analysis of interpolation effects in mutual information-based image registration. *IEEE Transactions on Medical Imaging*, 22: 1131–1140.
- Jiang, H., van Zijl, P. C. M., Kim, J., Pearlson, G. D., and Mori, S. (2006). DtiStudio: Resource program for diffusion tensor computation and fiber bundle tracking. *Computer Methods and Programs in Biomedicine*, 81: 106–116.
- Jones, D. K., Simmons, A., Williams, S. C. R., and Horsfield, M. A. (1999). Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine*, 42: 37–41.
- Junck, L., Moen, J. G., Hutchins, G. D., Brown, M. B., and Kuhl, D. E. (1990). Correlation methods for the centering, rotation, and alignment of functional brain images. *The Journal of Nuclear Medicine*, 31: 1220–1276.

- Keys, R. G. (1981). Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29: 1153–1160.
- Kim, B., Boes, J. L., Frey, K. A., and Meyer, C. R. (1997). Mutual information for automated unwarping of rat brain autoradiographs. *NeuroImage*, 5: 31–40.
- Kim, J., Fisher III, J. W., Tsai, A., Wible, C., Willsky, A. S., and Wells III, W. M. (2000). Incorporating spatial priors into an information theoretic approach for fMRI data analysis. In Delp, S. L., DiGoia, A. M., and Jaramaz, B., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000*, Vol. 1935 of *Lecture Notes in Computer Science*, pp. 62–71. Springer-Verlag, Berlin, Germany.
- Kindlmann, G. L., Weinstein, D. M., Lee, A. D., Toga, A. W., and Thompson, P. M. (2004). Visualization of anatomic covariance tensor fields. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Vol. 3, pp. 1842–1845, Piscataway, USA. IEEE Service Center.
- Kjems, U., Strother, S. C., Anderson, J., Law, I., and Hansen, L. K. (1999). Enhancing the multivariate signal of [¹⁵O] water PET studies with a new nonlinear neuroanatomical registration algorithm. *IEEE Transactions on Medical Imaging*, 18: 306–319.
- Kniss, J., Kindlmann, G., and Hansen, C. (2001). Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *Proceedings of Visualization '01*, pp. 255–262, Washington, USA. IEEE Computer Society.
- Kniss, J., Kindlmann, G., and Hansen, C. (2002). Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 8: 270–285.
- Kniss, J. M. (2006). Simian. [Online]. Available: <http://www.cs.utah.edu/~jmk/simian/>.
- Koch, M. A., Norris, D. G., and Hund-Georgiadis, M. (2002). An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage*, 16: 241–250.
- Lacroute, P. and Levoy, M. (1994). Fast volume rendering using a shear-warp factorization of the viewing transformation. *Computer Graphics*, 28: 451–458.
- Lau, Y. H., Braun, M., and Hutton, B. F. (2001). Non-rigid image registration using a median-filtered coarse-to-fine displacement field and a symmetric correlation ratio. *Physics in Medicine and Biology*, 46: 1297–1319.
- Lauritzen, M. (2005). Reading vascular changes in brain imaging: is dendritic calcium the key? *Nature Reviews Neuroscience*, 6: 77–85.
- Lauterbur, P. C. (1973). Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature*, 242: 190–191.
- Le Bihan, D. (1995). Molecular diffusion, tissue microdynamics and microstructure. *NMR in Biomedicine*, 8: 375–386.

- Le Bihan, D. (2003). Looking into the functional architecture of the brain with diffusion MRI. *Nature Reviews Neuroscience*, 4: 469–480.
- Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., and Laval-Jeantet, M. (1986). MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders. *Radiology*, 161: 401–407.
- Le Bihan, D. and van Zijl, P. (2002). From the diffusion coefficient to the diffusion tensor. *NMR in Biomedicine*, 15: 431–434.
- Lee, A. W. F., Sweldens, W., Schröder, P., Cowsar, L., and Dobkin, D. (1998). Maps: Multiresolution adaptive parameterization of surfaces. In *Proceedings of SIGGRAPH '98*, pp. 95–104, New York, USA. ACM Press.
- Lehmann, T. M., Gönner, C., and Spitzer, K. (1999). Survey: Interpolation methods in medical image processing. *IEEE Transactions on Medical Imaging*, 18: 1049–1075.
- Lehmann, T. M., Gönner, C., and Spitzer, K. (2001). Addendum: B-spline interpolation in medical image processing. *IEEE Transactions on Medical Imaging*, 20: 660–665.
- Levoy, M. (1988). Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8: 29–37.
- Levoy, M. (1990). Efficient ray tracing of volume data. *ACM Transactions on Graphics*, 9: 245–261.
- Lévy, B., Petitjean, S., Ray, N., and Maillot, J. (2002). Least squares conformal maps for automatic texture atlas generation. In *Proceedings of SIGGRAPH '02*, pp. 362–371, New York, USA. ACM Press.
- Little, J. A., Hill, D. L. G., and Hawkes, D. J. (1997). Deformations incorporating rigid structures. *Computer Vision and Image Understanding*, 66: 223–232.
- Logothetis, N. K. (2002). The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 357: 1003–1037.
- Lötjönen, J. and Mäkelä, T. (2001). Elastic matching using a deformation sphere. In Niessen, W. J. and Viergever, M. A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*, Vol. 2208 of *Lecture Notes in Computer Science*, pp. 541–548. Springer-Verlag, Berlin, Germany.
- Lötjönen, J., Reissman, P.-J., Magnin, I. E., and Katila, T. (1999). Model extraction from magnetic resonance volume data using the deformable pyramid. *Medical Image Analysis*, 3: 387–406.
- Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., and Suetens, P. (1997). Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16: 187–198.
- Maes, F., Vandermeulen, D., and Suetens, P. (1999). Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information. *Medical Image Analysis*, 3: 373–386.

- Maillot, J., Yahia, H., and Verroust, A. (1993). Interactive texture mapping. In *Proceedings of SIGGRAPH '93*, pp. 27–34, New York, USA. ACM Press.
- Maintz, J. B. A. and Viergever, M. A. (1998). A survey of medical image registration. *Medical Image Analysis*, 2: 1–36.
- Mangin, J.-F., Poupon, C., Clark, C., Le Bihan, D., and Bloch, I. (2001). Eddy-current distortion correction and robust tensor estimation for MR diffusion imaging. In Niessen, W. J. and Viergever, M. A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*, Vol. 2208 of *Lecture Notes in Computer Science*, pp. 186–194. Springer-Verlag, Berlin, Germany.
- Mansfield, P. (1977). Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 10: L55–L58.
- Markel, J. D. (1971). FFT pruning. *IEEE Transactions on Audio and Electroacoustics*, 19: 305–311.
- Masutani, Y. (2005a). dTV: diffusion tensor visualizer. [Online]. Available: <http://www.ut-radiology.umin.jp/people/masutani/dTV.htm>.
- Masutani, Y. (2005b). Volume-One. [Online]. Available: <http://www.volume-one.org/>.
- Masutani, Y., Aoki, S., Abe, O., Hayashi, N., and Otomo, K. (2003). MR diffusion tensor imaging: recent advance and new techniques for diffusion tensor visualization. *European Journal of Radiology*, 46: 53–66.
- Meijering, E. H., Niessen, W. J., and Viergever, M. A. (2001). Quantitative evaluation of convolution-based methods for medical image interpolation. *Medical Image Analysis*, 5: 111–126.
- Meijering, E. H., Zuiderveld, K. J., and Viergever, M. A. (1999). Image reconstruction by convolution with symmetrical piecewise n th-order polynomial kernels. *IEEE Transactions on Image Processing*, 8: 192–201.
- Meissner, M., Huang, J., Bartz, D., Mueller, K., and Crawfis, R. (2000). A practical evaluation of popular volume rendering algorithms. In *Proceedings of the 2000 IEEE Symposium on Volume Visualization*, pp. 81–90, New York, USA. ACM Press.
- Menon, R. S. and Kim, S.-G. (1999). Spatial and temporal limits in cognitive neuroimaging with fMRI. *Trends in Cognitive Sciences*, 3: 207–216.
- Menon, R. S., Luknowsky, D. C., and Gati, J. S. (1998). Mental chronometry using latency-resolved functional MRI. *Proceedings of the National Academy of Sciences of the United States of America*, 95: 10902–10907.
- Merboldt, K.-D., Hänicke, W., and Frahm, J. (1985). Self-diffusion NMR imaging using stimulated echoes. *Journal of Magnetic Resonance*, 64: 479–486.
- Meyer, C. R., Boes, J. L., Kim, B., Bland, P. H., Zasadny, K. R., Kison, P. V., Koral, K., Frey, K. A., and Wahl, R. L. (1997). Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformations. *Medical Image Analysis*, 1: 195–206.

- Mori, S., Crain, B. J., Chacko, V. P., and van Zijl, P. C. M. (1999). Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45: 265–269.
- Mori, S. and van Zijl, P. C. M. (2002). Fiber tracking: principles and strategies - a technical review. *NMR in Biomedicine*, 15: 468–480.
- Nagai, K. (1986). Pruning the decimation-in-time FFT algorithm with frequency shift. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 34: 1008–1010.
- Ogawa, S., Tank, D. W., Menon, R., Ellermann, J. M., Kim, S.-G., Merkle, H., and Ugurbil, K. (1992). Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences of the United States of America*, 89: 5951–5955.
- Pajevic, S. and Pierpaoli, C. (1999). Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42: 526–540.
- Parker, G. J. M. and Alexander, D. C. (2003). Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. In Taylor, C. and Noble, J. A., editors, *Information Processing in Medical Imaging*, Vol. 2732 of *Lecture Notes in Computer Science*, pp. 684–695. Springer-Verlag, Berlin, Germany.
- Parker, G. J. M. and Alexander, D. C. (2005). Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 360: 893–902.
- Parker, G. J. M., Haroon, H. A., and Wheeler-Kingshott, C. A. M. (2003). A framework for a streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. *Journal of Magnetic Resonance Imaging*, 18: 242–254.
- Parker, G. J. M., Wheeler-Kingshott, C. A. M., and Barker, G. J. (2002). Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging*, 21: 505–512.
- Parzen, E. (1962). On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33: 1065–1076.
- Pelizzari, C. A., Chen, G. T. Y., Spelbring, D. R., Weichselbaum, R. R., and Chen, C.-T. (1989). Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. *Journal of Computer Assisted Tomography*, 13: 20–26.
- Pierpaoli, C. and Basser, P. J. (1996). Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36: 893–906.
- Pierpaoli, C., Jezzard, P., Basser, P. J., Barnett, A., and Di Chiro, G. (1996). Diffusion tensor MR imaging of the human brain. *Radiology*, 201: 637–648.

- Pluim, J. P. W., Maintz, J. B. A., and Viergever, M. A. (2000a). Image registration by maximization of combined mutual information and gradient information. *IEEE Transactions on Medical Imaging*, 19: 809–814.
- Pluim, J. P. W., Maintz, J. B. A., and Viergever, M. A. (2000b). Interpolation artefacts in mutual information-based image registration. *Computer Vision and Image Understanding*, 77: 211–232.
- Pluim, J. P. W., Maintz, J. B. A., and Viergever, M. A. (2003). Mutual-information-based registration of medical images: A survey. *IEEE Transactions on Medical Imaging*, 22: 986–1004.
- Poupon, C., Clark, C. A., Frouin, V., Régis, J., Bloch, I., Le Bihan, D., and Mangin, J.-F. (2000). Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage*, 12: 184–195.
- Poupon, C., Mangin, J.-F., Clark, C. A., Frouin, V., Régis, J., Le Bihan, D., and Bloch, I. (2001). Towards inference of human brain connectivity from MR diffusion tensor data. *Medical Image Analysis*, 5: 1–15.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1988). *Numerical Recipes in C*. Cambridge University Press, Cambridge.
- Purcell, E. M., Torrey, H. C., and Pound, R. V. (1946). Resonance absorption by nuclear magnetic moments in a solid. *Physical Review*, 69: 37–38.
- Raichle, M. E. and Mintun, M. A. (2006). Brain work and brain imaging. *Annual Review of Neuroscience*, 29: 449–476.
- Reichenbach, S. E. and Geng, F. (2003). Two-dimensional cubic convolution. *IEEE Transactions on Image Processing*, 12: 857–865.
- Rezk-Salama, C., Engel, K., Bauer, M., Greiner, G., and Ertl, T. (2000). Interactive volume rendering on standard PC graphics hardware using multi-textures and multi-stage rasterization. In Spencer, S. N., editor, *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Workshop on Graphics Hardware*, pp. 109–118, New York, USA. ACM Press.
- Rohde, G. K., Aldroubi, A., and Dawant, B. M. (2003). The adaptive bases algorithm for intensity-based nonrigid image registration. *IEEE Transactions on Medical Imaging*, 22: 1470–1479.
- Rohlfing, T. and Maurer, Jr., C. R. (2001). Intensity-based non-rigid registration using adaptive multilevel free-form deformation with an incompressibility constraint. In Niessen, W. J. and Viergever, M. A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*, Vol. 2208 of *Lecture Notes in Computer Science*, pp. 111–119. Springer-Verlag, Berlin, Germany.
- Rohr, K. (1998). Image registration based on thin-plate splines and local estimates of anisotropic landmark localization uncertainties. In Wells, W. M., Colchester, A., and Delp, S., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI'98*, Vol. 1496 of *Lecture Notes in Computer Science*, pp. 1174–1183. Springer-Verlag, Berlin, Germany.

- Rohr, K., Stiehl, H. S., Sprengel, R., Beil, W., Buzug, T. M., Weese, J., and Kuhn, M. H. (1996). Point-based elastic registration of medical image data using approximating thin-plate splines. In Höhne, K. H. and Kikinis, R., editors, *Visualization in Biomedical Computing*, Vol. 1131 of *Lecture Notes in Computer Science*, pp. 297–306. Springer-Verlag, Berlin, Germany.
- Rohr, K., Stiehl, H.-S., Sprengel, R., Buzug, T. M., Weese, J., and Kuhn, M. H. (2001). Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions on Medical Imaging*, 20: 526–534.
- Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L. G., Leach, M. O., and Hawkes, D. J. (1999). Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18: 712–721.
- Sarrut, D. and Clippe, S. (2001). Geometrical transformation approximation for 2D/3D intensity-based registration of portal images and CT scan. In Niessen, W. J. and Viergever, M. A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*, Vol. 2208 of *Lecture Notes in Computer Science*, pp. 532–540. Springer-Verlag, Berlin, Germany.
- Schnabel, J. A., Rueckert, D., Quist, M., Blackall, J. M., Castellano-Smith, A. D., Hartkens, T., Penney, G. P., Hall, W. A., Liu, H., Truwit, C. L., Gerritsen, F. A., Hill, D. L. G., and Hawkes, D. J. (2001). A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In Niessen, W. J. and Viergever, M. A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*, Vol. 2208 of *Lecture Notes in Computer Science*, pp. 573–581. Springer-Verlag, Berlin, Germany.
- Schormann, T., Henn, S., and Zilles, K. (1996). A new approach to fast elastic alignment with applications to human brains. In Höhne, K. H. and Kikinis, R., editors, *Visualization in Biomedical Computing*, Vol. 1131 of *Lecture Notes in Computer Science*, pp. 337–342. Springer-Verlag, Berlin, Germany.
- Schwartz, E. L., Shaw, A., and Wolfson, E. (1989). A numerical solution to the generalized mapmaker’s problem: Flattening nonconvex polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11: 1005–1008.
- Seppä, M. (2007). High-quality two-stage resampling for 3-D volumes in medical imaging. *Medical Image Analysis*, 11: 346–360.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27: 379–423, 623–656.
- Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, 37: 10–21.
- Skinner, D. P. (1976). Pruning the decimation-in-time FFT algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 24: 193–194.
- Smit, T., Smith, M. R., and Nichols, S. T. (1990). Efficient sinc function interpolation technique for center padded data. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38: 1512–1517.

- Stejskal, E. O. and Tanner, J. E. (1965). Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42: 288–292.
- Studholme, C., Constable, R. T., and Duncan, J. S. (2000). Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion model. *IEEE Transactions on Medical Imaging*, 19: 1115–1127.
- Studholme, C., Hill, D. L. G., and Hawkes, D. J. (1996a). Automated 3-D registration of MR and CT images of the head. *Medical Image Analysis*, 1: 163–175.
- Studholme, C., Hill, D. L. G., and Hawkes, D. J. (1996b). Incorporating connected region labelling into automated image registration using mutual information. In Amini, A. A., Bookstein, F. L., and Wilson, D. C., editors, *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 23–31, Los Alamitos, CA. IEEE Computer Society Press.
- Studholme, C., Hill, D. L. G., and Hawkes, D. J. (1997). Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Medical Physics*, 24: 25–35.
- Studholme, C., Hill, D. L. G., and Hawkes, D. J. (1999). An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32: 71–86.
- Studholme, C., Novotny, E., Zupal, I. G., and Duncan, J. S. (2001). Estimating tissue deformation between functional images induced by intracranial electrode implantation using anatomical MRI. *NeuroImage*, 13: 561–576.
- Tanner, C., Schnabel, J. A., Chung, D., Clarkson, M. J., Rueckert, D., Hill, D. L. G., and Hawkes, D. J. (2000). Volume and shape preservation of enhancing lesions when applying non-rigid registration to a time series of contrast enhancing MR breast images. In Delp, S. L., DiGoia, A. M., and Jaramaz, B., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000*, Vol. 1935 of *Lecture Notes in Computer Science*, pp. 327–337. Springer-Verlag, Berlin, Germany.
- Taylor, D. G. and Bushell, M. C. (1985). The spatial mapping of translational diffusion coefficients by the NMR imaging technique. *Physics in Medicine and Biology*, 30: 345–349.
- Tench, C. R., Morgan, P. S., Blumhardt, L. D., and Constantinescu, C. (2002). Improved white matter fiber tracking using stochastic labeling. *Magnetic Resonance in Medicine*, 48: 677–683.
- Thévenaz, P., Blu, T., and Unser, M. (2000). Interpolation revisited. *IEEE Transactions on Medical Imaging*, 19: 739–758.
- Thévenaz, P. and Unser, M. (2000). Optimization of mutual information for multiresolution image registration. *IEEE Transactions on Image Processing*, 9: 2083–2099.
- Thirion, J.-P. (19981). Image matching as a diffusion process: an analogy with Maxwell’s demons. *Medical Image Analysis*, 2: 243–260.

- Thulborn, K. R., Waterton, J. C., Matthews, P. M., and Radda, G. K. (1982). Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochimica et Biophysica Acta*, 714: 265–270.
- Tsao, J. (2003). Interpolation artifacts in multimodality image registration based on maximization of mutual information. *IEEE Transactions on Medical Imaging*, 22: 854–864.
- Tsuchida, N., Yamada, Y., and Ueda, M. (1987). Hardware for image rotation by twice skew transformations. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 35: 527–532.
- Turk, G. (2001). Texture synthesis on surfaces. In *Proceedings of SIGGRAPH '01*, pp. 347–354, New York, USA. ACM Press.
- Turner, R. and Ordidge, R. J. (2000). Technical challenges of functional magnetic resonance imaging: The biophysics and technology behind a reliable neuroscientific tool for mapping the human brain. *IEEE Engineering in Medicine and Biology Magazine*, 19: 42–54.
- Tuy, H. K. and Tuy, L. T. (1984). Direct 2-D display of 3-D objects. *IEEE Computer Graphics and Applications*, 4: 29–33.
- Uluğ, A. M. and van Zijl, P. C. M. (1999). Orientation-independent diffusion imaging without tensor diagonalization: Anisotropy definitions based on physical attributes of the diffusion ellipsoid. *Journal of Magnetic Resonance Imaging*, 9: 804–813.
- Unser, M. (1999). Splines: A perfect fit for signal and image processing. *IEEE Signal Processing Magazine*, 16: 22–38.
- Unser, M., Aldroubi, A., and Eden, M. (1993a). B-spline signal processing: Part I-theory. *IEEE Transactions on Signal Processing*, 41: 821–833.
- Unser, M., Aldroubi, A., and Eden, M. (1993b). B-spline signal processing: Part II-efficient design and applications. *IEEE Transactions on Signal Processing*, 41: 834–848.
- Unser, M., Thévenaz, P., and Yaroslavsky, L. (1995). Convolution-based interpolation for fast, high-quality rotation of images. *IEEE Transactions on Image Processing*, 4: 1371–1381.
- van den Elsen, P. A., Maintz, J. B. A., Pol, E.-J. D., and Viergever, M. A. (1995). Automatic registration of CT and MR brain images using correlation of geometrical features. *IEEE Transactions on Medical Imaging*, 14: 384–396.
- Viola, P. and Wells, III, W. M. (1995). Alignment by maximization of mutual information. In *Proceedings of the 5th International Conference on Computer Vision '95*, pp. 16–23.
- Viola, P. and Wells, III, W. M. (1997). Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24: 137–154.

- Wang, Y. and Staib, L. H. (1998). Elastic model based non-rigid registration incorporating statistical shape information. In Wells, W. M., Colchester, A., and Delp, S., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI'98*, Vol. 1496 of *Lecture Notes in Computer Science*, pp. 1162–1173. Springer-Verlag, Berlin, Germany.
- Wansapura, J. P., Holland, S. K., Dunn, R. S., and Ball, W. S. (1999). NMR relaxation times in the human brain at 3.0 tesla. *Journal of Magnetic Resonance Imaging*, 9: 531–538.
- Warnking, J., Dojat, M., Guérin-Dugué, A., Delon-Martin, C., Olympieff, S., Richard, N., Chéhikian, A., and Segebarth, C. (2002). fMRI retinotopic mapping – step by step. *NeuroImage*, 17: 1665–1683.
- Weinstein, D., Kindlmann, G., and Lundberg, E. (1999). Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *Proceedings of Visualization '99*, pp. 249–253.
- Wells, W. M., Viola, P., Atsumi, H., Nakajima, S., and Kikinis, R. (1996). Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1: 35–51.
- Westin, C.-F., Maier, S. E., Khidhir, B., Everett, P., Jolesz, F. A., and Kikinis, R. (1999). Image processing for diffusion tensor magnetic resonance imaging. In Taylor, C. and Colchester, A., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI'99*, Vol. 1679 of *Lecture Notes in Computer Science*, pp. 441–453. Springer-Verlag, Berlin, Germany.
- Westin, C.-F., Maier, S. E., Mamata, H., Nabavi, A., Jolesz, F. A., and Kikinis, R. (2002). Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6: 93–108.
- Westover, L. (1990). Footprint evaluation for volume rendering. *Computer Graphics*, 24: 367–376.
- Wolfson, E. and Schwartz, E. L. (1989). Computing minimal distances on polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11: 1001–1005.
- Woods, R. P., Cherry, S. R., and Mazziotta, J. C. (1992). Rapid automated algorithm for aligning and reslicing PET images. *Journal of Computer Assisted Tomography*, 16: 620–633.
- Woods, R. P., Grafton, S. T., Holmes, C. J., Cherry, S. R., and Mazziotta, J. C. (1998). Automated image registration: I. General methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography*, 22: 139–152.
- Yaroslavsky, L. P. (1997). Efficient algorithm for discrete sinc interpolation. *Applied Optics*, 36: 460–463.
- Zhu, Y.-M. and Cochoff, S. M. (2002). Influence of implementation parameters on registration of MR and SPECT brain images by maximization of mutual information. *The Journal of Nuclear Medicine*, 43: 160–166.