

# Intrasubject Registration for Change Analysis in Medical Imaging

Marius Staring

## Colophon

This thesis was typeset by the author using L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub>.

Cover design: Marius Staring and Proefschriftmaken.nl

Intrasubject Registration for Change Analysis in Medical Imaging

M. Staring - Utrecht, University Medical Center Utrecht

PhD thesis Utrecht University - with a summary in Dutch

Copyright © 2008 by M. Staring. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission in writing from the author.

Financial support by ImagO and the Röntgen Stichting Utrecht for the publication of this thesis is gratefully acknowledged. A grant for publication of this thesis was kindly provided by Nucletron B.V. and Philips Medical Systems.

ISBN 978-90-8891-063-0

Printed by Proefschriftmaken.nl, Oisterwijk, The Netherlands

# Intrasubject Registration for Change Analysis in Medical Imaging

## Intrasubject Registratie voor het Analyseren van Veranderingen in Medische Beelden

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor  
aan de Universiteit Utrecht op gezag van de  
rector magnificus, prof. dr. J.C. Stoof,  
ingevolge het besluit van het college voor promoties  
in het openbaar te verdedigen op  
donderdag 9 oktober 2008 des middags te 2.30 uur

door

**Marinus Staring**

geboren op 14 maart 1977 te Den Helder

Promotor: Prof. dr. ir. M.A. Viergever

Co-promotor: Dr. J.P.W. Pluim

---

This work was carried out at the Image Sciences Institute, University Medical Center Utrecht (Utrecht, The Netherlands), under the auspices of ImagO, the Utrecht Graduate School for Biomedical Image Sciences. The project was financially supported by the Netherlands Organisation for Scientific Research (NWO), project number 612.065.205.





# Contents

<b>Colophon</b>	<b>ii</b>
<b>1 General Introduction</b>	<b>1</b>
<b>2 elastix: A Toolbox for Medical Image Registration</b>	<b>5</b>
2.1 Introduction . . . . .	6
2.2 Image Registration with <code>elastix</code> . . . . .	7
2.3 Experiments and Results . . . . .	12
2.4 Conclusion . . . . .	16
<b>3 Tissue-Dependent Filtering of the Deformation Field</b>	<b>17</b>
3.1 Introduction . . . . .	18
3.2 Method . . . . .	19
3.3 Experiments and Results . . . . .	23
3.4 Discussion and Conclusions . . . . .	31
<b>4 A Rigidity Penalty Term for Nonrigid Registration</b>	<b>33</b>
4.1 Introduction . . . . .	34
4.2 Method . . . . .	35
4.3 Experiments and Results . . . . .	40
4.4 Conclusions and Discussion . . . . .	49
<b>5 Image Subtraction for Change Detection in Ground-Glass Opacities in Chest CT</b>	<b>53</b>
5.1 Introduction . . . . .	54
5.2 Materials and Methods . . . . .	55
5.3 Results . . . . .	59
5.4 Discussion . . . . .	61
<b>6 Registration of Cervical MRI Using Multifeature Mutual Information</b>	<b>63</b>
6.1 Introduction . . . . .	64
6.2 Method . . . . .	66
6.3 Features . . . . .	69
6.4 Experiments . . . . .	70
6.5 Results . . . . .	73
6.6 Discussion and Conclusion . . . . .	75

<b>7 Summary</b>	<b>79</b>
7.1 Summary . . . . .	79
7.2 Clinical Relevance . . . . .	81
<b>References</b>	<b>93</b>
<b>Samenvatting</b>	<b>95</b>
<b>Publications</b>	<b>99</b>
<b>Dankwoord</b>	<b>103</b>
<b>Curriculum Vitae</b>	<b>105</b>



# General Introduction

# 1

**I**MAGE REGISTRATION is the process of aligning images, in other words the sequence of steps (algorithm) that attempts to overlay the structures visualised by the data. In this thesis we consider data acquired in the clinic. An example of a 2D image modality is x-ray. Nowadays, most acquisition protocols are 3D, from which Computed Tomography (CT) and Magnetic Resonance (MR) are well-known. Other examples are Single Photon Emission Computed Tomography (SPECT) and Positron Emission Tomography (PET). Some acquisition techniques are not only capable of visualising anatomy, but also function-related information, such as metabolism.

Further analysis of the acquired data often requires (nonrigid) image registration. The technique can be used for several applications, of which we give some examples:

**Combining information** Image modalities provide a wealth of information about the subject of interest. Combining this information often yields additional clinical insight not apparent in the separate images. An example is shown in Figure 1.1. One image modality provides detailed information about the anatomy (CT), the other modality (PET) provides functional information. To combine both information sources, the spatial relationship between the multi-modal images has to be found.

**Follow-up** Often, a single modality is used to follow a patient over time. A scan is taken every couple of months or every year. This is done to analyse the progress of a certain disease. An example is the growth or shrinkage of a cancer. Physicians not only need to know disease progress, but also the success of treatment. Did radiation therapy indeed shrink a prostate tumour? In this case we need to find the spatial relationship between images from the same modality.

**Contrast-enhancement** Also data sets acquired within minutes often need image registration. This is for example the case when vessels are analysed. Contrast material is injected in a vessel and a sequence of scans is made, each depicting part of the vasculature. Because of e.g. respiratory motion or heart beats the scans are not spatially aligned. Examples of images generated with the use of contrast material are given in Figure 4.3.

In the next chapter, more information can be found about the specific registration framework used in this thesis.

As indicated by the title of this thesis, we focus on data sets originating from the same patient: intrasubject registration. We concentrate on the use of image registration




Figure 1.1: Combining CT and PET data. The CT scan (top row) provides information about the patients anatomy. PET has a higher diagnostic accuracy for staging lung tumours compared to CT, but has a much lower spatial resolution. The PET scan (bottom row) clearly reveals a lung tumour. Image courtesy [1].

for the analysis of change in medical images. This area of application of image alignment demand some special requirements from the registration algorithm. A naive registration algorithm may have the following drawbacks:

1. All tissue is treated the same way, regardless of the possibly different properties. Specifically, rigid and nonrigid tissue are both deformed while performing a nonrigid registration. This way the progress of a disease may be concealed.

An example is shown in Figure 1.2. A patient suffering from a ground-glass opacity (a type of lung nodule, that may be benign or malignant) is followed over time using chest CT. Two scans, separated by a few months or a year, need to be compared to assess whether the nodule has grown. Growth would indicate malignancy, upon which clinical action is taken (surgery). If the nodule did not grow it may be a benign nodule, which for now does not require clinical intervention. Change can be visualised by subtracting registered scans. However, common nonrigid registration algorithms conceal the growth, as shown in Figure 1.2(a) and (c). When we make a distinction between the nodule and the other tissue, it is possible to visualise the growth. In this case, the clear rim in Figure 1.2(d) indicates malignancy and clinical action is required. Details can be found in Chapter 4 and 5.

2. In some registration problems, using only the image intensities may not be sufficient for obtaining a satisfactory alignment. This hinders the analysis of change.

The registration of MR images of women with cervical cancer is an example of a challenging registration problem. The size and position of organs around the cervix




Figure 1.2: Visualising nodule growth. Figures (a) and (c) show the results when all tissue is treated the same way, and (b) and (d) when the nodule is kept rigid, but all other tissue is deformed. (a) and (b) show the deformation field superimposed on the registration result. The subtraction image after registration is shown in (c) and (d). Note the undesired compression of the ground-glass opacity (located in the center of the image) in (a), resulting in a less clear rim in (c).

can vary considerably, which influences the position of the tumour. For success of radiation therapy it is crucial to know the exact location of the tumour, as well as the location of some surrounding organs. The tumour needs to receive a prescribed radiation dose, while the surrounding organs need to be spared as much as possible. Registration can be used to track the relevant organs, to study their motion over time, and to detect when positions have changed. This will make it possible to reduce the treatment margins, thereby reducing complications during and after therapy. As will be shown in Chapter 6, a standard algorithm using image intensity only did not give satisfactory registration results. An extension is proposed which uses multiple image features, and improves registration performance.

This thesis addresses the two problems given above.

## Outline of this Thesis

**Chapter 2** gives a general overview of image registration and describes the registration framework and software used in this thesis. The registration software was developed at the Image Sciences Institute jointly with Stefan Klein for our research projects. All experiments described in this thesis were performed with this software package, called `elastix`.

A first approach to tackle the rigid-nonrigid tissue problem uses adaptive filtering of the deformation field. The technique is reported in **Chapter 3**. Issues that remain with this approach are solved with another technique based on a local rigidity penalty term, see **Chapter 4**. The last technique is clinically employed for the detection of subtle changes in ground-glass opacities in the lung. The results of this study are presented in **Chapter 5**.

To address the problem of insufficient registration quality, in **Chapter 6** the registration cost function is extended with multiple image features, instead of using image intensity only, as is commonly done.

## Acknowledgements

This research was funded by the Netherlands Organisation for Scientific Research (NWO). This work also benefited from the use of the Insight Segmentation and Registration Toolkit (ITK), an open source software package developed as an initiative of the U.S. National Library of Medicine and available at <http://www.itk.org>.

# elastix: A Toolbox for Medical Image Registration

# 2

## Abstract

---

Medical image registration is an important task in medical image processing. It refers to the process of aligning data sets, possibly from different modalities (e.g., magnetic resonance (MR) and computed tomography (CT)), different time points (e.g., follow-up scans), and/or different subjects (in case of population studies). A large number of methods for image registration are described in the literature. Unfortunately, there is not one method that works for all applications. This chapter presents **elastix**: a publicly available computer program for intensity-based medical image registration. The software consists of a collection of algorithms that are commonly used to solve medical image registration problems. The modular design of **elastix** allows the user to quickly configure, test, and compare different registration methods for a specific application. The command-line interface enables automated processing of large numbers of data sets, by means of scripting. The usage of **elastix** for comparing different registration methods is illustrated with three example experiments, in which individual components of the registration method are varied.

---

Based upon: S. Klein<sup>a</sup>, M. Staring<sup>a</sup>, K. Murphy, M.A. Viergever, and J.P.W. Pluim, **elastix**: a toolbox for intensity-based medical image registration, *submitted*.

---

<sup>a</sup>In alphabetical order. Both authors contributed equally.

## 2.1 Introduction

IMAGE REGISTRATION is a frequently used technique in medical image processing. It is the task of finding the spatial relationship between two or more images. Areas of application include the alignment of data sets from different modalities [2], comparison of follow-up scans to a base-line scan [3] (see Chapter 4), alignment of pre- and postcontrast images [3–5], updating treatment plans for radiotherapy and surgery [6, 7], atlas-based segmentation [8, 9], creating models of anatomy [10], and aligning training images for classification [11, 12].

In registration, one image, which is called the *moving image*  $I_M(\mathbf{x})$ , is deformed to fit the other image, the *fixed image*  $I_F(\mathbf{x})$ . In other words, registration is the problem of finding a *coordinate transformation*  $\mathbf{T}(\mathbf{x})$  that makes  $I_M(\mathbf{T}(\mathbf{x}))$  spatially aligned with  $I_F(\mathbf{x})$ . The quality of alignment is defined by a cost function  $\mathcal{C}(\mathbf{T}; I_F, I_M)$ . The optimal coordinate transformation is estimated by minimising the cost function with respect to  $\mathbf{T}$ , usually by means of an iterative optimisation method embedded in a hierarchical (multiresolution) scheme. Extensive reviews on the subject of image registration are given in [13–17].

Application of an image registration method requires many choices to be made, such as the the optimisation method, the multiresolution strategy, the method of image interpolation to evaluate  $I_M(\mathbf{T}(\mathbf{x}))$ , the coordinate transformation model, and the definition of the cost function. Several possibilities for the optimisation method are discussed in [18, 19]. An overview of multiresolution strategies is given in [15]. Various image interpolation methods are compared in [20]. The degrees of freedom of the coordinate transformation  $\mathbf{T}$  determine the types of deformations that can be recovered. Whereas in many applications it may be sufficient to consider only rigid transformations (global translations and rotations) [21, 22], frequently a more flexible transformation model is needed, allowing for local deformations [2, 4, 23, 24]. For the cost function  $\mathcal{C}$  many options have been proposed in the literature. Commonly used intensity-based cost functions are the mean squared difference (MSD) [25, 26], normalised correlation (NC) [27, 28], mutual information (MI) [21, 29–31], and normalised mutual information (NMI) [5, 8, 32]. Sometimes, a regularisation term is added to the cost function, in order to penalise undesired deformations [3, 5, 24] (see Chapter 4). In medical image processing research it is often necessary to compare several options for each of the registration components. Given the large number of choices, this can be a tedious procedure.

This chapter presents a software package for medical image registration: `elastix`. The `elastix` software has a modular design, including several optimisation methods, multiresolution schemes, interpolators, transformation models, and cost functions. This allows the user to quickly compare different registration methods, in order to select a satisfactory configuration for a specific application. `elastix` has a command-line interface, which enables automated processing of large numbers of data sets, by means of scripting. The software is built upon a widely used open source library for medical image processing, the Insight Toolkit (ITK) [33].

In Section 2.2, the general registration framework of `elastix` is discussed, key features of the software are presented, and an overview of the available registration components




Figure 2.1: The basic registration components. The scheme is an extended version of the scheme introduced in [33].

is given. In Section 2.3, three examples of applications that can be handled with the software are given. In these experiments, the influence of three important registration components is demonstrated.

## 2.2 Image Registration with elastix

### 2.2.1 Registration Framework

Mathematically, the registration problem is formulated as an optimisation problem in which the cost function  $\mathcal{C}$  is minimised with respect to  $\mathbf{T}$ . The `elastix` software is based on a parametric approach, meaning that the number of possible transformations is limited by introducing a parameterisation of the transformation. The optimisation problem reads:

$$\hat{\mu} = \arg \min_{\mu} \mathcal{C}(\mathbf{T}_{\mu}; I_F, I_M), \quad (2.1)$$

where the subscript  $\mu$  indicates that the transform has been parameterised. The vector  $\mu$  contains the transformation parameters. The reader is referred to [17, 24] for an overview on nonparametric methods. The minimisation problem (2.1) is solved with an iterative optimisation method, usually in a multiresolution setting. A schematic overview of the basic registration components and their relations is given in Figure 2.1, which is a slightly extended version of the scheme introduced in [33].

### 2.2.2 Software Characteristics

The `elastix` software is structured according to the block scheme of Figure 2.1. For each component (transform, cost function, etc.) several choices are available. The user can configure a registration algorithm by specifying the names of the desired components in a parameter text file. Additional settings that some components may require can also be entered in this parameter file. Fixed and moving image file names are supplied as

command-line arguments, so that multiple image pairs can be registered using the same parameter settings.

All output of the registration, such as the deformed moving image  $I_M(\mathbf{T}_{\hat{\mu}}(\mathbf{x}))$  and intermediate progress information, is saved to disk. It is often necessary to apply the resulting transformation  $\mathbf{T}_{\hat{\mu}}$  to data sets other than the moving image. For example, in atlas-based segmentation methods [8, 9] the transformation is applied to a segmentation (label image) of the moving image. To that end, `elastix` outputs a text file that describes the transformation  $\mathbf{T}_{\hat{\mu}}$ . This text file can subsequently be passed to an accompanying program, called `transformix`, together with the image to be deformed. This program can also be used to evaluate the transformation at user-defined points, or to generate the deformation field.

A large part of the `elastix` code is based on the ITK [33]. The use of the ITK implies that the low-level functionality (image classes, memory allocation etc.) is thoroughly tested. Naturally, all image formats supported by the ITK are supported by `elastix` as well. The source code can be compiled on multiple operating systems, using various compilers, and supports both 32 and 64 bit systems. Executables and source code are publicly available from the website <http://elastix.isi.uu.nl>, under the BSD license. A manual for `elastix` and an example of usage can also be downloaded. The manual includes an example parameter file, describes in detail the various options that can be specified, and provides recommendations for image registration.

### 2.2.3 Registration Components

In the following subsections, more information is given about each component of the block scheme in Figure 2.1.

#### Cost Function

The cost function  $\mathcal{C}$  measures the similarity between the fixed image and the deformed moving image. An example is the MSD:

$$\text{MSD}(\mathbf{T}_{\mu}; I_F, I_M) = \frac{1}{N} \sum_{\mathbf{x} \in \Omega_F} (I_F(\mathbf{x}) - I_M(\mathbf{T}_{\mu}(\mathbf{x})))^2, \quad (2.2)$$

where  $\Omega_F$  denotes the fixed image domain, and  $N$  the number of voxels  $\mathbf{x}$  sampled from the fixed image domain. The sampler, which is responsible for selecting the samples  $\mathbf{x}$ , is discussed in more detail in Section 2.2.3.

The following metrics are currently supported by `elastix`: MSD, NC, MI, NMI, and the  $\kappa$ -statistic. MSD is only suited for two images with equal intensities, i.e. for images from the same modality. NC is less strict, it assumes an affine relation between the intensity values of the fixed and moving image. MI and NMI assume only a statistical relation between the intensities of the images. They are therefore suited not only for monomodal, but also for multimodal image pairs. The  $\kappa$ -statistic can be used for registering binary images. It measures the overlap of objects in the images. For each of the metrics above, a localised version can be constructed, as explained in [9], by selecting the appropriate sampler. This is described in Section 2.2.3.



Parameters such as the number of bins of the joint histogram, needed for MI and NMI, can be set in the aforementioned parameter file.

When a nonrigid transformation model is used, a regularisation term that penalises undesired deformations can be added to the cost function. An example is the incompressibility constraint described by [5], which penalises compression and expansion of structures. Other examples of regularisation terms are the bending energy of a thin plate [4] and the rigidity penalty term [3] (see Chapter 4). `elastix` supports these constraints, but currently in combination with MI only.

## Transformation

The parameterisation of the coordinate transformation  $T_{\boldsymbol{\mu}}$  determines the degrees of freedom of the deformation. An example is the affine transformation model, which allows for translation, rotation, scaling and skew of the images:

$$T_{\boldsymbol{\mu}}(\mathbf{x}) = A\mathbf{x} + \mathbf{t}, \quad (2.3)$$

where  $A$  is a matrix and  $\mathbf{t}$  represents the translation vector. The parameter vector  $\boldsymbol{\mu}$  is formed by the matrix elements  $a_{ij}$  and the translation vector. In 2D, this gives a vector of length 6:  $\boldsymbol{\mu} = (a_{11}, a_{12}, a_{21}, a_{22}, t_x, t_y)^T$ . In 3D,  $\boldsymbol{\mu}$  consists of 9 matrix elements and 3 translations.

The following transformation models are currently supported by `elastix`: translation, rigid (translation and rotation), similarity (rigid plus isotropic scaling), affine, and nonrigid. In the literature, several nonrigid transformation models have been proposed [4, 34, 35], each having its own advantages and disadvantages. In `elastix` a B-spline representation [4] has been implemented. The transformation is modelled as a weighted sum of B-spline basis functions, placed on a uniform control point grid. The B-spline basis functions have local support [36], which is beneficial for fast computation. The flexibility of the deformation is defined by the resolution of the control point grid, which has to be supplied by the user via the parameter file. Section 2.3.1 demonstrates the effect of using different transforms for an example application.

Frequently, nonrigid registration must be preceded by a rigid or affine registration, in order to achieve a rough initial alignment. `elastix` supports the concatenation of any sequence of transforms. The user may also supply an initial transformation, determined in advance, for example by manually clicking corresponding points.

## Optimisation

To solve (2.1), an iterative optimisation procedure is employed. In every iteration  $k$ , the current transformation parameters  $\boldsymbol{\mu}_k$  are updated by taking a step in the search direction  $\mathbf{d}_k$ :

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - a_k \mathbf{d}_k, \quad (2.4)$$

with  $a_k$  a scalar that determines the step size. A wide range of optimisation methods can be formulated in this way, each having different definitions of  $a_k$  and  $\mathbf{d}_k$  [19]. A common

choice for the search direction is the derivative of the cost function  $\partial\mathcal{C}/\partial\boldsymbol{\mu}$  evaluated at the current position  $\boldsymbol{\mu}_k$ . In this case, Eq. (2.4) boils down to a gradient descent method.

`elastix` includes all optimisation methods described in [19]: gradient descent, quasi-Newton, nonlinear conjugate gradient (several variants), evolution strategy, and a number of stochastic gradient descent methods (Kiefer-Wolfowitz, Robbins-Monro, and simultaneous perturbation). Also an exhaustive search routine is included, which is mainly useful for examining the cost function, as demonstrated in Section 2.3.2. The experimental results in [19] indicate that a stochastic gradient descent method (Robbins-Monro) is a good choice for many applications. It reduces the computation time per iteration by using only a small subset of the fixed image’s voxels for computing the cost function derivative. In each iteration, new samples must be selected randomly. This can be realised in `elastix` by selecting an appropriate sampler, which is explained in the next subsection.

## Sampling Strategies

To compute the cost function  $\mathcal{C}$  (and its derivative  $\partial\mathcal{C}/\partial\boldsymbol{\mu}$ ) a set of samples  $\boldsymbol{x} \in \Omega_F$  needs to be selected, as in (2.2). The sampler component in Figure 2.1 is responsible for this. The most straightforward strategy is to use all voxels from the fixed image, which has as an obvious downside that it is time-consuming for large images. A common approach is to use a subset of voxels, selected on a uniform grid, or sampled randomly. Another strategy is to pick only those points that are located on striking image features, such as edges.

`elastix` currently supports the use of all voxels, a subset of voxels selected on a uniform grid, random sampling of voxels, and random sampling off the voxel grid (at non-voxel locations). Random sampling off the grid has been shown to improve the smoothness of the cost function [37, 38]. In Section 2.3.2, we demonstrate this effect by comparing several sampling schemes on an example application. For all sampling strategies discussed above, the user may optionally supply a mask image, indicating regions of interest. In this way, one can force the sampler to pick only points near edges in the image, for example.

With random sampling, the `elastix` user can enforce the selection of new samples in every iteration  $k$  of the optimisation process. In this way, the stochastic optimisation methods described in [19] can be realised. The localised mutual information strategy, presented in [9], can be implemented by letting the sampler pick points in a small neighbourhood. A new neighbourhood is randomly selected in every iteration of the optimisation procedure.

## Interpolation

For computation of the cost function, the value  $I_M(\mathbf{T}_\boldsymbol{\mu}(\boldsymbol{x}))$  is evaluated at non-voxel positions, for which intensity interpolation is needed. Several methods for interpolation exist, varying in quality and speed, including nearest neighbour, linear and  $N$ -th order B-spline interpolation [36, 39]. `elastix` supports all interpolators mentioned above.




Figure 2.2: Two multiresolution strategies using a Gaussian pyramid ( $\sigma = 8.0, 4.0, 2.0$  voxels). The top row shows multiresolution with downsampling, the bottom row without. Note that in the top row the number of voxels in each dimension is halved every resolution, but the voxel size is doubled, so physically the images are of the same size.

## Hierarchical Strategies

Hierarchical (multiresolution) strategies are an important aspect of image registration. For an extensive overview, the reader is referred to [15]. `elastix` implements several hierarchical strategies.

The pyramid components in the block scheme of Figure 2.1 represent the multiresolution schemes for the *image data*. Two types of image pyramids are available in `elastix`: Gaussian pyramids with and without downsampling. Figure 2.2 illustrates the difference.

The second important multiresolution strategy, not apparent from Figure 2.1, is the gradual increase of *transformation model* complexity. During nonrigid registration, a hierarchical effect can be realised by starting with a coarse B-spline control point resolution and gradually refining the grid in subsequent resolutions, thereby introducing the capability to recover more fine-scale deformations.

More generally, any parameter setting can be subjected to a hierarchical strategy in `elastix`. For example, the number of joint histogram bins that is used for computing MI and NMI could be gradually increased, as was suggested in [30].

In Section 2.3.3 several multiresolution strategies are compared for the nonrigid registration of CT chest scans.

## 2.3 Experiments and Results

In this section, some applications of `elastix` are described, to illustrate its convenience for configuring, testing, and comparing different registration methods. Three key components of registration were studied: the transformation model in Section 2.3.1, the sampling technique in Section 2.3.2, and the multiresolution strategy in Section 2.3.3. The experiments demonstrate the impact these components can have on the registration results, and therefore stress the importance of a proper configuration for the application at hand.

### 2.3.1 Transformation Models

The effect of the type of transformation was investigated by comparing the registration performance of several transformation models.

To this end, a set of 50 clinical MR scans of the prostate was used, all originating from different patients. The scans were made by the Department of Radiotherapy of the University Medical Center Utrecht, as part of prostate cancer treatment planning. They were acquired on a Philips 3T scanner (Gyrosan NT Intera, Philips Medical Systems, Best, The Netherlands) using a balanced steady-state free precession sequence with fat suppression. The scans had a dimension of  $512 \times 512 \times 90$  voxels of size  $0.49 \times 0.49 \times 1.0$  mm.

Fifty interpatient registrations were performed by registering each MR scan with its predecessor in the 50-scan series. Interpatient registration of these scans is needed for atlas-based segmentation of the prostate [9]. The registration problem is challenging, since the anatomical variability between subjects is large. Also, the data suffer from several artefacts, as shown in [9]. For our experiments we used the same settings as in [9], with localised MI as a cost function (see Section 2.2.3), and a four-level Gaussian image pyramid with downsampling. The following transformation models were compared: translation, rigid, affine, and B-spline with different control point spacings: 64, 32, 16, 8, and 4 mm. The result of the registration with translations only was used as an initialisation for all other registrations. For the B-spline registrations, the control point grid was subjected to a multiresolution scheme: registration starts with a coarse control point resolution; with smoother versions of the images, the control point resolution is increased accordingly.

For all images a manual segmentation of the prostate was available, made by an experienced radiation oncologist. After registration with `elastix`, the transformation  $T_{\hat{\mu}}$  was applied to the prostate segmentation of the moving image, using `transformix`. The overlap with the segmentation of the fixed image was computed, using the Dice similarity coefficient (DSC) [40]:

$$\text{DSC}(X, Y) = \frac{2|X \cap Y|}{|X| + |Y|}, \quad (2.5)$$

where  $X$  and  $Y$  represent the two segmentations, and  $|\cdot|$  denotes the number of voxels within the segmentation. A DSC of 1 indicates perfect registration. A value of 0 means that the prostates had no overlap at all after registration.




Figure 2.3: The effect of the transformation model on the accuracy of registration, measured by the prostate overlap. The abbreviation BS- $\langle sp \rangle$  refers to a B-spline transformation with control point spacing  $sp$ , in mm.

The results are presented in Figure 2.3. For each transformation model, the DSC values of the 50 MR scans were summarised by a box-and-whiskers plot. A paired, two-sided Wilcoxon test was used to assess the median differences between adjacent columns. A star on top of a column indicates a significant difference ( $p < 0.05$ ) with respect to the previous column. The graph clearly shows that a nonrigid registration was essential in this application. The best results were obtained using a B-spline control point spacing of 8 mm. With this setting, the computation time was around 15 minutes per registration on a single processor Pentium 2.8 GHz personal computer.

### 2.3.2 Sampling Strategies

The grid effect is a well-known issue in image registration. It refers to the problem that the cost function contains irregularities at locations representing grid-aligning transformations, which can impede the registration process. It has often been studied in the context of interpolation artefacts [20]. In this section it is demonstrated that the sampling mechanism can solve this issue, by taking samples off the voxel grid, as suggested in [37, 38].

Brain images were taken from the “Retrospective Image Registration Evaluation” project [22]. We investigated the registration of a T1-weighted MR image (moving image)




Figure 2.4: The effect of different sampling strategies on the smoothness of the cost function; a) translation in the  $z$  direction, b) translation in the  $z$  direction after downsampling the MR image in the  $z$  direction, c) translation in the  $x$  direction.

to a PET image (both of patient 001). The PET image had a dimension of  $128 \times 128 \times 15$  voxels of size  $2.59 \times 2.59 \times 8.0$  mm. The MR image had a dimension of  $256 \times 256 \times 26$  voxels of size  $1.25 \times 1.25 \times 4.0$  mm.

The cost function (MI) was analysed using an exhaustive search in a single translation direction, with a step size of 0.1 mm. Linear interpolation was used to compute  $I_M(\mathbf{T}_\mu(\mathbf{x}))$ . Different sampling strategies were employed for computing the cost function: all voxels, random sampling at the voxel grid, and random sampling off the voxel grid.

In Figure 2.4(a) the cost function  $-\text{MI}(\mathbf{T}_\mu; I_F, I_M)$  is plotted as a function of the translation  $t_z$ , the direction with the largest voxel spacing. The two samplers that take samples on the voxel grid have a very irregular cost function. The irregularities show a pattern, related to the voxel sizes of the images in the  $z$  direction (8 mm for PET, 4 mm for MR). Every 8 mm a slice of the PET image maps outside the MR image. This causes the large discontinuities at  $t_z = 12$  mm and 20 mm for example. Every 4 mm, the cost function exhibits a small local maximum, caused by the aligning voxel grids of the images. The random sampler that takes samples off the grid clearly leads to a much smoother cost function.

The experiment was repeated after downsampling the MR image by a factor of 2 in the  $z$  direction. The voxel size of the MR image thus became equal to that of the PET image (8 mm). Figure 2.4(b) shows the cost function as a function of  $t_z$ . The irregularities follow a single pattern in this graph, with a peak at every 8 mm.

Figure 2.4(c) shows the result for translation in the  $x$  direction (obtained using the original non-downsampled MR image). The portion of voxels of the PET image that move simultaneously outside the MR image domain is smaller than in the  $z$  direction. Consequently, the grid effect is reduced. The cost function appears much smoother, also for the two samplers that take samples on the voxel grid, although small irregularities remain visible at multiples of the voxel spacing. This example shows that, in practice, it may not always be strictly necessary to sample voxels off the grid.

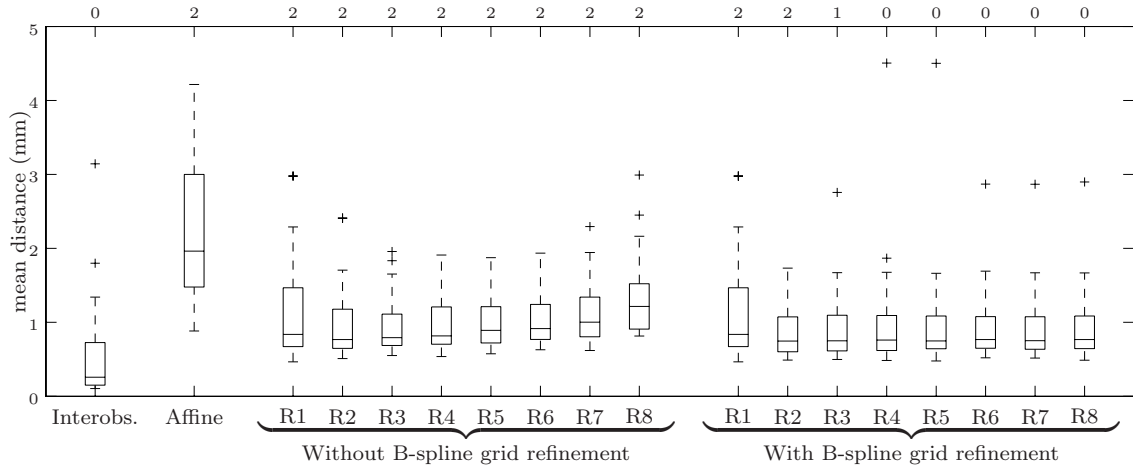


Figure 2.5: The effect of the multiresolution strategy on registration accuracy, expressed as the mean distance between corresponding points. ‘R1’ - ‘R8’ refer to the number of image resolution levels that was used. The numbers on top of the graph refer to the number of outliers with mean distance larger than 5 mm.

### 2.3.3 Multiresolution Strategies

The influence of the choice of multiresolution strategy is examined in this section. CT chest scans of 26 patients were taken from a lung cancer screening trial [41]. Each patient had a baseline and a follow-up scan, acquired 3-9 months apart. The scans were obtained at full inspiration and without contrast injection on a 16-detector-row scanner (Mx8000 IDT or Brilliance 16P, Philips Medical Systems, Best, The Netherlands). Images were of size  $512 \times 512$  in-plane, with the voxel size ranging from  $0.55 \times 0.55$  mm to  $0.8 \times 0.8$  mm. The number of slices varied from 383 to 529, with slice thickness 1 mm and slice spacing 0.7 mm. The images were downsampled by a factor of two in each dimension before registration, in order to decrease computational load.

For each patient the baseline and follow-up scans were registered using a nonrigid B-spline transformation. An affine registration was used for initialisation. The registration was performed with a Gaussian image pyramid (without downsampling) using  $R \in \{1, \dots, 8\}$  levels. Two experiments were performed for each value of  $R$ . Firstly, the resolution of the B-spline control point grid was kept at a constant value of 12 mm (isotropic) in all resolutions. Secondly, the grid was refined after each resolution, such that at the final resolution the control points were spaced 12 mm apart again. This yields 16 experiments on 26 image pairs, resulting in a total of 416 registrations. For the cost function MI was used. The Robbins-Monro stochastic optimisation method was applied, using 1000 iterations per resolution level. The image sampler was configured to select 2000 samples randomly in each iteration.

One hundred corresponding points in each baseline and follow-up scan were established by two independent observers using a semi-automatic algorithm [42]. The transformation  $T_{\hat{\mu}}$  was applied to the annotated points in the fixed image using `transformix`. To evaluate the registration accuracy, the mean distance between the resulting locations and the reference standard of the observer annotations was computed.

Figure 2.5 shows box-and-whisker plots of the mean distance to the annotations of one of the observers. The interobserver variability is shown in the left-most column. The first group R1-R8 displays the results without grid refinement. The second group shows the results with grid refinement. When no grid refinement was used, the registration quality improved until  $R = 3$ , but deteriorated for  $R > 3$ . Apparently, the dense B-spline grid yielded too much freedom on the heavily smoothed images. With grid refinement the results kept improving with increasing  $R$  up to  $R = 6$  (note the decreasing number of outliers above 5 mm). In practice, when considering the computation time, three or four resolutions with grid refinement seems to be a reasonable choice. With these settings the runtime was about 10 minutes on an AMD Opteron running at 2.4 GHz.

## 2.4 Conclusion

A software package, `elastix`, for medical image registration has been presented. Rather than implementing a single registration method, `elastix` is a collection of parametric intensity-based registration methods. Thanks to the modular design, the user can easily construct a registration algorithm, tailored to a specific application. Configuration of the registration method can be accomplished by writing a few lines in a parameter text file, without having to write any programming code. `elastix` has a command-line interface, which simplifies batch-processing of large numbers of data sets. Registration of large three-dimensional images can be done efficiently, thanks to the use of stochastic subsampling techniques.

The usage of `elastix` has been illustrated with three experiments. In the first experiment, eight transformation models were compared for the interpatient registration of 50 MR prostate scans. In the second experiment, we reproduced a result from the literature, showing that the so-called grid effect can be reduced by sampling the fixed image off the voxel grid. The third experiment demonstrated the importance of choosing a suitable hierarchical (multiresolution) strategy, by registering 26 chest CT image pairs with 16 different multiresolution configurations. These three investigations are just a few examples of the many possible comparative studies that one can perform with `elastix`.

The software has been used in several research projects, including [3, 9, 11, 19, 43, 44]. Both the executables and the source code are publicly available. The source code provides the users with the exact construction of the available algorithms, and allows them to enhance the functionality of `elastix` by adding their own algorithms. These features, in combination with the modular design, make `elastix` a useful tool for research on medical image registration.

## Acknowledgements

The authors thank dr. U.A. van der Heide for providing the prostate MR scans and dr. M. van Vulpen for the manual prostate segmentations.



# Tissue-Dependent Filtering of the Deformation Field

# 3

## Abstract

---

In present-day medical practice it is often necessary to nonrigidly align image data. Current registration algorithms do not generally take the characteristics of tissue into account. Consequently, rigid tissue, such as bone, can be deformed elastically, growth of tumours may be concealed, and contrast-enhanced structures may be reduced in volume.

We propose a method to locally adapt the deformation field at structures that must be kept rigid, using a tissue-dependent filtering technique. This adaptive filtering of the deformation field results in locally linear transformations without scaling or shearing. The degree of filtering is related to tissue stiffness: more filtering is applied at stiff tissue locations, less at parts of the image containing nonrigid tissue. The tissue-dependent filter is incorporated in a commonly used registration algorithm, using mutual information as a similarity measure and cubic B-splines to model the deformation field. The new registration algorithm is compared with this popular method.

Evaluation of the proposed tissue-dependent filtering is performed on 3D CT data of the thorax and on 2D Digital Subtraction Angiography (DSA) images. The results show that tissue-dependent filtering of the deformation field leads to improved registration results: tumour volumes and vessel widths are preserved rather than affected.

---

Based upon: M. Staring, S. Klein, and J.P.W. Pluim, Nonrigid Registration with Tissue-Dependent Filtering of the Deformation Field, *Physics in Medicine and Biology*, vol. 52, no. 23, pp. 6879 - 6892, December 2007.

### 3.1 Introduction

**N**ONRIGID image registration is an important technique in the field of medical image processing [14–16]. It performs the task of finding the spatial correspondence between two images, i.e. it searches for a one-to-one mapping from voxels in one image to voxels in the other image. This mapping is called the transformation. Commonly, first a rigid or affine registration is performed to capture the global transformation. The rigid or affine transformation is then used as a starting point for a nonrigid registration, which searches for local deformations. An example of a popular nonrigid registration algorithm is the one of [4].

A possible drawback of such a general nonrigid registration approach is that all tissue in the image volume is treated nonrigidly. The rigidity or stiffness of different tissue types is not taken into account. The implicit assumption that all tissue is deformable is violated in at least three situations:

- a. *truly rigid tissue*: Bone tissue that is present in the image volume must not be allowed to deform. Treating the bones as deformable tissue can lead, for example, to undesired thickening and bending of ribs [45].
- b. *temporal changes of tissue*: An example is the study of tumour progression in a patient. A successful nonrigid registration will achieve a perfect fit between tumours of two different time points, effectively concealing tumour growth or shrinkage. To visualise the tumour growth with a difference image, the tumours have to be kept rigid.
- c. *intensity changes of tissue*: Local intensity changes can be induced, for example, by injection of a contrast agent. [5] and [46] demonstrate that contrast-enhanced breast lesions in MR are compressed after intensity-based nonrigid registration. In [47] shrinkage is reported for contrast-enhanced vessels in CT Digital Subtraction Angiography (CT-DSA).

In cases (b) and (c) the tissue is not truly, i.e. physically, rigid. However, in order to prevent undesired volume changes of these structures, they should be regarded as undeformable.

Several methods are reported in the literature to address this issue. Some of these methods add a regularisation or penalty term to the registration, thereby constraining the transformation. [5] achieve incompressibility by penalising a deviation of the determinant of the Jacobian of the transformation from 1. The method does not distinguish between different tissue types, nor does it enforce rigidity of stiff tissue. [48] employ a rigidity penalty term to certain structures by penalising a deviation of the Jacobian from orthonormality. Other ways to enforce rigidity of structures have also been presented. [46] propose to couple the control points of a B-spline deformation to enforce rigidity on certain structures. It is assumed that the rotational part of the deformation can be captured by the initial rigid registration. Another approach is taken by [49], who constrain the nonlinear part of a deformation at rigid locations by multiplication with a weight function. Multiple rigid transformations are used by [35], each defined on different

parts of the image, but influencing each other. [50] model the deformation field using a three-component deformation model, consisting of rigid, elastic and fluid structures. [51] have employed a finite element method to introduce tissue-specific parameters. For the viscous fluid algorithm a locally adaptable regularisation technique is proposed that spatially varies the viscosity parameter [52].

A structure is kept rigid during registration if and only if the underlying deformation field represents a locally rigid transformation. We propose to achieve rigidity of certain structures by filtering the deformation field.

Filtering of deformation fields has been used in other situations, for example in optical flow like registration algorithms. In these cases the filtering is typically used to regularise the deformation field to get more plausible results. [53] uses a Gaussian kernel to smooth the deformation field. This kernel is not adapted to the underlying tissue type; it is fixed for the entire image volume. [54] use the Demons algorithm, which they adapt with a filtering of the deformation field. They derive and employ *a priori* and *a posteriori* smoothing weights for the filtering; these weights vary over the image, but do not depend on tissue type. Another method to smooth the deformation field is used by [55]. They regularise the deformation field by smoothing it with a Gaussian kernel with a standard deviation depending on the spatial position. They propose to take a large standard deviation in areas where little deformation is expected (e.g., rigid areas), and vice versa. At boundaries of rigid and nonrigid tissue the deformation field within the rigid objects is influenced by the nonrigid deformations of its surroundings. Both the extent and strength of this influence is quite large, because a large standard deviation is used. Also, a scaling of the deformation field within rigid objects is not prevented with this method. Therefore, that approach is not suitable for the problem we address. To ensure rigidity of the deformation field, an asymmetric kernel is needed with a large standard deviation inside the rigid object and a small one outside.

The filter we propose has the property that undesired deformations at rigid structures are addressed, while the surroundings are free to deform. The deformation field is adapted by filtering the deformation field in a tissue-dependent way: the filter is different for each voxel, depending on the stiffness of the underlying tissue. In this way some parts of the image can be kept rigid, whereas other parts are allowed to deform freely. In Section 3.2 the proposed tissue-dependent filter for regularising the deformation field is described in detail. Incorporation of the filter in the registration algorithm is also described in that section.

For evaluation two applications are considered: the study of tumour progression in CT thorax imaging (an example of case (b)), and the study of the vasculature at several regions in the human body in DSA imaging (an example of case (c)). The validation of the algorithm is described in Section 3.3, and the chapter is concluded in Section 3.4.

## 3.2 Method

Registration of a moving image  $I_M(\mathbf{x}) : \Omega_M \subset \mathbb{R}^D \mapsto \mathbb{R}$  to a fixed image  $I_F(\mathbf{x}) : \Omega_F \subset \mathbb{R}^D \mapsto \mathbb{R}$ , both of dimension  $D$ , is the problem of finding the deformation field  $\mathbf{d}$ , or

equivalently the transformation  $\mathbf{T}(\mathbf{x}) = \mathbf{x} + \mathbf{d}(\mathbf{x})$ , that spatially aligns  $I_M(\mathbf{x} + \mathbf{d}(\mathbf{x}))$  and  $I_F(\mathbf{x})$ . Registration can be defined as a minimisation problem:

$$\hat{\mathbf{d}} = \arg \min_{\mathbf{d}} \mathcal{C} [I_F(\mathbf{x}), I_M(\mathbf{x} + \mathbf{d}(\mathbf{x}))], \quad (3.1)$$

where  $\mathcal{C}$  denotes a cost function and  $\hat{\mathbf{d}}$  the optimal solution.

As discussed in the introduction, registration can result in a deformation field  $\hat{\mathbf{d}}(\mathbf{x})$  which is nonrigid at rigid structures. The filter we propose filters the deformation field adaptively, based on the content of a ‘stiffness coefficient’ image  $c(\mathbf{x})$ .

### 3.2.1 Filter Design

For rigid tissue the corresponding part of the deformation field should be linear without scaling or shearing. Given the deformation field  $\mathbf{d}$ , we achieve this by calculating a weighted mean  $\mathbf{m}(\mathbf{x})$  of  $\mathbf{d}$  over a neighbourhood  $\mathcal{N}_{\mathbf{x}}$  of  $\mathbf{x} \in \Omega_F$ :

$$\mathbf{m}(\mathbf{x}) \triangleq \sum_{\mathbf{x} \in \mathcal{N}_{\mathbf{x}}} c(\mathbf{T}(\mathbf{x})) \mathbf{d}(\mathbf{x}) / \sum_{\mathbf{x} \in \mathcal{N}_{\mathbf{x}}} c(\mathbf{T}(\mathbf{x})), \quad (3.2)$$

where  $c(\mathbf{x})$  is the stiffness coefficient image with values between 0 and 1. The stiffness coefficient image indicates for each voxel the relative stiffness of the underlying tissue. In case  $\sum_{\mathbf{x} \in \mathcal{N}_{\mathbf{x}}} c(\mathbf{T}(\mathbf{x})) = 0$ , we choose  $\mathbf{m}(\mathbf{x}) = \mathbf{d}(\mathbf{x})$ , to avoid division by zero. In order to control the degree of filtering, we define a filtered deformation field  $\mathcal{F}(\mathbf{d}(\mathbf{x}))$  by assigning a value close to the locally weighted mean deformation  $\mathbf{m}(\mathbf{x})$  when the stiffness coefficient  $c(\mathbf{x})$  is high, and a value close to the original deformation  $\mathbf{d}(\mathbf{x})$  for low stiffness coefficients:

$$\mathcal{F}(\mathbf{d}(\mathbf{x})) \triangleq (1 - c(\mathbf{T}(\mathbf{x}))) \mathbf{d}(\mathbf{x}) + c(\mathbf{T}(\mathbf{x})) \mathbf{m}(\mathbf{x}). \quad (3.3)$$

For neighbourhoods on the edge of rigid and nonrigid structures, the weighted mean  $\mathbf{m}(\mathbf{x})$  is determined mainly by the voxels having a high stiffness coefficient. In other parts of the image Equation (3.2) yields an averaging, also in the nonrigid parts. By defining the tissue-dependent filtering as in (3.3), it is controlled to what extent the weighted mean is used, such that it is only employed at rigid parts of the image. The result is a more homogeneous, linear deformation field at regions with a high  $c(\mathbf{x})$ , where tissue is more stiff. At regions with a low  $c(\mathbf{x})$  the deformation field is unaffected by the filtering, and therefore those regions can deform freely, as desired.

To obtain the stiffness coefficient image  $c(\mathbf{x})$  it is necessary to process the moving image, for example by segmenting the relevant structures and assigning a stiffness coefficient (a value within  $[0, 1]$ ) to the voxels. For CT images the Hounsfield units can be rescaled to values in  $[0, 1]$ , since stiffer material usually implies a higher attenuation value. In this work we focus on the registration method, and therefore we have opted for a simple manual segmentation.

We are interested in the deformed moving image. Therefore, we only have to take care that the transformation of that image, as found by the registration, obeys restrictions of




Figure 3.1: The original 2D deformation field (a), stretches the underlying tissue both vertically and horizontally. Applying the tissue-dependent filter, with  $c(\mathbf{x}) = 1$  everywhere, once to this field results in (b). This field still stretches tissue, but less. After applying the adaptive filter several times, the result is a homogeneous deformation field (c), yielding a rigid transformation of the underlying tissue. (d) illustrates in 1D the adaptive filter applied to a scaling. The dashed line ( $k = 0$ ) shows an unfiltered deformation function that scales the data ( $\mathbf{d}(\mathbf{x}) = 0.2\mathbf{x}$ ). The dashed-dotted and the solid lines refer to applying the adaptive filter 1 and 10 times, respectively, using  $l = 5$ . Above the graph the stiffness coefficients  $c(\mathbf{T}(\mathbf{x}))$  are shown: 0 referring to complete nonrigid tissue and 1 to rigid tissue. In this case  $c(\mathbf{T}(\mathbf{x}))$  is the same for all three ( $k = 0, 1, 10$ ) transformations  $\mathbf{T}$ . The adaptive filter works from the edge of the rigid structure to the inside, and converges after 10 iterations.

rigidity. This means that only the relevant structures in the moving image have to be segmented.

The tissue-dependent filter  $\mathcal{F}$  can be applied to the deformation field several times in succession to increase the power of the filter. We define:

$$\mathcal{F}^k(\mathbf{d}(\mathbf{x})) \triangleq \mathcal{F}(\mathcal{F}^{k-1}(\mathbf{d}(\mathbf{x}))), \quad (3.4)$$

$$\mathcal{F}^0(\mathbf{d}(\mathbf{x})) \triangleq \mathbf{d}(\mathbf{x}). \quad (3.5)$$

When applying  $\mathcal{F}$  multiple times ( $k > 1$ ), the orientation and magnitude of the deformation field in a neighbourhood become more similar. This effect is shown in Figure 3.1.

The neighbourhood  $\mathcal{N}_{\mathbf{x}}$  in (3.2) is another parameter that can be controlled. It is chosen to be a square in 2D and a cube in 3D, of size  $l \times l$  and  $l \times l \times l$  voxels, respectively. For large  $l$  the adaptive filter will linearise larger parts of the image at once, thereby increasing the power of the filter. However, for large  $l$ , the deformation fields of two separate rigid structures lying close together might both be within the neighbourhood. This way, the deformation field of both rigid structures evolves to the average of them, which is undesired. Therefore, a small neighbourhood of  $l = 3$  is used in this study. The power of the tissue-dependent filter is controlled with the number of successive applications  $k$  of the filter, see (3.4). In the following section a second parameter  $m$  is introduced, that also controls the power of the filter.

Figure 3.1(d) illustrates that the effect of the tissue-dependent filter is first observed near the boundary of rigid structures. After successive application of the filter, the results

are propagated to the inside of such regions. Strong edges in the stiffness coefficient image at the boundary of rigid and nonrigid tissue can result in sharp transitions in the deformation field after filtering. Invertibility of the deformation field is not guaranteed. This can be easily achieved by defining some rules on the displacement field [55]. However, no problems were encountered in the experiments and therefore this step is omitted. The tissue-dependent filter is based on a weighted mean, therefore, the filtered deformation field will eventually, for large  $k$ , evolve to a translational field, not allowing rotations of rigid objects.

The computational complexity of the tissue-dependent filter is linearly dependent on the size of the neighbourhood  $l^D$ , the number of voxels  $N$  in the image, and the number of successive applications  $k$ :  $\mathcal{O}(l^D N k)$ .

### 3.2.2 Tissue-Dependent Nonrigid Registration

This section describes how the tissue-dependent filter  $\mathcal{F}$  is incorporated in the registration. The registration algorithm is largely based on the papers of [4], [2], and [30].

The methods described by these authors use the well-known *mutual information* between the fixed image  $I_F$  and the moving image  $I_M$  as a cost function  $\mathcal{C}$ . In this study the implementation of [30] is adopted, who use B-spline Parzen windows to describe the joint histogram. This way the cost function is a continuous, differentiable function of the transformation parameters  $\mathbf{p}$ . The derivative of the cost function  $\partial\mathcal{C}/\partial\mathbf{p}$  can therefore be computed analytically. The deformation field  $\mathbf{d}$  is parameterised by cubic B-splines [36] with parameters  $\mathbf{p}^{\text{BS}}$ . The B-splines modelling the deformation field have an inherent smoothness, which is desired in registration problems. They also have local support, which is beneficial both for modelling local transformations, and for fast computation. The derivative  $\partial\mathbf{d}/\partial\mathbf{p}^{\text{BS}}$ , needed for computing  $\partial\mathcal{C}/\partial\mathbf{p}$ , is available analytically.

An iterative optimisation strategy is used to solve the minimisation problem (3.1):

$$\mathbf{d}_{t+1}(\mathbf{x}) = \mathbf{d}_t(\mathbf{x}) + \mathbf{\Delta}_t(\mathbf{x}; \mathbf{p}^{\text{BS}}), \quad (3.6)$$

where  $\mathbf{\Delta}_t$  is an incremental deformation field at iteration  $t$ , as found by the optimiser. Incorporating our tissue-dependent filtering of the deformation field in the optimisation scheme is straightforward:

$$\mathbf{d}_{t+1}(\mathbf{x}) = \mathcal{F}^k(\mathbf{d}_t(\mathbf{x}) + \mathbf{\Delta}_t(\mathbf{x}; \mathbf{p}^{\text{BS}})). \quad (3.7)$$

The resulting deformation  $\mathbf{d}_{t+1}(\mathbf{x})$  is a filtered version of a B-spline based deformation field for  $t > 1$  and it cannot be parameterised by B-splines. It is described by a vector field. Only the additional deformation field  $\mathbf{\Delta}_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$  is modelled by B-splines, which has fewer restrictions than modelling the entire transformation with B-splines. The optimisation is performed on  $\mathbf{\Delta}_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$ , i.e. on the B-spline coefficients only. This is computationally advantageous compared to optimising over the entire deformation field.

The tissue-dependent filtering does not necessarily have to be applied after every optimisation iteration. To decrease the power of the filter it can be applied after every  $m$  iterations. Summarising, the proposed nonrigid registration is given by:

1. Perform  $m > 0$  iterations of the optimiser, resulting in a new additional deformation field  $\Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$ .
2. Calculate the deformation field  $\tilde{\mathbf{d}}_{t+1}(\mathbf{x}) = \mathbf{d}_t(\mathbf{x}) + \Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$ .
3. Filter the deformation field, i.e., calculate  $\mathbf{d}_{t+1}(\mathbf{x}) = \mathcal{F}^k(\tilde{\mathbf{d}}_{t+1}(\mathbf{x}))$  according to Equations (3.2), (3.3), and (3.4),(3.5).

Repeat step (i) - (iii) until convergence.

A stochastic gradient descent optimiser is used to determine the additional deformation field  $\Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$  in the following way:  $\Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}}) = -a_t \partial \mathcal{C} / \partial \mathbf{p}^{\text{BS}}$ , where  $a_t$  is the size of the step taken in the direction  $-\partial \mathcal{C} / \partial \mathbf{p}^{\text{BS}}$ . The stochastic gradient descent optimiser uses an *approximation* of the derivative of the mutual information with respect to the B-spline parameters  $\mathbf{p}^{\text{BS}}$ . Approximation is done by using only a small random subset of voxels from the fixed image. This has been shown to accelerate registration significantly, without compromising registration accuracy [19]. At each iteration a step is taken towards the minimum of the cost function  $\mathcal{C}$ . The size of this step  $a_t$  is determined by a decreasing function of the iteration number  $t$ . This function is of the form  $a_t = a / (t + A)^\alpha$ , where  $a$ ,  $A$ , and  $\alpha$  are user-defined constants. Following the suggestions in [56], in this chapter  $A = 100.0$  and  $\alpha = 0.602$  are used. The parameter  $a$  is related to the expected magnitude of the deformation, and is tuned for each application. As a stopping condition a user-defined number of iterations is used, upon which convergence is assumed. In order to avoid local minima, a multiresolution approach is taken. A Gaussian image pyramid is used with a subsampling factor of two. Also a multigrid approach is taken: when the image resolution is doubled, the B-spline control point spacing is halved. Prior to the nonrigid registration an affine registration is performed, in order to capture the global transformation between fixed and moving image.

### 3.3 Experiments and Results

In order to evaluate the effectiveness of the proposed tissue-dependent nonrigid registration, it is compared with a general nonrigid registration approach based only on B-splines. This general approach does not apply the tissue-dependent filtering of the deformation field, but is similar to the proposed algorithm in all other respects (see the description in Section 3.2.2). The methods will be referred to as ‘BS’ (B-spline only based) and ‘BSF’ (B-splines with tissue-dependent filtering of the deformation field).

The two methods are illustrated on a synthetic example (Section 3.3.1), and compared on clinical data, viz. 3D CT follow-up data of the thorax containing lung tumours (Section 3.3.2), and 2D Digital Subtraction Angiography (DSA) image data, see Section 3.3.3. The applications are examples of cases (b) and (c) of the list given in the Introduction. Both applications consider structures that are not physically rigid, but to prevent undesired volume changes, they should be regarded as undeformable. Special attention is given to the influence of the parameters  $k$  and  $m$ .

In all experiments with BSF, the size of the neighbourhood  $\mathcal{N}_{\mathbf{x}}$  in Equation (3.2) is  $3 \times 3$  voxels in 2D and  $3 \times 3 \times 3$  in 3D. For both BS and BSF, the number of bins for

calculation of the mutual information was set to 32. All experiments were performed with software developed by the authors (<http://elastix.isi.uu.nl>). This registration package is largely based on the Insight Segmentation and Registration Toolkit [33]. Filtering the deformation field took approximately 1 s. and 1.7 s. for  $k = 1$  and  $k = 10$ , respectively, on 2D  $512^2$  data. For 3D  $256^3$  data this took 140 and 275 s. approximately. An AMD Opteron 2218 running at 2.6 GHz. was used, without multi-threading enabled.

### 3.3.1 2D Synthetic Example

A 2D synthetic example image was constructed to demonstrate the behaviour of the non-rigid registration algorithm with and without tissue-dependent filtering. Figures 3.2(a) and 3.2(b) show the fixed and moving image, respectively, both of size  $128 \times 128$  pixels. The central white structure represents a rigid structure, for which the stiffness coefficient image  $c(\mathbf{x})$  was set to 1.0. The two larger white structures represent some nonrigid image content, for which  $c(\mathbf{x}) = 0.0$ . A standard nonrigid registration was performed using three resolutions. At each resolution 350 iterations of the optimiser were used. A multigrid approach was taken, with a B-spline grid spacing at the final resolution of 64 pixels. No affine or rigid registration was performed prior to nonrigid registration. The result of the registration is shown in Figure 3.2(c). It is clear that the rigid inner structure has thickened. The B-spline fails to keep the inner rigid structure rigid, because the control points are pulled apart by the contraction of the two outer white structures. The experiment is repeated with BSF with low filtering power ( $k = 1, l = 3, m = 100$ ), resulting in a still thickened inner structure, see Figure 3.2(d). The power is increased by applying the filter more often ( $m = 10$ ), and by extending the neighbourhood size to  $l = 5$ . Both approaches result in a perfect registration. In Figure 3.2(e) a horizontal profile through the inner rigid structure of the  $x$ -component of the deformation field is shown for the three BSF versions. BSF (1,100,3) still thickens the inner structure, while the other two show no displacement at the centre.

### 3.3.2 3D CT Thorax Data with Lung Tumours

Nonrigid registration is a valuable tool for following disease progress of patients over time. A possible way to detect changes is by analysing the deformation field found by the nonrigid registration, as is done by [57]. Another way to evaluate this progress is to visually inspect the difference between a first scan, taken at time  $t_0$ , and a registered follow-up scan, taken at  $t_1$ . For the case of patients suffering from lung tumours, there might be tumour growth between the first scan and the follow-up. Because of nonrigid motion occurring between scans, a rigid registration is not sufficient to achieve good alignment of the anatomy. However, standard nonrigid registration methods will minimise the difference between tumours at different time points, effectively concealing tumour growth (see Figure 3.3). Therefore, the tumours should not be allowed to deform by the nonrigid registration algorithm. It is assumed that the tumour volume does not change between scans due to differences in the lung inspiration levels, because the protocol for acquiring the data was such that inspiration levels are approximately equal. In this






Figure 3.2: BS registration compared with the adaptive filtering of the deformation field. (a) The fixed image  $I_F$ , (b) the moving image  $I_M$ , (c) the result with BS, (d) the result with BSF ( $k = 1$ ,  $m = 100$ ,  $l = 3$ ). When the neighbourhood size is increased to  $l = 5$ , or when more filtering is applied ( $m = 10$ ), the resulting images are equal to the fixed image. In (e) a profile in the horizontal direction of the  $x$ -component of the deformation field is depicted. In the legend the parameters  $(k, l, m)$  refer to the corresponding BSF registration. The two vertical lines indicate the rigid part.

section, BS registration of CT thorax follow-up scans of patients suffering from lung tumours is compared with BSF registration. Rigidity is evaluated with tumour volume measurements.

### Data Description

Registration was performed on CT follow-up data sets of the thorax of five patients having lung tumours. For each patient two images were taken at different time points. The data were acquired with a 16-slice spiral CT scanner (Mx8000 IDT 16, Philips Medical Systems, Best, The Netherlands). The images have an in-plane resolution of  $512 \times 512$  pixels. The number of slices varies for the data sets, ranging from 400 to 550. The in-




Figure 3.3: BS registration compared with BSF for the CT thorax follow-up data. Difference images of the registration result with the fixed  $t_0$  image are shown. (a) for affine, (b) for BS, and (c) for BSF registration. The white box shows a tumour location.

Table 3.1: Average lung overlap. Between brackets the parameters  $(k, m)$  are given.

before registration	affine	BS	BSF (10,1)	BSF (1,1)	BSF (1,10)
$0.64 \pm 0.22$	$0.92 \pm 0.06$	$0.97 \pm 0.02$	$0.97 \pm 0.02$	$0.97 \pm 0.02$	$0.97 \pm 0.02$

plane voxel size is around  $0.7 \times 0.7$  mm. Slice thickness was always 1.0 mm and slices were reconstructed every 0.7 mm. Before registration, each data set was downsampled with a factor of two in each dimension by discarding odd rows, columns and slices, to reduce computer memory and computational load. The five data sets contain thirty-six tumours in total, with an average volume of  $2.8 \pm 3.4$  ml for the first scan  $t_0$  and  $5.9 \pm 7.2$  ml for the follow-up  $t_1$ . No new tumours had developed at  $t_1$ .

### Experiment Setup

The CT image taken at time  $t_0$  is set to be the fixed image, the CT image taken at time  $t_1$  the moving image. In order to get a coarse alignment between fixed and moving image an affine registration was performed first. For both the affine and the nonrigid registration three resolutions were employed. For the nonrigid registration the B-spline grid spacing at the final level was 16 voxels. In the three resolutions 100, 100, and 300 iterations were used for the optimiser, respectively. In every iteration 5000 samples were selected to calculate (the derivative of) the mutual information, see Section 3.2.2. The parameter  $a$  that defines the step size for the optimisation algorithm was chosen 100000.0, 70000.0, and 50000.0 for the three resolutions. For BSF, a crude manual segmentation of the tumours was used to define  $c(\mathbf{x})$ , setting  $c(\mathbf{x})$  to 1.0 for voxels within the tumour and to 0.0 elsewhere. The BSF method was tested with varying strength of the adaptive filtering, with  $(k, m) = (1, 10)$ ,  $(1, 1)$ , and  $(10, 1)$ .

### Results

Accuracy of the registration is measured by calculating the lung overlap of the registered image and the fixed image. For this purpose, automatic lung segmentations were made with an algorithm based on the method by [58], described in detail in [59]. The segmentation does not include the lung vasculature and the tumours. The overlap measure is defined as

$$\text{overlap} \triangleq \frac{2|L_1 \cap L_2|}{|L_1| + |L_2|}, \quad (3.8)$$

where  $L_i$  is the set of all voxels within the lung, and  $|L_i|$  denotes the size of set  $L_i$ . The average and standard deviation values of lung overlap were calculated over all data, and are reported in Table 3.1. The results show that a good lung overlap was achieved with both nonrigid registration algorithms. This is confirmed by visual inspection of the results, see Figure 3.3 for an example.




Figure 3.4: Deformation fields for the CT thorax follow-up data around a tumour, following from the (a) BS, and the (b) BSF registration. Where BS compresses the tumour ( $r_{BS} = 0.69$ ), BSF retains rigidity ( $r_{BSF} = 0.98$ ).

For evaluation of the rigidity of the tumours, a second manual segmentation of the tumours in the moving image is used. The second segmentation was performed in a precise manner, unlike the segmentation used for BSF. Tumour volume measurements were performed to evaluate if the registration is volume preserving, a condition for rigidity. Volume preservation is expressed by the ratio  $r$  between the tumour volume after registration, denoted by  $v_{reg}$ , and the tumour volume at  $t_1$  (since  $t_1$  is the moving image), denoted by  $v_{t_1}$ :  $r_{reg} = v_{reg}/v_{t_1}$ . The tumour segmentation after registration is determined by applying the coordinate transformation to the tumour segmentation of the moving image. Linear interpolation was used in this step. If a nonrigid registration is volume preserving for a tumour, the volume ratio will be 1.0; if a tumour is compressed  $r_{reg} < 1.0$ . For ratios it is better to use the *geometric* mean and standard deviation, instead of their arithmetic counterparts. This can be easily seen with a small example. Say, we have two ratios 0.5 and 2.0. The arithmetic mean of those two ratios is 1.25, whereas the geometric mean equals 1.0, rating a two times increase in volume equal to a two time decrease. The geometric mean is defined as  $\mu_g = \sqrt[n]{\prod_{i=1}^n r_{reg}^i}$ . From the definition of the geometric standard deviation it follows that  $\sigma_g \geq 1$ . The geometric mean volume ratios and standard deviations are reported in Table 3.2, where the symbol  $\times/$  is used to indicate the distinction with the arithmetic mean and standard deviation. Geometric means were calculated for four volume ratio groups and for all ratios together. The tumours were grouped according to tumour growth  $v_{t_1}/v_{t_0}$ . The third group, for example, is the group of tumours with tumour growth between  $3/2$  and 3. Table 3.2 shows that volume was much better preserved with BSF, compared to BS. For a low strength of the adaptive filter (BSF (1,10)) volume-preservation is low, which increases by increasing the power (BSF (1,1) and BSF (10,1)). Volume is by definition preserved for rigid registration. Residual volume differences of the BSF method can partly be explained by the fact that the tissue-dependent nonrigid registration obtains a trade-off between the rigidity of tumours and the maximisation of similarity. Especially when there are large tumour volume differences between  $t_0$  and  $t_1$  (the fourth group in Table 3.2), volume preservation is harder to obtain. Examples of the deformation fields after BS and BSF registration are given in Figure 3.4. Unlike BS, the deformation field for BSF is clearly rigid at the tumour location.

Table 3.2: Geometric mean tumour volume ratios. Between brackets the parameters  $(k, m)$  are given.

group	$v_{t_1}/v_{t_0}$	$r_{\text{rigid}}$	$r_{\text{BS}}$	$r_{\text{BSF}} (1,10)$	$r_{\text{BSF}} (1,1)$	$r_{\text{BSF}} (10,1)$
(0, 1]		1.00 ×/ 1.00	1.06 ×/ 1.03	1.01 ×/ 1.01	1.00 ×/ 1.00	0.99 ×/ 1.00
(1, 3/2]		1.00 ×/ 1.00	1.09 ×/ 1.11	1.03 ×/ 1.03	1.02 ×/ 1.02	1.02 ×/ 1.02
(3/2, 3]		1.00 ×/ 1.00	0.89 ×/ 1.04	0.93 ×/ 1.04	0.97 ×/ 1.02	0.99 ×/ 1.01
(3, ∞)		1.00 ×/ 1.00	0.88 ×/ 1.05	0.93 ×/ 1.04	0.96 ×/ 1.02	0.98 ×/ 1.01
all		1.00 ×/ 1.00	0.95 ×/ 1.08	0.96 ×/ 1.04	0.99 ×/ 1.02	0.99 ×/ 1.02

### 3.3.3 Digital Subtraction Angiography

Digital Subtraction Angiography (DSA) is an established modality for visualising blood vessels in the human body. During image acquisition patient motion often occurs, due to breathing, heart beat, activity in the intestines, or movement of the body. This motion results in artifacts in the subtraction images. In Figure 3.5(a) a subtraction image is shown after rigid registration; the artifacts in the original subtraction image are even larger. In order to correct for patient motion after acquisition, a nonrigid registration is needed. Typically, a sequence of images is taken, where different parts of the vasculature are visible at different times. To see the whole imaged vasculature, all the images from the specific sequence have to be registered to some fixed image. The first image, acquired just before the arrival of the contrast bolus and known as the baseline image, is used as the fixed image. As reported in the literature [5,46], nonrigid registration of images containing contrast-enhanced structures can lead to significant compression of those structures. Switching the fixed and the moving image and simply using BS does not guarantee that compression is avoided, and requires the inverse of the transformation. Therefore, a nonrigid registration is required that can treat the vasculature different from other tissue, to maintain vasculature size.

#### Data Description

The 2D digital X-ray angiography image data were acquired with an Integris V3000 C-arm imaging system (Philips Medical Systems, Best, The Netherlands). In total, twenty-six image sequences of twelve different patients were obtained. The image sequences were of size  $512 \times 512$  pixels for twenty-two data sets, and  $1024 \times 1024$  pixels for four data sets; they contain about ten images each. Intensities in the DSA images range approximately from 100 to 950, with an arithmetic mean and standard deviation of  $550 \pm 180$ . Images were taken of different locations in the body: abdomen (10), brain (5), hip and foot (4), heart (1), neck (5) and lungs (1). The first image in each sequence was taken before arrival of the contrast bolus, in the following images contrast agent is visible in parts of the vasculature.




Figure 3.5: BS registration compared with BSF for the DSA data. Difference images of the registration result with the fixed image are shown. (a) for rigid, (b) for BS, and (c) for BSF (2,1) registration. Note the substantial change in vessel radius of the BS method, compared to the other algorithms. Whereas the BS method compresses the vasculature, the BSF method is much more capable of retaining vessel widths.

### Experiment Setup

The baseline image was taken to be the fixed image. For our experiments one image from each sequence was registered to the fixed image. The image showing the most vasculature was manually selected as the moving image. To get a coarse alignment between fixed and moving image, a rigid registration was performed prior to nonrigid registration. Two resolutions were used, with a B-spline grid spacing of 32 and 16 pixels at the first and final resolution level, respectively. For the two resolutions 600 and 300 iterations were used, respectively. At every resolution 5000 samples were used to calculate (the derivative of) the mutual information, see Section 3.2.2. The parameter  $a$  that defines the step size for the optimisation algorithm was chosen 6000.0, and 3500.0 for the two resolutions. For the BSF registration, a crude manual segmentation of the vessels was used to define  $c(\mathbf{x})$ , setting  $c(\mathbf{x})$  to 1.0 for voxels within the vasculature and to 0.0 otherwise. The BSF method was tested with the parameter settings  $(k, m) = (1, 5), (1, 1), (2, 1),$  and  $(5, 1)$ .

### Results

The Root Mean Square Difference (RMSD) of  $I_M(\mathbf{x} + \mathbf{d}(\mathbf{x}))$  with  $I_F(\mathbf{x})$  at the background was calculated to see if the nonrigid registration indeed reduces motion artifacts. The background is defined as everything within the cone beam, but outside the manual vessel segmentation. Differences at vessel locations are not taken into account, since these differences are meant to be large. Arithmetic means and standard deviations of the RMSD were calculated for all 26 images. The results are reported in the top row of Table 3.3. Rigid registration reduced the motion artifacts. Both BS and BSF nonrigid registration methods improved substantially on that. In Figure 3.5 the difference image after rigid registration, and after the two nonrigid registration methods (BS and BSF (2,1)) is shown for an example DSA image. Visual inspection confirms the reduction in motion artifacts, as measured by the RMSD.

Table 3.3: Results for DSA data. Arithmetic means and standard deviations for the RMSD are displayed in the top row. Geometric means and standard deviations for the vessel diameter ratios are shown in the bottom row. Between brackets the parameters  $(k, m)$  are given.

	no registration	rigid	BS	BSF (1,5)	BSF (1,1)
RMSD	$14.0 \pm 6.0$	$13.5 \pm 5.7$	$12.1 \pm 4.2$	$11.8 \pm 4.0$	$11.9 \pm 4.0$
diameter	$1.00 \times / 1.00$	$1.00 \times / 1.00$	$0.85 \times / 1.16$	$0.91 \times / 1.11$	$0.96 \times / 1.06$

	BSF (2,1)	BSF (5,1)
...	$11.9 \pm 4.0$	$11.9 \pm 4.0$
	$0.97 \times / 1.04$	$0.98 \times / 1.03$




Figure 3.6: A box-and-whisker plot for the difference  $\mathcal{D}$  in background motion artifacts. The box represents the first and third quartiles. The horizontal line in this box is the median. The whiskers indicate the extent of the rest of the data, except for outliers. Outliers are points falling more than 1.5 times the interquartile range past the ends of the box, and are denoted by a dot.

In order to compare the reduction in motion artifacts between the different registration methods, the difference of the RMSDs is defined as  $\mathcal{D}_{i,j} = \text{RMSD}_i - \text{RMSD}_j$ , where  $i \neq j$  are the different registration methods. This difference was calculated for several combinations of registration methods and for all DSA data, resulting in the differences  $\mathcal{D}_{\text{rigid,noreg}}$ ,  $\mathcal{D}_{\text{BS,rigid}}$ ,  $\mathcal{D}_{\text{BSF,rigid}}$ , and  $\mathcal{D}_{\text{BSF,BS}}$ , where BSF is BSF (2,1). In Figure 3.6 a box-and-whisker plot for these differences in motion artifacts is shown. Rigid registration reduced the motion artifacts slightly, compared to no registration (top row). The second and third row show that nonrigid registration indeed improved substantially on rigid registration (rigid vs. BS and rigid vs. BSF both have  $p < 0.05$  in a two-tailed paired  $t$ -test). The bottom row indicates that the RMSD for the BS and BSF method are similar.

Rigidity of the vasculature is evaluated by manually measuring the vessel diameter  $vd$  at several locations. Six locations were selected for each of the 26 images, yielding a total

of 156 diameter measurements. As in the previous section, ratios  $r$  are used to evaluate vessel compression. The vessel diameter after registration is compared to the diameter before registration:  $r_{reg} = vd_{reg}/vd_{noreg}$ . The geometric mean of the vessel diameter ratios is reported in the bottom row of Table 3.3. Standard nonrigid registration (BS) clearly resulted in compression of the vessels, whereas the tissue-dependent filtering (BSF (2,1)) retained the vessel widths. When the adaptive filtering strength is decreased (BSF (1,1) and BSF (1,5)) rigidity-preservation is decreased, and vice versa (BSF (5,1)).

### 3.4 Discussion and Conclusions

In common nonrigid registration approaches all tissue is treated equally, thereby deforming rigid tissue. We propose a method that takes the rigidity of different tissue types into account, using a tissue-dependent filtering of the deformation field. Such a filter is described in this chapter, and incorporated in a nonrigid registration framework: tissue-dependent nonrigid registration.

In contrast to other work on filtering of the deformation field, we designed the filter specifically to maintain rigidity of tissue. All other work we are aware of uses filtering to obtain *smoothness* of the deformation, sometimes also in a tissue-dependent manner. The proposed filter does not affect the deformation field outside the region that is to be kept rigid, thereby allowing the neighbouring tissue of the rigid structures to deform freely. Methods that add a penalty term to the cost function to achieve rigidity need to weigh the penalty term against similarity. The choice of this weight is somewhat nonintuitive and a matter of tuning. Varying the strength of the proposed tissue-dependent filtering can be done with two intuitive parameters ( $k$  and  $m$ ). The results show that by increasing  $k$  and/or decreasing  $m$  the method is better capable of retaining rigidity, and vice versa. For very large  $k$  the deformation field will converge to a translational field. The proposed method also offers possibilities to move rigid objects that are in close proximity, independently of each other by adding an object label to the filtering process.

A segmentation of the rigid structures is needed to define which regions are to be kept rigid. In this chapter we used rough manual segmentations, which is labour intensive, and therefore not very practical in a clinical setting. Since a rough segmentation is sufficient, a (semi-) automatic segmentation method will usually be available.

The proposed tissue-dependent nonrigid registration was evaluated on CT thorax follow-up data and Digital Subtraction Angiography (DSA) data. It was compared against a standard B-spline based registration approach. From the experiments on the CT thorax follow-up data it is observed that tissue-dependent nonrigid registration is better in terms of preserving tumour volume than a standard B-spline based approach. The DSA data show a clear improvement of the proposed method over the B-spline method in retaining vessel widths.

Based on the results and on visual inspection of the data, it is concluded that the tissue-dependent nonrigid registration algorithm is indeed able to model locally rigid transformations, thereby improving registration results.





# A Rigidity Penalty Term for Nonrigid Registration

# 4

---

## Abstract

Medical images that are to be registered for clinical application often contain both structures that deform and ones that remain rigid. Nonrigid registration algorithms that do not model properties of different tissue types may result in deformations of rigid structures. In this chapter we propose a local rigidity penalty term, which is included in the registration function in order to penalise deformation of rigid objects. This term can be used for any representation of the deformation field capable of modelling locally rigid transformations. By using a B-spline representation of the deformation field, a fast algorithm can be devised. We compare the proposed method with an unconstrained nonrigid registration algorithm. It is evaluated on clinical 3D CT follow-up data of the thorax and on 2D DSA image sequences. The results show that nonrigid registration using the proposed rigidity penalty term is capable of nonrigidly aligning images, while keeping user-defined structures locally rigid.

---

Based upon: M. Staring, S. Klein, and J.P.W. Pluim, A Rigidity Penalty Term for Nonrigid Registration, Medical Physics, vol. 34, no. 11, pp. 4098 - 4108, November 2007.

## 4.1 Introduction

**I**MAGE REGISTRATION is an important technique in the field of medical imaging. In many clinical situations several images of a patient are made in order to analyse the patient's situation. These images are acquired with, for example, X-ray scanners, Magnetic Resonance Imaging (MRI) scanners, Computed Tomography (CT) scanners, and Ultrasound scanners, which provide knowledge about the anatomy of the subject. Combination of patient data, mono- or multi-modal, often yields additional clinical information not apparent in the separate images. For this purpose, the spatial relation between the images has to be found. Image registration is the task of finding a spatial one-to-one mapping from voxels in one image to voxels in the other image. Good reviews on the subject are given in [14–16].

Popular nonrigid image registration algorithms, as described by Rueckert *et al.* [4] and Mattes *et al.* [2], do not take the rigidity of different tissue types into account. This can lead to undesired effects in three situations:

- *The image contains various tissue types, each with its own mechanical stiffness.* An example is the presence of rigid objects like bones or surgical instruments in a soft tissue region. Not taking into account the spatially varying stiffness during the registration likely results in unwanted distortions of these rigid objects [45].
- *Structural changes over time (e.g. tumour growth) need to be visualised.* Differences that are not of interest (e.g. due to breathing or heart beat) should be compensated for by the nonrigid registration, while the relevant differences induced by changes in the objects of interest should be retained. This would lead to a subtraction image from which the relevant changes are immediately clear.
- *The visibility of structures varies between acquisitions.* For instance, in pairs of pre- and post-contrast images, like angiography images, vessels are visible in one image and not or only partially in the other. Intensity-based registration algorithms are designed to minimise the difference between two images, resulting in compression of the contrast-enhanced structures. Rohlfing *et al.* [5] and Tanner *et al.* [46] report shrinkage of contrast-enhanced breast lesions in MRI. Rohlfing *et al.* [47] report shrinkage for contrast-enhanced vessels in CT-DSA.

In all three situations the objects of interest should be considered as locally rigid by the registration algorithm.

Other possibilities to detect structural changes over time exist. For example, the deformation field as found by nonrigid registration can be analysed. This is commonly done by looking at the determinant of the Jacobian of the transformation, see [57, 60] and references therein. A discussion on the advantages and disadvantages of the two approaches can be found in the literature [61, 62].

Several methods to constrain deformations have been described in the literature. One that is well known and widely used is to employ a regularisation term to penalise undesired transformations. Typical regularisation terms include the bending energy of a thin plate [4], the linear elasticity constraint [24, 63], and the incompressibility constraint

[5, 64]. Methods that specifically enforce *rigidity* on structures have also been proposed. Tanner *et al.*[46] propose to couple the control points of a B-spline deformation to impose rigidity on certain structures. Another approach is taken by Little *et al.*[49], who use modified basis functions to describe the deformation. At rigid locations the deformation is constrained. Arsigny *et al.*[35] use multiple rigid transformations on different parts of the image, which they fuse in such a way that the transformation is invertible. Recently, some approaches were published [48, 65] in which rigidity is enforced by penalising deviation of the Jacobian of the transformation from orthonormality.

In this chapter we propose a novel penalty term that is capable of locally penalising nonrigid transformations, which we call a rigidity penalty term. It is based on three criteria a transformation must meet in order to be rigid: the transformation should be affine, and the rotation matrix of the transformation should be orthonormal and proper. Our rigidity penalty term is a weighted combination of three terms, each imposing one of these conditions. While parts of the proposed rigidity penalty proved to be similar or identical to penalty terms used by other authors [4, 5, 48, 65], we demonstrate that imposing the complete rigidity penalty term performs better than each of the terms separately. First results of the method were published previously [66]. In Section 4.2 we construct the rigidity penalty term and proof its validity. Nonrigid registration using the rigidity penalty term is compared against a standard unconstrained nonrigid registration algorithm in Section 3.3. Special attention is paid to the sensitivity of the results to the choice of the parameters controlling the rigidity penalty term. We end with a discussion and conclusion in Section 4.4.

## 4.2 Method

Registration of a moving image  $I_M(\mathbf{x}) : \Omega_M \subset \mathbb{R}^d \mapsto \mathbb{R}$  to a fixed image  $I_F(\mathbf{x}) : \Omega_F \subset \mathbb{R}^d \mapsto \mathbb{R}$ , both of dimension  $d$ , is the problem of finding a displacement  $\mathbf{u}(\mathbf{x})$  that makes  $I_M(\mathbf{x} + \mathbf{u}(\mathbf{x}))$  spatially aligned to  $I_F(\mathbf{x})$ . The quality of alignment is defined by a distance or similarity measure  $\mathcal{S}$ , such as the sum of squared differences (SSD), the correlation ratio, or the mutual information (MI) measure.

Because this problem is ill-posed, a regularisation or penalty term  $\mathcal{P}$  is often introduced that constrains  $\mathbf{u}$ . The registration problem is formulated as an optimisation problem in which the cost function  $\mathcal{C}$  is minimised w.r.t.  $\mathbf{u}$ , with:

$$\mathcal{C}[\mathbf{u}; I_F, I_M] = -\mathcal{S}[\mathbf{u}; I_F, I_M] + \alpha\mathcal{P}[\mathbf{u}], \quad (4.1)$$

where  $\alpha$  weighs similarity against regularity. This formalism is used in many other papers[4, 5, 24, 48, 63, 64]. Note that at the minimum of  $\mathcal{C}$  the derivatives of the similarity measure and the regularisation term are not necessarily zero. Rather, a balance is found between the two, which is influenced by the parameter  $\alpha$ . Therefore, the penalty term cannot be regarded as a hard constraint; it is sometimes referred to as a soft constraint.

The regularisation term  $\mathcal{P}$  can not only be considered as a way to achieve well-posedness, but also as a way to enforce desirable properties on the transformation. We propose a regularisation term  $\mathcal{P}^{\text{rigid}}[\mathbf{u}; I_M]$  that penalises deformations of rigid objects,

which we call the rigidity penalty term. This penalty term can be weighted locally, so that some parts of the image are restricted to rigid movement, while other parts may be penalised partially or may deform freely.

In the following sections we describe the general registration algorithm (Section 4.2.1), derive the rigidity penalty term (Section 4.2.2), and describe how this penalty term can be efficiently computed if the deformation field field is parameterised by B-splines (Section 4.2.3).

### 4.2.1 Registration Algorithm

We employ a registration framework largely based on the papers of Rueckert *et al.* [4] and Mattes *et al.* [2]. The deformation field  $\mathbf{u}$  is parameterised by cubic B-splines [4], with parameters  $\boldsymbol{\mu}$ . The similarity measure  $\mathcal{S}$  is the mutual information measure implemented according to Thévenaz and Unser [30]:

$$MI(\boldsymbol{\mu}; I_F, I_M) = \sum_{\iota \in L_M} \sum_{\kappa \in L_F} p(\iota, \kappa; \boldsymbol{\mu}) \log_2 \left( \frac{p(\iota, \kappa; \boldsymbol{\mu})}{p_M(\iota; \boldsymbol{\mu}) p_F(\kappa)} \right), \quad (4.2)$$

where  $L_F$  and  $L_M$  are the sets of histogram bin centres of the fixed and moving image, respectively,  $p$  is the joint discrete probability, and  $p_F$  and  $p_M$  are the marginal discrete probabilities. B-spline Parzen windows are used to estimate the joint probabilities:

$$p(\iota, \kappa; \boldsymbol{\mu}) = \frac{1}{|I_F|} \sum_{\mathbf{x}_i \in I_F} w_F(\kappa/\sigma_F - I_F(\mathbf{x}_i)/\sigma_F) \cdot w_M(\iota/\sigma_M - I_M(\mathbf{x}_i + \mathbf{u}_\boldsymbol{\mu}(\mathbf{x}_i))/\sigma_M), \quad (4.3)$$

with  $w_F$  and  $w_M$  the fixed and moving Parzen windows,  $\sigma_F$  and  $\sigma_M$  the histogram bin widths, and  $\mathbf{u}_\boldsymbol{\mu}$  the B-spline deformation field. The marginal probabilities are obtained by summing (4.3) over  $\kappa$  or  $\iota$ .

For the optimisation of the cost function  $\mathcal{C}$ , an iterative stochastic gradient descent optimiser is used:

$$\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}_k - a_k \frac{\partial \mathcal{C}}{\partial \boldsymbol{\mu}}, \quad (4.4)$$

where  $\frac{\partial \mathcal{C}}{\partial \boldsymbol{\mu}} = -\frac{\partial \mathcal{S}}{\partial \boldsymbol{\mu}} + \alpha \frac{\partial \mathcal{P}}{\partial \boldsymbol{\mu}}$ , and  $a_k > 0$  is the size of the step taken in the negative direction of the derivative. Thévenaz and Unser [30] derived an analytical expression for the derivative  $\partial \mathcal{S} / \partial \boldsymbol{\mu}$ . In this chapter the analytical expression is evaluated approximately. Instead of evaluating the derivative over all voxels of the fixed image, i.e. taking all  $\mathbf{x}_i \in I_F$  in Equation (4.3), only a randomly chosen subset of  $I_F$  is used. This random subset is renewed every iteration. Klein *et al.* [67] showed that, using this stochastic gradient descent optimiser, the computation time per iteration can be significantly decreased, without affecting the rate of convergence and final precision. In this chapter, a decaying function of the iteration number  $k$  is used for computing the gain factor:  $a_k = a / (k + A)^\gamma$ , where  $a > 0$ ,  $A \geq 1$ , and  $0 \leq \gamma \leq 1$  are user-defined constants. This function, and practical guidance for choosing the parameters, is suggested by Spall [56]. A multiresolution approach is taken to avoid local minima, using a Gaussian pyramid with a subsampling factor of two in each dimension.

## 4.2.2 Construction of the Rigidity Penalty Term

In this section we derive three conditions that must hold for a transformation  $\mathbf{u}(\mathbf{x}) + \mathbf{x}$  to be a rigid transformation. These three conditions are combined in one penalty term, our rigidity penalty term  $\mathcal{P}^{\text{rigid}}[\mathbf{u}; I_M]$ , constructed such that deviation from these three conditions is penalised. For the sake of clarity the penalty term is derived in 2D. The results can be readily extended to 3D, or even nD.

For a displacement field  $\mathbf{u}$  to be rigid, it must hold that

$$\mathbf{u}(\mathbf{x}) + \mathbf{x} = R\mathbf{x} + \mathbf{t}, \quad (4.5)$$

with  $R$  and  $\mathbf{t}$  a rotation matrix and a translation vector, respectively. Three conditions on  $\mathbf{u}(\mathbf{x}) + \mathbf{x}$  can be derived:

**affine** A rigid transformation is an affine function in  $\mathbf{x}$ , giving the ‘affinity’ conditions  $\text{AC}_{kij}(\mathbf{x})$ , which state that the second order derivatives of  $\mathbf{u}$  to  $\mathbf{x}$  have to be zero:

$$\text{AC}_{kij}(\mathbf{x}) = \frac{\partial^2 u_k(\mathbf{x})}{\partial x_i \partial x_j} = 0, \quad (4.6)$$

for all  $k, i, j = 1, 2$ , not counting duplicates. Rueckert *et al.*[4] penalises deviation from this constraint to enforce smoothness.

**orthonormality** For the matrix  $R$  to be a rotation matrix it must be orthonormal. This defines the orthonormality conditions  $\sum_{k=1}^2 r_{ki}r_{kj} = \delta_{ij}$ , for all  $i, j = 1, 2$ , with  $r_{ij}$  the elements of  $R$ , and  $\delta_{ij}$  the Kronecker delta function. From Equation (4.5) it follows that  $\frac{\partial u_i}{\partial x_j} = r_{ij} - \delta_{ij}$ , for all  $i, j = 1, 2$ . Hence, the orthonormality conditions  $\text{OC}_{ij}$  can be rewritten as

$$\text{OC}_{ij}(\mathbf{x}) = \sum_{k=1}^2 \left( \frac{\partial u_k(\mathbf{x})}{\partial x_i} + \delta_{ki} \right) \left( \frac{\partial u_k(\mathbf{x})}{\partial x_j} + \delta_{kj} \right) - \delta_{ij} = 0, \quad (4.7)$$

for all  $i, j = 1, 2$ , again not counting duplicates. Note that, since  $r_{ij} = \frac{\partial u_i}{\partial x_j} + \delta_{ij}$  (Equation (4.5)), the rotation matrix  $R$  is the Jacobian of the transformation  $\mathbf{u}(\mathbf{x}) + \mathbf{x}$ . Therefore, the OC term is equal to forcing orthonormality of the Jacobian, which is used in Loeckx *et al.*[48] and Ruan *et al.*[65].

**properness** A matrix  $R$  satisfying the orthonormality conditions can still be proper or improper, meaning that the determinant can still be either 1 or -1, respectively. An improper orthonormal matrix corresponds to a rotation with an inversion (mirroring). Therefore we need to impose the properness condition  $\text{PC}(\mathbf{x}) = \det(R) - 1 = 0$ , with the elements  $r_{ij}$  of the matrix  $R$  again expressed in derivatives of  $\mathbf{u}$  to  $\mathbf{x}$ . Note that, since  $R$  is the Jacobian of the transformation  $\mathbf{u}(\mathbf{x}) + \mathbf{x}$ , this condition basically amounts to an incompressibility constraint, see also [5].

We define the rigidity penalty term  $\mathcal{P}^{\text{rigid}}[\mathbf{u}; I_M]$  to be the sum of all these conditions squared. In order to distinguish between rigid and nonrigid tissue, the total penalty term is weighted by a so-called rigidity coefficient  $c(\mathbf{x}) \in [0, 1]$  of the tissue type at position  $\mathbf{x}$ . This results in the following expression:

$$\mathcal{P}^{\text{rigid}}[\mathbf{u}; I_M] \triangleq \frac{1}{\sum_{\mathbf{x}} c(\mathbf{x} + \mathbf{u}(\mathbf{x}))} \sum_{\mathbf{x}} c(\mathbf{x} + \mathbf{u}(\mathbf{x})) \left\{ \begin{aligned} & c_{\text{AC}} \sum_{k,i,j} \text{AC}_{kij}(\mathbf{x})^2 + c_{\text{OC}} \sum_{i,j} \text{OC}_{ij}(\mathbf{x})^2 + c_{\text{PC}} \text{PC}(\mathbf{x})^2 \end{aligned} \right\}. \quad (4.8)$$

The weights  $c_{\text{AC}}$ ,  $c_{\text{OC}}$ , and  $c_{\text{PC}}$  determine the relative strength of each of the three terms. The rigidity coefficient  $c(\mathbf{x})$  is set to zero for pixels  $\mathbf{x}$  in completely nonrigid tissue, and to one for rigid tissue. For other tissue types a value of  $c(\mathbf{x})$  is chosen between zero and one. The rigidity coefficient image can be constructed by performing a manual or (semi-) automatic segmentation of structures of interest, after which a rigidity coefficient can be assigned to each segment. For the case of CT images the Hounsfield units themselves may be used, rescaled to the range  $[0, 1]$ , since more rigid tissue usually has a higher attenuation value. This chapter concerns the registration method and therefore we use a simple manual segmentation. The rigidity coefficient image only has to be defined on the moving image. The moving image deforms to the fixed image, so specific parts of the moving image should transform in a rigid fashion. Therefore, the regions that correspond to rigid structures have to be defined on the moving image. When computing the derivative  $\partial \mathcal{P} / \partial \mu$  we neglect the derivative of  $c(\mathbf{x} + \mathbf{u}(\mathbf{x}))$  to  $\mathbf{u}$ , the same way as overlap changes of images are sometimes neglected in registration [30]. Hence, discontinuity of  $c(\mathbf{x})$  is not an issue here.

The following theorem states the validity of the proposed rigidity constraint. For clarity, we assume that  $c(\mathbf{x}) > 0$ , on a connected subregion  $\Omega_S$  of the moving image.

**Theorem 1.**  $\mathcal{P}^{\text{rigid}}[\mathbf{u}; I_M] = 0$  if and only if the transformation  $\mathbf{u}(\mathbf{x}) + \mathbf{x}$  is rigid, provided that  $c(\mathbf{x}) > 0, \forall \mathbf{x} \in \Omega_S \subset \Omega_M$ .

*Proof.* The *if*-part is trivial, so we prove the *only if*-part. Let  $\mathcal{P}^{\text{rigid}}[\mathbf{u}; I_M] = 0$ , which is at every location  $\mathbf{x}$  a sum of three nonnegative terms. Therefore, each of the three terms is zero. Since the first term is zero, all second order derivatives are zero, making  $\mathbf{u}(\mathbf{x})$  affine. Therefore the displacement can be written as  $\mathbf{u}(\mathbf{x}) = \tilde{R}\mathbf{x} + \mathbf{t}$ . Given that the second term is zero it holds that

$$\sum_{k=1}^2 (\tilde{r}_{ki} + \delta_{ki}) (\tilde{r}_{kj} + \delta_{kj}) = \delta_{ij}, \quad \forall i, j = 1, 2, \quad (4.9)$$

since  $\frac{\partial u_i}{\partial x_j} = \tilde{r}_{ij}$ . Now splitting up  $\tilde{R} = R - I_d$  (giving a displacement  $\mathbf{u}(\mathbf{x}) = R\mathbf{x} - \mathbf{x} + \mathbf{t}$ ), gives  $\tilde{r}_{ij} = r_{ij} - \delta_{ij}$ . Substituting this in Equation (4.9) results in  $\sum_{k=1}^2 r_{ki}r_{kj} = \delta_{ij}, \forall i, j = 1, 2$ , in which we recognise the orthonormality conditions for  $R$ . A similar argument holds for the third term, resulting in the properness condition. Therefore  $R$  is a rotation matrix, and  $\mathbf{u}(\mathbf{x}) + \mathbf{x}$  represents a rigid transformation.  $\square$

For every application the parameters  $c_{AC}$ ,  $c_{OC}$ , and  $c_{PC}$  have to be tuned. A suitable first estimate for those values is the energy of the three terms after unconstrained registration, relative to each other. The energy of AC is defined as:

$$E_{AC} = \frac{1}{\sum c(\mathbf{x} + \mathbf{u}(\mathbf{x}))} \sum_{\mathbf{x}} c(\mathbf{x} + \mathbf{u}(\mathbf{x})) \sum_{k,i,j} AC_{kij}(\mathbf{x})^2. \quad (4.10)$$

$E_{OC}$  and  $E_{PC}$  are defined similarly. We have found that a ratio in the order  $c_{AC} : c_{OC} : c_{PC} = 100 : 1 : 10$  is a good starting point. The results did not appear to be overly sensitive to the choice of the Section 4.3.

Provided that no mirroring occurs the OC term is a sufficient condition for local rigidity [48, 65]. In the experiments OC per se is explicitly compared with the complete penalty term.

### 4.2.3 Using B-splines

The proposed rigidity penalty term is not dependent on a parameterisation of the displacement field. We employ a B-spline parameterisation, because of the computational benefits. Parameterising the displacement field  $\mathbf{u}(\mathbf{x})$  by cubic B-splines yields in 2D:

$$u_i(x_1, x_2) = \sum_{l \in V_{\mathbf{x}}} \mu_{li} \beta^3(x_1 - x_{l1}) \beta^3(x_2 - x_{l2}), \quad (4.11)$$

for all  $i = 1, 2$ , with  $\beta^3(x)$  the cubic B-spline polynomial,  $\mu_{li}$  the B-spline coefficients, and  $V_{\mathbf{x}}$  the set of all control points within the compact support of the B-spline at  $\mathbf{x}$ . It is well known [36] that the derivatives of  $\mathbf{u}(\mathbf{x})$  can also be expressed in terms of the B-spline coefficients, using the rule  $\frac{d\beta^n(x)}{dx} = \beta^{n-1}(x + \frac{1}{2}) - \beta^{n-1}(x - \frac{1}{2})$ . Therefore,  $\mathcal{P}^{\text{rigid}}[\mathbf{u}, I_M]$  can be expressed in terms of the B-spline coefficients  $\mu_{li}$ . We evaluate the rigidity constraint over the control points only, which imposes local rigidity of the control point grid. Rigidity of the control points in the B-spline support region of an object guarantees rigidity of the entire object. Evaluating over the control points only, combined with the compact support of B-splines, gives us an efficient way to calculate  $\mathcal{P}^{\text{rigid}}[\mathbf{u}, I_M]$  and its derivatives, needed for gradient descent like optimisers.

In order to define a rigid transformation at some point  $\mathbf{x}$ , all control points within the compact support of the B-spline have to be kept rigid. The nonrigid deformation in the neighbourhood of a rigid structure is therefore also penalised to some extent. The precision with which a rigid region can be defined, is determined by the density of the B-spline control point grid and the order of the B-spline.

### 4.2.4 Synthetic Example

To illustrate the effect of the three terms in Equation (4.8) we have constructed a synthetic example of size 512 by 512 pixels, see Figure 4.1. The ellipsoid in the fixed image in Figure 4.1(a) is rotated and anisotropically scaled compared to the moving image. It represents a tumour that has grown. The rectangle is only rotated and represents a rigid object, such as a bone. The bar is an example of a structure visible in only one of the

images, representing a situation before and after contrast uptake in a vessel. The desired behaviour of a registration algorithm is that all three objects are kept rigid. Therefore,  $c(\mathbf{x}) = 1.0$  on the three objects and zero elsewhere. The resulting rigidity coefficient image is dilated with a kernel of radius 8.0 pixels (two control point spacings). The other parameters in Equation (4.8) for this experiment are  $\alpha = 1.0$ ,  $c_{AC} = 250.0$ ,  $c_{OC} = 2.0$ , and  $c_{PC} = 10.0$ . A B-spline control point spacing of 4.0 pixels is used.

In the remainder of this chapter the abbreviation ‘NRP’ is used for registration with the rigidity penalty term  $\mathcal{P}^{\text{rigid}}[\mathbf{u}, I_M]$  and ‘NRU’ for unconstrained nonrigid registration (only a similarity term). When only one of the conditions in Equation (4.8) is used, the method is referred to as ‘AC’, ‘OC’, or ‘PC’.

The NRU algorithm, see Figure 4.1(b), registers the ellipsoid and the rectangle, and tries to eliminate the bar, since it is not visible in the fixed image. Note that the outer border of the rectangle is registered correctly, but the deformation field within the object is not rigid at all. When only AC is applied as in Figure 4.1(c), the deformations are affine at the three objects, but not rigid, since scaling is not penalised. The orthonormality condition nicely preserves rigidity, see Figure 4.1(d). The PC term is volume preserving (Figure 4.1(e)), but volume preservation is not equal to rigidity. Finally, in Figure 4.1(f), the combination of all terms is used. Rigidity of the three objects is achieved, and the deformation field looks slightly smoother than that of OC alone. Note, that the major axes of the ellipsoids are not aligned in some cases. This is due to the fact that the similarity measure has a constant value for a range of inclination angles, as long as the smaller ellipsoid fits in the larger one entirely.

When the three objects are placed somewhat closer together OC alone is unable to reach the optimum, resulting in a failed registration (the ellipsoid and the rectangle are not rigid, the rectangles do not match completely), see Figure 4.1(g). The use of all terms gives a smoother deformation field, and also a good registration, see Figure 4.1(h). Note that when the objects are placed closer together some folding appears. This is a consequence of the use of B-splines: invertibility of the deformation field is not guaranteed. Since the rigidity penalty term is independent of the representation of the deformation field, it can also be applied in combination with a diffeomorphic transformation model.

In Table 4.1 the energies of the three components are given for the several registrations of Figure 4.1. Together they form a measure of rigidity, since Theorem 1 states that each of them is zero in case of rigidity. The three conditions help each other in achieving rigidity, OC is the most dominant of the three, and the total rigidity term yields the best results.

### 4.3 Experiments and Results

In order to evaluate the effectiveness of the rigidity penalty term (4.8), NRP was compared to NRU. Additionally, all registrations were performed with only the AC, OC or PC term as penalty. An appropriate setting for  $\alpha$  in (4.1) and for the relative strength of each of the terms of  $\mathcal{P}^{\text{rigid}}[\mathbf{u}, I_M]$ , controlled by  $c_{AC}$ ,  $c_{OC}$  and  $c_{PC}$ , was determined experimentally. Sensitivity of the results to the relative strength of each term was also investigated.






Figure 4.1: Illustrating the different parts of the penalty term. In (a) the fixed and moving image are shown, together with their difference. The other images show the result after registration using the various methods. (g) and (h) show the result of OC and NRP, respectively, when the objects are placed a little closer together.

The methods were compared on clinical data, viz. 3D CT follow-up data of the thorax containing lung tumours (Section 4.3.1), and 2D Digital Subtraction Angiography (DSA) (DSA) image data, see Section 4.3.2.

Throughout this chapter a  $32 \times 32$  joint histogram was used to estimate the mutual information. The parameters  $A$  and  $\gamma$ , needed to compute the gain factor  $a_k$  in (4.4), were set to  $A = 100.0$  and  $\gamma = 0.602$ , see [56]. All experiments were performed with software developed by the authors (<http://elastix.isi.uu.nl>). This registration package is largely based on the Insight Segmentation and Registration Toolkit [33]. The com-

Table 4.1: The energy  $E$  of the three conditions.

method	$E_{AC}$	$E_{OC}$	$E_{PC}$
NRU	$1.9 \cdot 10^{-3}$	$3.4 \cdot 10^{+1}$	$3.6 \cdot 10^0$
AC	$5.6 \cdot 10^{-6}$	$3.0 \cdot 10^{+1}$	$3.6 \cdot 10^0$
OC	$3.6 \cdot 10^{-5}$	$8.6 \cdot 10^{-3}$	$2.9 \cdot 10^{-3}$
PC	$4.8 \cdot 10^{-4}$	$6.3 \cdot 10^{-1}$	$2.9 \cdot 10^{-4}$
NRP	$1.4 \cdot 10^{-6}$	$1.1 \cdot 10^{-3}$	$1.3 \cdot 10^{-4}$

putation time for the rigidity penalty term scales linearly with the number of B-spline parameters, and is about 0.06 s for 10 000 parameters in 2D and about 3 s for 100 000 parameters in 3D, on a standard computer (AMD Opteron 250 running on 2.4 GHz). Calculation of the derivative of the mutual information in 2D and 3D requires 0.11 s and 1 s, respectively.

### 4.3.1 3D CT Thorax Data With Lung Tumours

A possible way to follow disease progress over time is to visually inspect the difference between a first scan, taken at  $t_0$ , and a registered second follow-up scan, taken at  $t_1$ . For the case of patients with lung tumours, this difference will indicate tumour growth between two scans. Because of differences in lung inspiration levels, rigid registration is not sufficient to achieve good alignment of the anatomy. However, standard nonrigid registration methods will minimise the difference between tumours at different time points, effectively concealing tumour growth, see Figure 4.2(b). Therefore, the tumours should be considered rigid tissue by the nonrigid registration algorithm. Rigidity is evaluated with tumour volume measurements.

#### Data Description

Registration was performed on CT follow-up data sets of the thorax of five patients having lung tumours. For each patient two or three images of different time points were available, such that in total seven registrations were performed. The data were acquired with a Philips 16-slice spiral CT scanner (Mx8000 IDT 16). The images have slices of 512 by 512 voxels. The number of slices varies for the data sets, ranging from 400 to 550. The in-plane voxel size is around  $0.7 \times 0.7$  mm. Slice thickness is always 1.0 mm and slices were reconstructed every 0.7 mm. Before registration, each data set was downsampled with a factor of two in each dimension to reduce computer memory and computational load. Downsampling was performed by discarding odd rows, columns and slices. The five data sets contain thirty tumours in total at time  $t_0$ , with an average volume of 2.5 ml for the first scan  $t_0$  and 5.1 ml for the follow-up  $t_1$ . No new tumours had developed at  $t_1$ .




Figure 4.2: Comparison of NRU and NRP for a slice taken from 3D CT thorax images. The tumours, located within the box (see (d)), are to be kept rigid. (a) and (d): CT slice at time  $t_0$  (the fixed image) and time  $t_1$  (the moving image), respectively; (b) and (e): difference of the registration result with the fixed image, for NRU and NRP, respectively; (c) and (f): the resulting deformation field near the tumours for NRU and NRP, respectively.

### Experiment Setup

The CT image taken at time  $t_0$  was set to be the fixed image. The CT image taken at time  $t_1$  was used as the moving image. A coarse alignment between fixed and moving image was obtained by a rigid registration. For the rigid registration three resolutions were used, for the nonrigid registration four. For nonrigid registration the resolution of the B-spline grid is adapted each resolution: for the four resolutions the grid spacing was set to 64, 32, 16, and 8 voxels, respectively. The number of optimisation iterations for the first three resolutions was 300, for the last one 1200. For every iteration 5000 samples were used to calculate the derivative of the mutual information. The parameter  $a$  in the gain factor (see below (4.4)) was set to 150000, 120000, 70000 and 20000, for the four resolutions, respectively. For the nonrigid registrations using the rigidity penalty term, a crude manual segmentation of the tumours was used to define  $c(\mathbf{x})$ , setting  $c(\mathbf{x})$  to 1.0 for voxels within the tumour and to 0.0 elsewhere. The weight  $\alpha$  was set to 4.0, and the weights of the three terms were chosen  $c_{AC} = 100.0$ ,  $c_{OC} = 1.0$ , and  $c_{PC} = 2.0$ . In the final resolution the rigidity coefficient image was dilated with a radius of 16 voxels (two B-spline control points) to achieve (almost) complete rigidity of the tumours, as

Table 4.2: Quantitative results for the 3D CT thorax data. The last five columns show the geometric mean and standard deviation of the tumour volume ratios, grouped according to true tumour growth  $v_{t_1}/v_{t_0}$ . The number of lung or tumours segmentations the means and standard deviations are based on, is indicated by  $n$ .

method	lung overlap	$\frac{v_{t_1}}{v_{t_0}}$ : all	$\frac{v_{t_1}}{v_{t_0}}$ : (0, 1]	$\frac{v_{t_1}}{v_{t_0}}$ : (1, $\frac{3}{2}$ ]	$\frac{v_{t_1}}{v_{t_0}}$ : ( $\frac{3}{2}$ , 3]	$\frac{v_{t_1}}{v_{t_0}}$ : (3, $\infty$ )
		$r_{reg}$	$r_{reg}$	$r_{reg}$	$r_{reg}$	$r_{reg}$
rigid	0.93 $\pm$ 0.05	0.99 $\times$ /1.05	0.99 $\times$ /1.01	0.96 $\times$ /1.08	0.99 $\times$ /1.03	1.02 $\times$ /1.05
NRU	0.97 $\pm$ 0.01	0.78 $\times$ /1.24	0.96 $\times$ /1.04	0.93 $\times$ /1.06	0.73 $\times$ /1.15	0.69 $\times$ /1.34
AC	0.98 $\pm$ 0.01	0.89 $\times$ /1.16	1.04 $\times$ /1.05	1.01 $\times$ /1.11	0.85 $\times$ /1.12	0.84 $\times$ /1.16
OC	0.99 $\pm$ 0.01	0.96 $\times$ /1.08	1.02 $\times$ /1.03	1.00 $\times$ /1.06	0.96 $\times$ /1.09	0.91 $\times$ /1.04
PC	0.98 $\pm$ 0.01	0.95 $\times$ /1.06	0.97 $\times$ /1.03	1.00 $\times$ /1.04	0.94 $\times$ /1.06	0.91 $\times$ /1.03
NRP	0.99 $\pm$ 0.01	0.98 $\times$ /1.05	1.02 $\times$ /1.04	1.00 $\times$ /1.05	0.97 $\times$ /1.03	0.95 $\times$ /1.06
$n$	7	30	3	6	13	8

discussed in Section 4.2.3.

## Results

As can be seen from Figure 4.2(b) and 4.2(c), NRU fails to keep the tumours rigid. Therefore, the difference in size due to growth can not be appreciated. Using a crude segmentation of the tumours, NRP succeeds in keeping the tumours rigid, see the deformation field in Figure 4.2(f). From the difference image in Figure 4.2(e) it is immediately clear that the tumours have grown, whereas the rest of the image was registered with equal accuracy as for the unconstrained nonrigid registration.

Accuracy of the registration was measured by calculating the lung overlap of the registered image with the fixed image. For this purpose, automatic lung segmentations were made with an algorithm based on the method by Hu *et al.* [58], described in detail in [59]. The overlap measure is defined as

$$\text{overlap} \triangleq \frac{2|L_1 \cap L_2|}{|L_1| + |L_2|}, \quad (4.12)$$

where  $L_i$  is the set of all voxels within the lung, and  $|L_i|$  denotes the size of set  $L_i$ . The averages and standard deviations of lung overlap were calculated over all data, and are reported in Table 4.2. The results show that good lung overlap was achieved with all nonrigid registration algorithms. This is confirmed by visual inspection of the results.

For evaluation of the rigidity of the tumours, precise manual segmentations of the tumours were used. Tumour volume measurements were performed to see if the registration is volume preserving, a condition for rigidity. Tumour volume before registration at  $t_1$  is denoted by  $v_{t_1}$ ; tumour volume after registration with algorithm *reg* is denoted by  $v_{reg}$ . The volume after registration was computed by applying the resulting transformation to the manual segmentation of the tumour in the moving image, and subsequently calculating the volume of the transformed segmentation. Volume preservation is expressed by the

Table 4.3: The arithmetic means and standard deviations of the energy  $E$  of the three conditions for the 3D CT thorax data.

method	$E_{AC}$	$E_{OC}$	$E_{PC}$
NRU	$1.6 \cdot 10^{-4} \pm 8.3 \cdot 10^{-5}$	$1.4 \cdot 10^{-1} \pm 1.4 \cdot 10^{-1}$	$2.8 \cdot 10^{-2} \pm 3.1 \cdot 10^{-2}$
AC	$1.5 \cdot 10^{-5} \pm 1.7 \cdot 10^{-5}$	$9.7 \cdot 10^{-2} \pm 1.1 \cdot 10^{-1}$	$2.3 \cdot 10^{-2} \pm 3.8 \cdot 10^{-2}$
OC	$2.8 \cdot 10^{-4} \pm 5.1 \cdot 10^{-4}$	$7.7 \cdot 10^{-3} \pm 1.2 \cdot 10^{-2}$	$2.2 \cdot 10^{-2} \pm 5.6 \cdot 10^{-2}$
PC	$7.1 \cdot 10^{-3} \pm 1.1 \cdot 10^{-2}$	$6.7 \cdot 10^0 \pm 1.1 \cdot 10^1$	$1.5 \cdot 10^{-3} \pm 1.3 \cdot 10^{-3}$
NRP	$5.1 \cdot 10^{-6} \pm 6.9 \cdot 10^{-6}$	$5.2 \cdot 10^{-3} \pm 7.6 \cdot 10^{-3}$	$3.6 \cdot 10^{-4} \pm 6.2 \cdot 10^{-4}$

ratio  $r_{reg}$  between the tumour volume after registration and at  $t_1$  (since  $t_1$  is the moving image):  $r_{reg} = v_{reg}/v_{t_1}$ . If a nonrigid registration is volume preserving for a tumour, then  $v_{reg} = v_{t_1}$ , and  $r_{reg} = 1.0$ ; if a tumour is compressed  $r_{reg} < 1.0$ . For ratios it is better to use the *geometric mean*  $\mu_g$  and the *geometric standard deviation*  $\sigma_g$ , instead of their arithmetic counterparts. This can be easily seen from a small example. Say, we have two ratios 0.5 and 2.0. The arithmetic mean of those two ratios is 1.25, whereas the geometric mean equals 1.0, rating a two time increase in volume equal to a two time decrease. From the definition of the geometric standard deviation it follows that  $\sigma_g \geq 1$ .

The geometric mean volume ratios and standard deviations are reported in Table 4.2, where the symbol  $\times/$  is used to indicate the distinction with the arithmetic mean and standard deviation. Geometric means were calculated for four volume ratio groups and for all ratios together. The tumours were grouped according to true tumour growth  $v_{t_1}/v_{t_0}$ . The 6th column in Table 4.2, for example, is about the group of tumours with true tumour growth between  $3/2$  and 3.

It can be appreciated from Table 4.2 that volume was much better preserved when applying the rigidity penalty term, compared to unconstrained nonrigid registration. To evaluate the difference in rigidity between NRP and NRU, a statistical test was performed on the logarithm of  $r_{reg}$ . Wilcoxon signed-ranks tests<sup>1</sup> show that NRP is significantly ( $p < 0.001$ ,  $n = 30$ ) more rigid than NRU. Volume is by definition preserved for rigid registration. However, the corresponding column in Table 4.2 shows that the volume measurements do *not* result in the perfect value  $1.0 \times/ 1.0$ . This indicates that part of the residual volume difference for all methods can be explained by interpolation artifacts due to resampling when computing  $v_{reg}$ . For the nonrigid registration methods, the residual volume difference can partially also be explained by the aforementioned balance (Section 4.2) between the intensity-based similarity measure, which yields a force that keeps compressing the tumour, and the rigidity penalty term. This balance can of course be influenced by the parameter  $\alpha$ . From Table 4.2 it is observed that the separate conditions each are less capable of preserving volume than the combined penalty term. The OC term is again the most dominant term, and performs even better than the PC that penalises volume changes. The AC term is not volume preserving, since it does not

<sup>1</sup>The differences of NRP with NRU, are not normally distributed ( $p < 0.0001$  for Shapiro-Wilk tests). Therefore, we employ the Wilcoxon signed-ranks test.

penalise scalings. In Table 4.3, the energies of the different components of the penalty term are reported. It shows that when only one of the conditions is used, the energies of the other terms also decline. The three terms aid each other in achieving rigidity. To evaluate the difference between NRP and OC for rigidity and smoothness (i.e.  $E_{AC}$ ), a statistical test was performed. For rigidity the test was again performed on the logarithm of  $r_{reg}$ . Wilcoxon signed-ranks tests (the data is not normally distributed:  $p < 0.0001$  for Shapiro-Wilk tests) show that NRP is significantly ( $p < 0.01$ ,  $n = 30$ ) more rigid than OC; no significant difference was found between the  $E_{AC}$ 's of NRP and OC ( $p = 0.016$ ,  $n = 7$ ). We conclude that nonrigid registration with the total penalty term NRP is best capable of achieving rigidity.

To investigate the sensitivity of the results to the relative importance of each of the three terms of  $\mathcal{P}^{rigid}[\mathbf{u}, I_M]$ , the experiments were repeated while varying the weights  $c_{AC}$ ,  $c_{OC}$  and  $c_{PC}$ . Each of the weights was subsequently increased and decreased by a factor of two, resulting in six additional experiments. The results are given in Table 4.4, where the three numbers after NRP refer to the setting of the parameters  $c_{AC}$ ,  $c_{OC}$  and  $c_{PC}$ , respectively. For example NRP:  $1\frac{1}{2}1$  refers to a registration with the parameter  $c_{OC}$  chosen half the original value of 1.0. All other parameter settings were taken equal to NRP above. A Wilcoxon signed-ranks test was used to investigate the difference with NRP in terms of rigidity and smoothness ( $E_{AC}$ ). No significant difference for rigidity, at a confidence level of  $p = 0.01$ , can be found between NRP and its perturbations, except for NRP:  $11\frac{1}{2}$ , which is significantly less rigid. No significant difference is found for NRP:  $11\frac{3}{4}$ . At the same confidence level, no significant difference was found regarding smoothness. We conclude that the results are quite robust to the weights of the AC, OC and PC terms relative to each other.

### 4.3.2 Digital Subtraction Angiography

Digital Subtraction Angiography (DSA) is an established modality for visualising blood vessels in the human body. During image acquisition patient motion often occurs, due to breathing, heart beat, activity in the intestines, or movement of the body. This motion results in artifacts in the subtraction images. In Figure 4.3(c) an example of a subtraction image is shown. Rigid registration is not sufficient in these cases, as shown by the difference image after rigid registration, Figure 4.3(d), since 3D rigid motion cannot always be corrected with a 2D rigid transformation of the projection image, and since patient motion is usually of a nonrigid nature.

Typically, in DSA imaging a sequence of images is taken, where different parts of the vasculature are visible at different times. To see the entire imaged vasculature, all the images from this sequence have to be registered to some fixed image. The first image, acquired just before the arrival of the contrast bolus and known as the baseline image, was used as the fixed image. As reported in the literature [5, 46], nonrigid registration of images containing contrast-enhanced structures can lead to significant compression of those structures. See Figure 4.3(e) for an example of this behaviour. Switching the fixed and the moving image and simply using NRU does not guarantee that compression is avoided. Another disadvantage is that the inverse of the transformation is required.

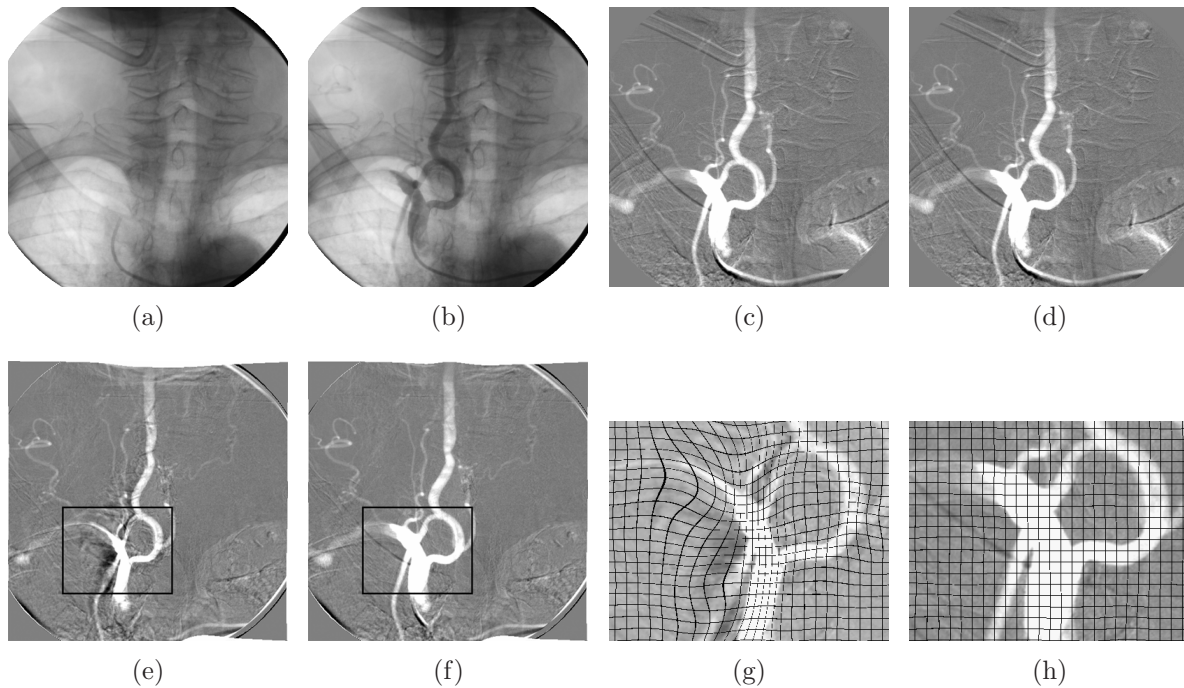


Figure 4.3: Comparison of different registration algorithms for 2D DSA images. (a) DSA baseline image: the fixed image, (b) DSA image after injection of the contrast bolus: the moving image, (c) - (f) are difference images of the fixed image, (c) with the moving image, (d) with the result of rigid registration, (e) with the result of NRU, (f) with the result of NRP. (g) and (h) depict parts of the resulting deformation field for NRU and NRP, respectively. The rectangle in Figure (e) and (f) denotes the part of the deformation field that is depicted.

Therefore, a nonrigid registration is required that can treat the vasculature differently from other tissue, to maintain vasculature size.

## Data Description

Two-dimensional digital X-ray angiography image data were acquired with an Integris V3000 C-arm imaging system (Philips). In total, twenty-six image sequences of twelve different patients were obtained. The image sequences are of size 512 by 512 pixels for twenty-two data sets, and 1024 by 1024 pixels for four data sets; they contain about ten images each. Intensities in the DSA images range approximately from 100 to 950, with an arithmetic mean and standard deviation of about  $550 \pm 180$ . Images were taken of different locations in the body: abdomen (10), brain (5), hip and foot (4), heart (1), neck (5) and lungs (1). The first image in each sequence was taken before arrival of the contrast bolus, the following images each show a part of the vasculature.

## Experiment Setup

The baseline image was taken to be the fixed image. For our experiments one image from each sequence was registered to its fixed image. The image showing the most vasculature was manually selected. To get a coarse alignment between fixed and moving image, a rigid registration was performed prior to nonrigid registration. For nonrigid registration, two resolutions were used, with a B-spline grid spacing of 16 pixels for both resolution levels. For both resolutions, 600 iterations were used. For every resolution 5000 samples were used to calculate the derivative of the mutual information. The parameter  $a$  was set to 6000.0 and 3500.0, for the respective resolutions. For the rigidity penalty term, a crude manual segmentation of the vessels was used to define  $c(\mathbf{x})$ , setting  $c(\mathbf{x})$  to 1.0 for voxels within the vasculature and to 0.0 otherwise. The weight  $\alpha$  was set to 8.0, and the weights of the three terms were chosen  $c_{AC} = 250.0$ ,  $c_{OC} = 1.0$ , and  $c_{PC} = 10.0$ . The rigidity coefficient image was dilated with a radius of 16.0 pixels to achieve complete rigidity of the vasculature.

## Results

The Root Mean Square Difference (RMSD) of the background was calculated to verify that the nonrigid registration indeed reduces motion artifacts. The background is defined as everything within the cone beam, but outside the manual vessel segmentation. Arithmetic means and standard deviations of the RMSD were calculated for all 26 images. The results are reported in the second column of Table 4.5 and show that rigid registration reduces the motion artifacts only slightly. All nonrigid registration methods improve on that.

In order to further compare the reduction in motion artifacts between the different registration methods, the difference of the RMSDs is calculated as  $\mathcal{D}_{i,j} = \text{RMSD}_i - \text{RMSD}_j$ , where  $i \neq j$  are the different registration methods. This difference was calculated for several combinations of registration methods and for all DSA data. A statistical  $t$ -test was performed to test the hypothesis that the difference has a zero mean. For the difference  $\mathcal{D}_{\text{rigid}, \text{noreg}}$  this hypothesis is rejected ( $p < 0.01$ ,  $n = 26$ ). The difference of any of the nonrigid registration algorithms with rigid registration was also significant ( $p < 0.005$ ). No significant difference was found between NRU and the several nonrigid registration algorithms with a penalty term ( $p > 0.2$ ). It is clear that rigid registration reduces the motion artifacts slightly, compared to no registration. Nonrigid registration improves substantially on rigid registration.

Visual inspection confirms the reduction in motion artifacts, as measured by the RMSD. An example of this is shown in Figure 4.3, where the difference images before registration, after rigid registration, and after NRU and NRP are shown. Compression is clearly visible for NRU (Figure 4.3(e)), whereas the vessels for NRP in Figure 4.3(f) are similar to the image without registration, Figure 4.3(c). The deformation fields in Figure 4.3(g) and 4.3(h) show that NRU compresses the vessels, whereas the rigidity penalty term preserves rigidity.

Rigidity of the vasculature was evaluated by manually measuring the vessel diameter  $vd$  at several locations, see Figure 4.4 for an example. Six locations were selected for






Figure 4.4: An example of six locations where the vessel diameter is measured.

each of the 26 images, yielding a total of 156 diameter measurements. Similar to the tumour volumes in the previous section, ratios  $r$  are used to evaluate vessel compression. The vessel diameter after registration is compared to the diameter before registration:  $r_{reg} = vd_{reg}/vd_{noreg}$ . The geometric mean and standard deviation of the vessel diameter ratios are reported in the third column of Table 4.5. Unconstrained nonrigid registration (NRU) severely compresses the vasculature, which is avoided with the use of  $\mathcal{P}^{rigid}[\mathbf{u}, I_M]$  (NRP). A Wilcoxon signed-ranks test confirms that this difference is significant ( $p < 0.001$ ,  $n = 156$ ). The three conditions each preserve the vessel diameter to a large extent. The AC term performs slightly worse than the other two, and the combination again gives a slight improvement, both in mean and standard deviation. Judging from the last three columns in Table 4.5 the three conditions help each other in achieving rigidity. Again, a Wilcoxon signed-ranks test was performed (the data is not normally distributed:  $p < 0.0001$  for Shapiro-Wilk tests) on the difference  $\log r_{NRP} - \log r_{OC}$ , and on the difference in smoothness ( $E_{AC}$ ). A significant difference is found for both rigidity ( $p < 0.01$ ,  $n = 156$ ) and smoothness ( $p < 0.01$ ,  $n = 26$ ), showing that the combination of all terms performs best.

Experiments were repeated with  $c_{AC}$ ,  $c_{OC}$  and  $c_{PC}$  halved and doubled. The results are shown in the last six rows of Table 4.5. Changing the relative importance of the three terms does not change the vessel rigidity significantly (using a Wilcoxon signed-ranks test). Smoothness is sometimes significantly different from NRP, see column  $E_{AC}$  in Table 4.5, but is still in the order of  $10^{-7}$ .

## 4.4 Conclusions and Discussion

We have proposed a novel method to keep user-defined structures locally rigid, while performing nonrigid registration. This is achieved by including a rigidity penalty term, derived from three rigidity conditions, in the registration cost function.

The method was evaluated on 2D and 3D clinical data. It was compared against

an unconstrained nonrigid registration approach. From the experiments on CT thorax follow-up data it is observed that tumour volume is preserved when applying the rigidity penalty term, in contrast with unconstrained nonrigid registration. The results on the DSA data show that vessel width is also much better retained with the rigidity penalty term.

The rigidity penalty term consists of three terms: an affinity, an orthonormality, and a properness condition. If used separately, both the properness and the orthonormality condition are to a large extent capable of achieving volume preservation. The affinity condition aids in regularising the problem. The affinity and properness conditions are by themselves not capable of achieving rigidity. The orthonormality condition is the most dominant of the three terms. The three terms aid each other in achieving rigidity, giving better results than a single condition. Also, the results are quite robust to changing the relative importance of each of the terms. In conclusion, the combination of all terms is the best choice.

Unlike the algorithm described by Tanner *et al.* [46], the proposed method does not rely on the assumption that the rotation of the rigid object was captured by the initial registration. The rigidity penalty term is suitable for any transformation capable of modelling locally rigid transformations, and does therefore not rely on particular basis functions to describe the transformation [49], or on its specific design [35]. Compared to approaches that only take into account the orthonormality of the Jacobian of the transformation [48, 65], we have shown that the proposed method gives better results in terms of achieving rigidity and smoothness of the resulting deformation field. Since a B-spline parameterisation of the deformation field is used, the ability to model local deformations is limited by the B-spline control point spacing. Therefore, segmentation errors in the rigidity coefficient image are a minor problem, as long as they remain smaller than approximately half the control point spacing.

We conclude that nonrigid registration using the proposed rigidity penalty term is capable of nonrigidly aligning images, while keeping user-defined structures locally rigid.

Table 4.4: Sensitivity of the CT results to the relative importance of the three terms. (+) or (-) means that the method is significantly more respectively less rigid or smooth than NRP using a Wilcoxon signed-ranks test at confidence level  $p = 0.01$ . (=) means that no significant difference is found.

method	lung overlap	$r_{reg}$ : all	$E_{AC}$	$E_{OC}$	$E_{PC}$
NRP	$0.99 \pm 0.01$	$0.98 \times / 1.05$	$5.1 \cdot 10^{-6} \pm 6.9 \cdot 10^{-6}$	$5.2 \cdot 10^{-3} \pm 7.6 \cdot 10^{-3}$	$3.6 \cdot 10^{-4} \pm 6.2 \cdot 10^{-4}$
NRP: $\frac{1}{2}11$	$0.97 \pm 0.02$	$0.97 \times / 1.06$ (=)	$9.1 \cdot 10^{-6} \pm 1.2 \cdot 10^{-5}$ (=)	$5.4 \cdot 10^{-3} \pm 7.8 \cdot 10^{-3}$ (=)	$3.9 \cdot 10^{-4} \pm 6.5 \cdot 10^{-4}$
NRP: $1\frac{1}{2}1$	$0.97 \pm 0.02$	$0.98 \times / 1.05$ (=)	$5.7 \cdot 10^{-6} \pm 7.0 \cdot 10^{-6}$ (=)	$9.9 \cdot 10^{-3} \pm 1.3 \cdot 10^{-2}$ (=)	$4.2 \cdot 10^{-4} \pm 6.7 \cdot 10^{-4}$
NRP: $11\frac{1}{2}$	$0.97 \pm 0.02$	$0.96 \times / 1.05$ (-)	$4.8 \cdot 10^{-6} \pm 6.2 \cdot 10^{-6}$ (=)	$5.4 \cdot 10^{-3} \pm 7.8 \cdot 10^{-3}$ (=)	$6.7 \cdot 10^{-4} \pm 1.2 \cdot 10^{-3}$
NRP: 211	$0.97 \pm 0.02$	$0.97 \times / 1.06$ (=)	$3.1 \cdot 10^{-6} \pm 4.2 \cdot 10^{-6}$ (=)	$5.1 \cdot 10^{-3} \pm 7.5 \cdot 10^{-3}$ (=)	$3.8 \cdot 10^{-4} \pm 6.5 \cdot 10^{-4}$
NRP: 121	$0.97 \pm 0.02$	$0.98 \times / 1.05$ (=)	$4.1 \cdot 10^{-6} \pm 5.9 \cdot 10^{-6}$ (=)	$2.5 \cdot 10^{-3} \pm 3.9 \cdot 10^{-3}$ (=)	$2.7 \cdot 10^{-4} \pm 4.8 \cdot 10^{-4}$
NRP: 112	$0.97 \pm 0.02$	$0.97 \times / 1.06$ (=)	$5.3 \cdot 10^{-6} \pm 7.1 \cdot 10^{-6}$ (=)	$5.1 \cdot 10^{-3} \pm 7.4 \cdot 10^{-3}$ (=)	$1.7 \cdot 10^{-4} \pm 2.8 \cdot 10^{-4}$

Table 4.5: Quantitative results for the DSA data. Arithmetic means and standard deviations of the RMSD are displayed in the second column. Geometric means and standard deviations of the vessel diameter ratios are shown in the third. The remaining columns show the average values of the three conditions after registration, together with their standard deviation. The last six rows show the sensitivity of the DSA results to the relative importance of the three terms. (+) or (-) means that a method is significantly more respectively less rigid or smooth than NRP using a Wilcoxon signed-ranks test at confidence level  $p = 0.01$ . (=) means that no significant difference is found.

method	RMSD	$r_{reg}$	$E_{AC}$	$E_{OC}$	$E_{PC}$
noreg	14.02 ± 5.97				
rigid	13.53 ± 5.73	1.00×/1.00			
NRU	11.94 ± 4.10	0.84×/1.17	$3.7 \cdot 10^{-4} \pm 2.6 \cdot 10^{-4}$	$1.9 \cdot 10^{-1} \pm 1.2 \cdot 10^{-1}$	$4.3 \cdot 10^{-2} \pm 2.6 \cdot 10^{-2}$
AC	12.06 ± 4.38	0.96×/1.05	$3.5 \cdot 10^{-6} \pm 2.2 \cdot 10^{-6}$	$1.8 \cdot 10^{-2} \pm 1.1 \cdot 10^{-2}$	$3.9 \cdot 10^{-3} \pm 2.3 \cdot 10^{-3}$
OC	12.07 ± 4.57	0.98×/1.04	$3.0 \cdot 10^{-5} \pm 3.0 \cdot 10^{-5}$	$6.7 \cdot 10^{-4} \pm 2.8 \cdot 10^{-4}$	$1.1 \cdot 10^{-4} \pm 5.3 \cdot 10^{-5}$
PC	12.13 ± 4.44	0.98×/1.07	$1.3 \cdot 10^{-4} \pm 1.3 \cdot 10^{-4}$	$2.5 \cdot 10^{-2} \pm 2.2 \cdot 10^{-2}$	$1.4 \cdot 10^{-5} \pm 1.6 \cdot 10^{-5}$
NRP	12.18 ± 4.66	0.99×/1.02	$2.5 \cdot 10^{-7} \pm 1.2 \cdot 10^{-7}$	$3.1 \cdot 10^{-4} \pm 1.5 \cdot 10^{-4}$	$7.2 \cdot 10^{-6} \pm 3.7 \cdot 10^{-6}$
NRP: $\frac{1}{2}11$	12.17 ± 4.66	0.99×/1.02 (=)	$5.0 \cdot 10^{-7} \pm 2.5 \cdot 10^{-7}$ (-)	$3.3 \cdot 10^{-4} \pm 1.6 \cdot 10^{-4}$	$7.6 \cdot 10^{-6} \pm 3.9 \cdot 10^{-6}$
NRP: $1\frac{1}{2}1$	12.15 ± 4.64	1.00×/1.02 (=)	$3.7 \cdot 10^{-7} \pm 1.7 \cdot 10^{-7}$ (-)	$6.3 \cdot 10^{-4} \pm 3.0 \cdot 10^{-4}$	$8.9 \cdot 10^{-6} \pm 5.2 \cdot 10^{-6}$
NRP: $11\frac{1}{2}$	12.17 ± 4.66	0.99×/1.02 (=)	$2.6 \cdot 10^{-7} \pm 1.3 \cdot 10^{-7}$ (=)	$3.3 \cdot 10^{-4} \pm 1.6 \cdot 10^{-4}$	$1.7 \cdot 10^{-5} \pm 8.3 \cdot 10^{-6}$
NRP: 211	12.19 ± 4.67	0.99×/1.02 (=)	$1.4 \cdot 10^{-7} \pm 6.7 \cdot 10^{-8}$ (+)	$2.9 \cdot 10^{-4} \pm 1.5 \cdot 10^{-4}$	$7.2 \cdot 10^{-6} \pm 3.7 \cdot 10^{-6}$
NRP: 121	12.21 ± 4.69	1.00×/1.02 (=)	$1.9 \cdot 10^{-7} \pm 1.0 \cdot 10^{-7}$ (+)	$1.4 \cdot 10^{-4} \pm 7.7 \cdot 10^{-5}$	$5.9 \cdot 10^{-6} \pm 3.0 \cdot 10^{-6}$
NRP: 112	12.19 ± 4.67	1.00×/1.02 (=)	$2.7 \cdot 10^{-7} \pm 1.5 \cdot 10^{-7}$ (=)	$3.1 \cdot 10^{-4} \pm 1.6 \cdot 10^{-4}$	$2.8 \cdot 10^{-6} \pm 1.6 \cdot 10^{-6}$

# Image Subtraction for Change Detection in Ground-Glass Opacities in Chest CT

# 5

---

## Abstract

---

**Purpose:** To study the impact of image subtraction of registered images on the detection of change in pulmonary ground-glass nodules, identified on chest CT scans of a lung screening trial.

**Materials and Methods:** A cohort of 33 individuals (25 men, 8 women; age range 51 to 75 years) with 37 focal ground-glass opacities (GGO) were recruited from a lung cancer screening. For every participant, at least one follow-up scan was available (total number of pairs, 84). Pairs of scans of the same nodule were registered nonrigidly, to generate a subtraction. Four observers rated size and density change of the GGO between pairs of scans by visual comparison alone and with availability of an additional subtraction image and indicated their confidence. An independent experienced chest radiologist established a standard of reference for nodule change having all reader data, clinical data and follow-up examinations available. Exclusion of nodule pairs for which the standard of reference could not be securely established, left 59 and 58 pairs for evaluation of size and density change, respectively. Weighted kappa statistics ( $\kappa_w$ ) were used to assess inter-observer agreement and agreement with the standard of reference. Statistical significance was tested with a kappa  $z$ -test.

**Results:** When the subtraction image was available, the average inter-observer  $\kappa_w$  improved from 0.52 to 0.66 for size change and from 0.47 to 0.57 for density change. Average agreement with the standard of reference improved from 0.61 to 0.76 for size change and from 0.53 to 0.64 for density change. The effect was especially seen when observer confidence without the subtraction image was low: agreement improved from 0.26 to 0.57 and from 0.19 to 0.47 in those cases.

**Conclusion:** Image subtraction improves the detection of subtle changes in pulmonary ground-glass opacities, and decreases inter-observer variability.

---

Based upon: M. Staring, J.P.W. Pluim, B. de Hoop, S. Klein, B. van Ginneken, H. Gietema, G. Nossent, C. Schaefer-Prokop, S. van de Vorst, and M. Prokop, "Influence of Image Subtraction on the Assessment of Volume and Density Change in Ground-Glass Opacities in Chest CT", *submitted*.

## 5.1 Introduction

**D**IFFERENTIATION of benign and malignant nodules in lung cancer screening trials [41,68–70] is largely based on the detection of subtle density and size changes after a short period of time. Various software tools are available for automatic or semiautomatic volumetry. Multiple studies could prove the superiority of those software tools over manual measurements of diameters [71] and visual assessment. Until now these software tools are only approved for solid nodules and not for ground-glass or part solid nodules.

The prognostic significance of these so-called ground-glass opacities (GGOs) is not yet fully understood. They often represent atypical adenomatous hyperplasia, bronchoalveolar carcinoma (BAC), or minimally invasive adenocarcinoma [72–74]. The Early Lung Cancer Action Project (ELCAP) group reported a malignancy rate of 34% in nonsolid lesions and a 63% malignancy rate in part-solid nodules during their lung cancer screening trial [75]. The malignancy type in the part-solid or nonsolid nodules was predominantly BAC or adenocarcinoma with bronchoalveolar features [75].

The decision to treat a patient with a GGO depends on the level of suspicion for malignancy. Current ACCP guidelines [76] recommend that a surgical biopsy is used to establish a histopathologic diagnosis in all patients with suspected bronchoalveolar carcinoma (BAC), a carcinoma known to frequently present as a GGO or part-solid nodule. The two most important criteria for increasing suspicion and further invasive work-up are change in size, and appearance of a solid component. In case of a BAC, the solid part visible on CT represents the collapse of the alveoli, subsequent formation of fibrotic focus and proliferation of tumor cells. Therefore, the presence of a solid part within the nodule on CT generally indicates invasive growth [77]. Consequently, the early detection of size and density change of a GGO may result in an early detection of cancer, and is commonly taken as an indication for surgical resection [77–79]. Nakamura et al. reported that early detection and treatment of GGOs improves the prognosis of lung cancer [80].

Detection of growth or density change of GGOs, however, is frequently difficult, especially when the lesion has a complex shape or when the position of the lesion varies between scans due to different levels of inspiration. While techniques for automated volume measurements have been extensively studied for solid nodules [81–84], such routines are just emerging for nonsolid GGOs [85]. The current clinical practice to detect change of GGOs is a side-by-side visual comparison of follow-up scans, aided by manual diameter measurements.

In this chapter we study an alternative approach that provides a subtraction image along with the original images to more readily appreciate changes. To reduce artifacts in the subtraction image due to motion, and to get a meaningful subtraction image, the scans were registered before subtraction.

The goal of this study was to investigate whether subtraction of registered images improves detection of subtle changes in pulmonary ground-glass opacities.

## 5.2 Materials and Methods

### 5.2.1 Study Participants

A cohort of 33 individuals (25 men, 8 women; age range 51 - 75 years; mean 59 years) with 37 focal ground-glass lesions (29 participants with 1 nodule, 4 with 2; 24 nonsolid and 13 part-solid nodules) were recruited from a lung cancer screening study (the NELSON study [41,86]). The trial was approved by the Dutch Ministry of Health and by the institutional review board of each participating center. Written informed consent had been obtained from each participant. All participants smoked a minimum of 16 cigarettes/day for 25 years or 11 cigarettes/day for 30 years. Participants that were scanned between April 2004 and February 2007, with a GGO visible in at least two consecutive scans, were included in our study. A ground-glass opacity was defined as a focal area of increased lung opacity that did not distort or obscure the underlying lung markings (bronchi and lung vasculature). The nodule locations were determined in the NELSON study, and were available for this study. Following these criteria we included 16 patients with 2 scans, 14 with 3, and 3 with 4 scans; a total of 86 scans. The time between the follow-ups was typically 3 or 9 months (mean: 6 months), with a minimum of 2 and a maximum of 13 months.

Mean initial nodule size as determined by manually measuring the maximum nodule diameter in sagittal, coronal or transversal direction was 14 mm (standard deviation: 6.7 mm; range: 6-34 mm; 11 in the range 5-10 mm, 21 in 10-20 mm, 5 nodules  $\geq$  20 mm).

### 5.2.2 CT Imaging

Imaging data were acquired with low-dose CT, using 16-detector-row spiral CT scanners (Mx8000 IDT or Brilliance 16P, Philips Medical Systems, Cleveland, OH). Data were acquired in approximately 12 seconds in helical mode with  $16 \times 0.75$  mm collimation and 15 mm table feed per rotation (pitch 1.3). A caudo-cranial scan direction was chosen to minimize breathing artifacts. Participants were asked to take a deep breath and to hold their breath. No intravenous contrast material was injected. Exposure settings were 30 mAs at 120 kVp for patients weighting less than 80 kg (15/33), and 30 mAs at 140 kVp for those weighting more than 80 kg (18/33), without dose modulation. All participants received the same radiation dose for all follow-up scans. The in-plane voxel size ranged from  $0.53 \times 0.53$  mm to  $0.84 \times 0.84$  mm. Transverse images were reconstructed at a thickness of 1.0 mm and a 0.7 mm increment. All scans were reconstructed with a  $512 \times 512$  matrix and a moderately soft filter (Philips B).

### 5.2.3 Registration

Image registration is the process of spatially aligning two images, so that corresponding structures overlap. The CT scans were registered using the mutual information measure [30]. Details of the registration method are provided by Staring et al. [3] (see Chapter 4). The follow-up scan was deformed to match the image acquired at an earlier date. For




Figure 5.1: Example of a GGO that grew on an eight-month follow-up scan. The observers compared baseline (a) to follow-up (b). In one test set, no subtraction image was provided, in the other test set the subtraction image (c) was available. For the standard of reference, a registered follow-up data set (d) was also available in addition to data sets (a) - (c). Note that the nodule is not displayed exactly identical on (a) and (b), because of small differences in inspirational depth. Also note the white rim on (c), indicative of tumour growth.

improved registration performance, only the lungs were registered. The lungs were automatically segmented using an algorithm described by Sluimer et al. [59]. The resulting lung segmentation was used as a mask during the registration. To prevent the nodules from deforming in a nonrigid fashion, which would effectively conceal nodule growth, a local rigidity penalty term was employed [3] (see Chapter 4). This method allows the scans to be registered nonrigidly, while the nodule is not compressed or enlarged. The region that has to be kept rigid needs to be defined, but only roughly. Therefore, an ellipsoid is manually placed on the ground-glass opacity. The subtraction image is constructed by subtracting the baseline image from the deformed follow-up image. This way, growth is observed as a white rim, shrinkage as a dark rim. Density increase is observed in the subtraction image as a white color, density decrease as dark. The region outside the lungs was hidden on the subtraction image, to focus observer attention on the lungs.

#### 5.2.4 Observer Study

Pairs of scans of the same GGO were presented on a computer screen, with or without an additional subtraction image. Figure 5.1 shows an example of the data. To optimise the display, the 3D CT data were cropped to a region of  $80 \times 80 \times 80$  mm around the nodule, ensuring that the observers did not have to locate the nodule. A subtraction image was created after registering the image pairs using the method described above.

Observers were asked to independently rate volume and density change of the ground-glass opacity, and to rate the confidence in their decision. Volume changes were rated on a three point scale as either shrinkage, no change or growth. Density changes were also evaluated on a three point scale as either less dense, no change or more dense. Confidence was rated on a five-point scale, from highly unconfident to highly confident. A user interface was developed to perform the observer study (see Figure 5.2). As in a clinical situation, the observer could also make manual diameter measurements of the






Figure 5.2: A screen shot of the interface. The middle image is a follow-up of the data set on the left. In this case, the observer is also provided with the subtraction image after registration.

GGO, if warranted. Before the start of the study a short (10-15 minute) training session was performed on solid nodule data to familiarise the observer with the software.

Per patient all combinations of image pairs of the same nodule, not counting duplicates, were shown to the observers. For instance, if three time points were available, the pairs  $(t_0, t_1)$ ,  $(t_0, t_2)$  and  $(t_1, t_2)$  were evaluated. This resulted in one pair of scans for 18 nodules, three pairs of scans for 16 nodules, and 6 pairs of scans for 3 nodules, a total of 84 combinations. Each pair was shown once without the subtraction image and once with. The second scan of the pair was shown in its original format, i.e. not registered to the first scan, as in current clinical practice.

The 84 image pairs were shown to observers in random order, determined by a computer program. The observers were blinded to all patient information, it was unknown to the observers which pairs corresponded to the same nodule. All image pairs were scored twice by each observer, once with and once without the subtraction image. Each observer scored in two reading sessions, which were at least four weeks apart. Per session, all image pairs were scored, half of them with subtraction image and half without. In the second session, all image pairs that had been evaluated without subtraction image in the first session were now presented with subtraction image and vice versa. Per image pair the reading time was measured automatically, a fact the observer was informed of.

Four observers with a broad range of experience performed in the study: a radiology

resident with 3 years of experience in lung cancer screening in CT ( $> 2000$  scans), a pulmonologist with more than 15 years of experience with chest CT, a chest radiologist with more than 15 years of experience in CT, and a CT technician with special training in evaluating and reporting cancer screening CTs ( $> 3000$  scans in 2 years).

Lesions were included from a still ongoing screening trial. For most of these lesions, histopathological standard is not (yet) available. The standard of reference was therefore established by an independent chest radiologist with more than 15 years of experience, who did not participate in the observer study. He was provided with all available information: in addition to the original data and the subtraction images, he was given a follow-up scan registered to the baseline scan to better appreciate subtle changes. He had access to the anonymous opinions of all observers, to take into account the consensus of the observers. He knew the order in time of the scans and did not have to evaluate nodule pairs in random order. This way he could more easily compare multiple scans of the same nodule, and follow disease progress. This reference observer also had to indicate his confidence.

Despite availability of all this information, the reference observer was not able to confidently establish the presence of changes in all cases. To minimize impact of uncertainty on data analysis, only cases for which the standard of reference could be established with the highest confidence (rating 5) were included for further data analysis. Consequently, 59 out of 84 nodule pairs were eligible for evaluation of size change and 58 pairs were eligible for evaluation of density change. This final group consisted of 30 nodules (20 GGO and 10 part-solid).

## 5.2.5 Statistical Analysis

Agreement with the standard of reference and between observers was analysed with the weighted kappa statistic  $\kappa_w$  [87, 88]. Weights are taken inversely proportional to the difference in ratings, see [87–89]. To statistically compare  $\kappa_w$  with a value, the kappa z-test was employed. Note, that this test is not paired and that it is therefore conservative. A paired test, if it would be available, could indicate statistical significance sooner.

Besides measuring the overall effect of image subtraction, we also investigated the effect for difficult and least difficult cases separately, as indicated by observer confidence. Cases for which an observer was highly confident (marked 5) are subgroup 1, cases marked as 4 or lower are subgroup 2. For each observer, the agreement with the standard of reference was calculated for both subgroups separately.

Differences in confidence between the two methods were assessed by a two-sided paired Wilcoxon test.

A value of  $p < 0.05$  was considered to indicate a statistically significant difference. Statistical analyses were performed with Microsoft Office Excel 2003 software (Microsoft, Redmond, WA).

Table 5.1: Overall agreement with the standard of reference. Agreement  $\kappa_w$  with the standard of reference for 59 and 58 cases for size and density assessment, respectively. Between brackets the significance is given: \* ( $p < 0.05$ ) and \*\* ( $p < 0.001$ ) indicate a significant difference in agreement  $\kappa_w$  of the subtraction method compared to the standard method. ns indicates no significant difference ( $p = 0.05$ ).

		observer 1	observer 2	observer 3	observer 4	mean
size	standard	0.64	0.61	0.58	0.60	0.61
	subtraction	0.72 (ns)	0.70 (ns)	0.87 (**)	0.74 (*)	0.76
density	standard	0.57	0.52	0.52	0.51	0.53
	subtraction	0.68 (ns)	0.60 (ns)	0.74 (*)	0.54 (ns)	0.64

## 5.3 Results

According to the reference standard, 5/30 (17%) GGOs showed regression over time and no change was detected in 10/30 (33%) GGOs. Fifteen nodules (50%) showed an increase in size, while 8 of these GGOs (53%) additionally showed an increase in density. The observers needed on average 75s to score a case, both with and without subtraction.

The four observers made 14, 48, 4 and 16 (on average 21) diameter measurements, respectively, in all 84 image pairs when no subtraction image was provided. With a subtraction image, these numbers changed to 0, 0, 6 and 10 (on average 4), respectively.

The average observer confidence ratings for growth detection were 4.3 (out of 5) for the standard technique and 4.4 with the subtraction image available (significant increase for two observers:  $p = 0.79, 0.028, 0.79, 0.011$ ). For detection of density change the average confidence was 4.5 for both methods (no significant differences).

### 5.3.1 Agreement with the Standard of Reference

The observers' scores never differed more than one category from the standard of reference, with respect to changes of density and size of the GGOs; in other words it never occurred that observer and standard of reference stated the opposite (e.g., density increase and decrease, shrinkage and growth, respectively). The agreement of the individual observers with the standard of reference is reported in Table 5.1. Following the terminology of Landis and Koch [89] ( $0.4 < \kappa_w = 0.6$ : moderate agreement;  $0.6 < \kappa_w = 0.8$ : substantial agreement;  $\kappa_w > 0.8$ : almost perfect agreement), the mean agreement with the standard of reference was substantial for size change detection, and moderate for density change detection with the standard method. The agreement increased with the use of a subtraction image, on average from 0.61 to 0.76 for size change detection, and from 0.53 to 0.64 for density change detection. For size change detection the increase is significant for two raters ( $p = 0.33, 0.29, < 0.001, 0.029$ ); for density change detection the increase is significant for one observer ( $p = 0.20, 0.44, 0.009, 0.78$ ).

Table 5.2 lists the impact of image subtraction for the two subsets of nodules dependent on the readers' confidence. For subgroup 1 lesions, on average, the agreement only

Table 5.2: Agreement with the standard of reference, differentiation. Agreement  $\kappa_w$  with the standard of reference. Subgroup 1 are those cases with observer confidence 5. Subgroup 2 are the cases with confidence 1 - 4. Between brackets the significance is given: \* ( $p < 0.05$ ) and \*\* ( $p < 0.001$ ) indicate a significant difference in agreement  $\kappa_w$  of the subtraction method compared to the standard method. ns indicates no significant difference ( $p = 0.05$ ).

subgroup 1		observer 1	observer 2	observer 3	observer 4	mean
size	standard	0.92	0.87	0.79	0.88	0.86
	subtraction	0.85 (ns)	0.87 (ns)	0.93 (*)	1.00 (ns)	0.91
density	standard	0.72	0.72	0.75	0.83	0.76
	subtraction	0.68 (ns)	0.83 (ns)	0.81 (ns)	0.79 (ns)	0.78
subgroup 2		observer 1	observer 2	observer 3	observer 4	mean
size	standard	0.14	0.46	0.09	0.33	0.26
	subtraction	0.50 (*)	0.59 (ns)	0.70 (**)	0.51 (*)	0.57
density	standard	0.20	0.41	-0.11	0.28	0.19
	subtraction	0.61 (*)	0.47 (ns)	0.50 (*)	0.32 (ns)	0.47

marginally increased from 0.86 to 0.91 for size change and from 0.76 to 0.78 for density change detection. One of the observers achieved a significantly higher agreement with the reference standard with the subtraction image for the detection of size change ( $p = 0.004$ ). For all other readers, we did not find significant differences ( $p = 0.31, 0.97, 1.00$ , respectively, for size change and  $p = 0.73, 0.38, 0.51, 0.69$ , respectively, for density change).

For subgroup 2 lesions, agreement with the standard of reference improved substantially with the availability of the subtraction image. On average, the agreement increased from 0.26 to 0.57 and from 0.19 to 0.47 with the use of image subtraction. The improvements were significant for three of the four observers for the detection of size change ( $p = 0.010, 0.25, < 0.001, 0.046$ ), and for two of the four observers for the detection of density change ( $p = 0.002, 0.67, 0.001, 0.67$ ).

### 5.3.2 Inter-Observer Agreement

We calculated inter-observer agreement separately for both reading conditions, i.e. with and without availability of the subtraction image. Results are listed in Figure 5.3. With four readers, there were six combinations of reader pairs, for each of the two reading conditions. For the detection of size change, five of the six reader combinations showed an increase in agreement with the availability of subtraction image, differences were significant for three reader pairs ( $p = 0.027*, 0.099, < 0.001*, 0.43, 0.001*, 0.40$ ). For the detection of density change, also five of the six reader pairs showed an increase in agreement with the availability of the subtraction image, differences were significant




Figure 5.3: Inter-observer agreement for 59 and 58 cases for size and density assessment, respectively.  $O_iO_j$  indicates the agreement between observer  $i$  and  $j$  when using a subtraction image (WS, gray bar), and separately, without using the subtraction image (NS, white bar). At the top,  $p$ -values are given. Statistically significant improvement of  $\kappa_w$  is denoted by an asterisk.

for two reader pairs ( $p = 0.003^*, 0.060, 0.040^*, 0.14, 0.24, 0.41$ ). Both for size and density assessment the mean agreement between observers is higher with the subtraction method, compared to the standard visual comparison method.

## 5.4 Discussion

While for the evaluation of solid nodules, automatic volumetry represents an established and accepted tool for lesion characterization, these software tools are only under development for semi-solid and ground glass lesions. Current clinical practice is side-by-side comparison of follow-up scans, assisted by diameter measurements, each coming with its own constraints as is known from multiple studies with solid nodules.

This chapter presents the results of a study on the impact of a subtraction image on observer performance for the detection of change of GGOs on chest CT scans. The use of image subtraction was compared to the current clinical practice.

Since in most of these lesions no absolute standard of reference in terms of histology or verified volumetry was available, we used the judgment of an experienced radiologist

who had all clinical, reading and imaging data available, as a standard of reference. This standard has of course its limitations but represented - in our opinion - a sufficient superiority compared to the individual readings. Lesion pairs that could not be classified with high confidence even with all information available, were excluded from further data analysis. This was done to minimize the uncertainty on the data analysis, which is inevitably introduced by a non-perfect standard. Exclusion of nodules may lead to a bias because difficult nodules are omitted from evaluation. Separate analysis of these cases, however, showed that it was virtually impossible for observers to determine size change, indicated by very low weighted kappa values ( $\leq 0.19$ ), both for inter-observer variability and for agreement with the observer that established the standard of reference. This indicates that the establishment of a meaningful standard of reference for these difficult nodules may not be possible, if histology is not available.

We found that the use of a subtraction image improves the detection of changes. When observers had a high level of confidence even without image subtraction, the agreement with the standard of reference was high, both with and without subtraction. For a lower level of confidence, image subtraction improved observer performance substantially. Taking inter-observer agreement and agreement with the standard of reference as performance measures, we found that image subtraction gives a more reproducible and accurate method to detect change in size and density of a GGO.

Diameter measurements of lung tumor size on CT scans are often inconsistent [90–92]. In our study, the perceived need to perform diameter measurements varied strongly between observers: some used diameter measurements in 57% (48/84), others in only 5% (4/84) of images pairs. When the subtraction image was supplied, the number of diameter measurements dropped substantially, indicating that the subtraction image replaced the perceived need for diameter measurements. Reading of the additional subtraction image had no impact on reading time.

Automatic or semi-automatic volumetric or densitometric tools are also under development for GGOs. So far they suffer especially from insufficient automatic segmentation. Manual segmentation is an option but may suffer from reproducibility problems and is labor-intensive. Saito et al. [93] used density slicing (the selection of voxels with values in a range of intensities) to segment a GGO. Sumikawa et al. [94] employed a combination of adaptive thresholding and morphological pruning. A more sophisticated method based on voxel classification together with a vessel suppression technique was proposed by Zhou et al. [85]. Our technique provides an alternative that is based on visual image evaluation, and uses subtraction to enhance subtle differences.

Application of the subtraction method in clinical practice currently requires manual placement of an ellipsoid around the nodule. However, the only conditions are that the ellipsoid contains the nodule and excludes most of the surrounding parenchyma; thus, a rough delineation will suffice. After placement of the ellipsoid the scans need to be registered and subtracted, which requires about 10 minutes with the current implementation.

We conclude that image subtraction after registration for the follow-up of ground-glass opacities improves observer performance for the detection of nodule size and density change on chest CT, especially when the level of confidence of observers is low without the subtraction image, and decreases inter-observer variability.

# Registration of Cervical MRI Using Multifeature Mutual Information

# 6

## Abstract

---

Radiation therapy for cervical cancer can benefit from image registration in several ways, for example by studying the motion of organs, or by (partially) automating the delineation of the target volume and other structures of interest. In this chapter, the registration of cervical data is addressed using mutual information (MI) of not only image intensity, but also features that describe local image structure. Three aspects of the registration are addressed to make this approach feasible. Firstly, instead of relying on a histogram-based estimation of mutual information, which poses problems for a larger number of features, a graph-based implementation of  $\alpha$ -mutual information ( $\alpha$ -MI) is employed. Secondly, the analytical derivative of  $\alpha$ -MI is derived. This makes it possible to use a stochastic gradient descent method to solve the registration problem, which is substantially faster than non-derivative-based methods. Thirdly, the feature space is reduced by means of a principal component analysis, which also decreases the registration time. The proposed technique is compared to a standard approach, based on the mutual information of image intensity only. Experiments are performed on 93 T2-weighted MR clinical data sets acquired from 19 patients with cervical cancer. Several characteristics of the proposed algorithm are studied on a subset of 19 image pairs (one pair per patient). On the remaining data (36 image pairs, one or two pairs per patient) the median overlap is shown to improve significantly compared to standard MI from 0.75 to 0.81 for the bladder ( $p = 2 \cdot 10^{-5}$ ) and from 0.76 to 0.77 for the rectum ( $p = 9 \cdot 10^{-5}$ ).

---

Based upon: M. Staring, U.A. van der Heide, S. Klein, M.A. Viergever, and J.P.W. Pluim, Registration of Cervical MRI Using Multifeature Mutual Information, *submitted*.

## 6.1 Introduction

CERVICAL CANCER is a common type of cancer in women, with 493,000 incidences and 274,000 deaths worldwide in 2002 [95]. If surgery is not possible, external-beam radiation therapy is the treatment of choice [96], in combination with chemotherapy [97]. During therapy, a patient is irradiated several times in succession. In our hospital, a patient is irradiated 5 times a week for a duration of 5 weeks. Several factors are important for the success of radiotherapy. The first factor concerns the dose targeting. The region that should receive a high dose is the clinical target volume (CTV), which includes the gross tumour volume (GTV) and a region around the GTV with suspicion of micro-metastases. Neighbouring tissue such as the bladder and rectum should receive a dose as low as possible, to reduce complications during and after therapy. Therefore, accurate localisation of the cervix and surrounding structures is needed to effectively aim the dose delivery. A second factor is the ability to detect whether treatment plans need to be adapted to changes in the anatomy, for example due to response of the tumour to treatment. Adaptation of the treatment plans can improve sparing of bowel and rectum [98]. Pötter *et al.* [97] expect that replanning can have a minimising effect on treatment related morbidity. A third important factor in radiation therapy concerns the treatment margins used to accommodate uncertainties in the position of the target volumes [99]. The uncertainty can be high due to the large movement of the bladder, rectum and intestines. Quantification of the geometrical changes of these structures may help to reduce the treatment margin [6, 51, 100, 101].

Image registration can aid in all these cases. Segmentation of the relevant structures can be achieved by an atlas-based registration approach, as has been done for the prostate [9]. Another approach is to automatically update the manually created treatment plan of the first day, to that of the current day, instead of creating new manual segmentations. Automatic updates can be realised by performing a nonrigid intrasubject registration, and propagating the segmentation of the first day to the current day. Even if the quality of these automatic segmentations is not sufficient for direct clinical use, it can be used as an initialisation for annotators, which can save time and reduce observer variability [102]. To detect the need for adapting the radiotherapy treatment plans, one can register the current image with the baseline image. The volume of the GTV in combination with its motion are indicators for the necessity to update the treatment plans [103]. Geometrical changes can be quantified by the transformation resulting from image registration.

While traditionally CT images are used for radiotherapy, MRI is increasingly added, since it is superior to CT for staging, delineation of the relevant organs, and for measuring cervical carcinoma size and uterine extension, see [104–106]. Factors that hinder the success of image registration are MR imaging artifacts, the highly anisotropic voxels, and the large interfraction variability of the position, shape and size of the bladder, rectum and intestines in particular. In Figure 6.1 examples of these limiting factors are given.

In this chapter, we present an intensity-based nonrigid registration method to automatically align inpatient MR images of the cervix, for the purposes mentioned above. A standard approach, based on the mutual information of the intensity of the MR images only, may not be sufficient to overcome all of the limiting factors. Therefore, we propose



to use an algorithm that additionally incorporates features, in our case features that describe local image structure. The feature information is utilised by a multidimensional  $\alpha$ -mutual information measure ( $\alpha$ -MI), and is shown to improve the registration quality.

In order to make  $\alpha$ -MI feasible for a large number of features, we address several aspects of the method:

- Multidimensional  $\alpha$ -mutual information is implemented by computing the length of entropic graphs [107–110]. Commonly, mutual information is implemented by the estimation of a joint histogram. The joint histogram is of dimension  $2d$ , with  $d$  the number of features. Therefore, it suffers from the so-called curse of dimensionality: more and more samples are needed to fill the joint histogram up to a sufficient level, so that MI can be reliably estimated. The need to estimate the joint histogram can be eliminated, by directly relating entropy to the data, without estimating the probability density function. The choice for the entropic graphs implementation makes it possible to use multidimensional mutual information for a larger number of features.
- The analytical derivative of  $\alpha$ -MI with respect to the transformation parameters is derived. With the analytical derivative the registration problem can be solved using a stochastic gradient descent method. This is much faster, compared to a finite difference optimisation routine or a non-derivative-based optimiser.
- The feature space is reduced by employing a principal component analysis (PCA) on the total set of features. This also decreases the registration time substantially.

$\alpha$ -Mutual information for image registration was introduced by Hero *et al.* [107, 108]. It has been used on 2D data [109, 111], sometimes using a non-derivative-based optimiser [112]. In [113] a descent direction for a Euclidean minimal spanning tree is derived, and some results for rigid registration are given. Oubel *et al.* [114] apply  $\alpha$ -joint entropy to 3D tagged MR sequences of the heart, for which they also introduce an analytical derivative. The method was evaluated on five sequences, using only two features: the intensity from two different viewing directions.

In Section 6.2 details about the registration framework are given,  $\alpha$ -MI is explained, and the analytical derivative is deduced. The features used for the registration of the cervical data are discussed in Section 6.3. For tuning and evaluation, T2-weighted MR data sets are available for 19 patients. Registration parameters are tuned on a subset of 19 image pairs. Evaluation is performed on the remaining 36 image pairs by comparing automatic segmentations with a manually established ground truth, see Section 6.4 and 6.5.




Figure 6.1: Examples of limiting factors for image registration. Figures (a) and (b) show two scans of the same patient, intensity windowed for optimal display of the bladder. (b) demonstrates an intensity inhomogeneity over the bladder. (c) shows a slice in the  $x$  direction to illustrate data anisotropy. In (d) - (f) three scans of another patient are shown, exhibiting large variations in bladder filling (compare (d) with (e)) and rectum filling (compare (d,e) with (f)).

## 6.2 Method

### 6.2.1 Registration Framework

Registration of a  $q$ -dimensional moving image  $I_M : \Omega_M \subset \mathbb{R}^q \rightarrow \mathbb{R}$  to a fixed image  $I_F : \Omega_F \subset \mathbb{R}^q \rightarrow \mathbb{R}$  can be formulated as an optimisation problem:

$$\hat{\boldsymbol{\mu}} = \arg \min_{\boldsymbol{\mu}} \mathcal{S}(\mathbf{T}_{\boldsymbol{\mu}}; I_F, I_M, \Omega_F), \quad (6.1)$$

with  $\mathbf{T}_{\boldsymbol{\mu}} : \Omega_F \rightarrow \Omega_M$  the transformation modelled by the vector of parameters  $\boldsymbol{\mu} \in \mathbb{R}^M$ , and  $\mathcal{S}$  a suitable cost function (similarity measure) that is minimal when  $I_F(\mathbf{x})$  and  $I_M(\mathbf{T}_{\boldsymbol{\mu}}(\mathbf{x}))$  are aligned.

### 6.2.2 Mutual Information

The mutual information [21,115] of image intensities is defined as:  $\text{MI} = H(I_F) + H(I_M) - H(I_F, I_M)$ , with  $H(I) = -\sum p_i \log p_i$  the entropy of the intensities of image  $I$ ,  $H(I_F, I_M)$  the joint entropy, and  $p_i$  the intensity probabilities. It is usually implemented by constructing a joint histogram, which estimates the joint intensity probabilities. Details

about the specific implementation used in this chapter can be found in [30], in which the analytical derivative of MI can also be found.

### 6.2.3 Entropy, Entropic Graphs and $\alpha$ -Mutual Information

Mutual information is commonly calculated on the intensities only, i.e. on one-dimensional signals, by estimation of a joint histogram. For higher-dimensional signals, however, this computation method poses problems, as outlined in the Introduction. Therefore, we opt for an implementation of  $\alpha$ -mutual information introduced by Hero *et al.* [107, 108], based on entropic graphs.

Define  $\mathbf{z}(\mathbf{x}_i) = [z_1(\mathbf{x}_i), \dots, z_d(\mathbf{x}_i)]$  to be a vector of dimension  $d$  containing all feature values at point  $\mathbf{x}_i$ . For example, the first index  $z_1$  is the intensity of the image, and  $z_2$ ,  $z_3$  and  $z_4$  contain the values of the spatial derivatives of that image. It was shown by Beardwood *et al.* [116] (see also [117]) that the length of the graph that connects the feature points  $\{\mathbf{z}_i\}$ , is related to the Rényi  $\alpha$ -entropy [118]. It can be shown, using l'Hôpital's rule, that the limit  $\alpha \rightarrow 1$  of  $\alpha$ -entropy equals the Shannon entropy. Not all graphs are suited to compute  $\alpha$ -entropy, but for example the minimal spanning tree and the  $k$ -Nearest Neighbour ( $k$ NN) graph are [108]. We choose to use the  $k$ NN graph, because of its computational attractiveness [108].

Let  $\mathbf{z}^f(\mathbf{x}_i)$  be the feature vector of the fixed image at a point  $\mathbf{x}_i$ , and  $\mathbf{z}^m(\mathbf{T}_\mu(\mathbf{x}_i))$  that of the moving image at the transformed point  $\mathbf{T}_\mu(\mathbf{x}_i)$ . Let  $\mathcal{Z}_f = \{\mathbf{z}^f(\mathbf{x}_1), \dots, \mathbf{z}^f(\mathbf{x}_N)\}$  be the collection of  $N$  feature vectors drawn from the fixed image, and  $\mathcal{Z}_m$  that of the moving feature vectors. Let  $\mathbf{z}^{fm}(\mathbf{x}_i, \mathbf{T}_\mu(\mathbf{x}_i))$  be the concatenation of the two feature vectors:  $[\mathbf{z}^f(\mathbf{x}_i), \mathbf{z}^m(\mathbf{T}_\mu(\mathbf{x}_i))]$ , with corresponding collection  $\mathcal{Z}_{fm}$ . Three  $k$ NN graphs can be constructed on the three collections, where the total distance of a feature vector  $\mathbf{z}$  to its  $k$  nearest neighbours is given by:

$$\Gamma_i^f = \sum_{p=1}^k \|\mathbf{z}^f(\mathbf{x}_i) - \mathbf{z}^f(\mathbf{x}_{ip})\|, \quad (6.2)$$

$$\Gamma_i^m(\boldsymbol{\mu}) = \sum_{p=1}^k \|\mathbf{z}^m(\mathbf{T}_\mu(\mathbf{x}_i)) - \mathbf{z}^m(\mathbf{T}_\mu(\mathbf{x}_{ip}))\|, \quad (6.3)$$

$$\Gamma_i^{fm}(\boldsymbol{\mu}) = \sum_{p=1}^k \|\mathbf{z}^{fm}(\mathbf{x}_i, \mathbf{T}_\mu(\mathbf{x}_i)) - \mathbf{z}^{fm}(\mathbf{x}_{ip}, \mathbf{T}_\mu(\mathbf{x}_{ip}))\|, \quad (6.4)$$

with  $\mathbf{z}^f(\mathbf{x}_{ip})$ ,  $\mathbf{z}^m(\mathbf{T}_\mu(\mathbf{x}_{ip}))$  and  $\mathbf{z}^{fm}(\mathbf{x}_{ip}, \mathbf{T}_\mu(\mathbf{x}_{ip}))$  the  $p$ -th nearest neighbour of  $\mathbf{z}^f(\mathbf{x}_i)$ ,  $\mathbf{z}^m(\mathbf{T}_\mu(\mathbf{x}_i))$  and  $\mathbf{z}^{fm}(\mathbf{x}_i, \mathbf{T}_\mu(\mathbf{x}_i))$ , respectively. Note that the neighbours of the three graphs in general do not correspond. A graph-based estimator for  $\alpha$ -MI is defined as [108]:

$$\alpha\text{-}\widehat{\text{MI}}(\boldsymbol{\mu}; \mathcal{Z}_f, \mathcal{Z}_m, \mathcal{Z}_{fm}) = \frac{1}{\alpha - 1} \log \frac{1}{N^\alpha} \sum_{i=1}^N \left( \frac{\Gamma_i^{fm}(\boldsymbol{\mu})}{\sqrt{\Gamma_i^f \Gamma_i^m(\boldsymbol{\mu})}} \right)^{2\gamma}, \quad (6.5)$$

with  $\gamma = d(1 - \alpha)$ , and  $0 < \alpha < 1$  a user-defined constant.

## 6.2.4 Optimisation

To solve the minimisation problem (6.1), an iterative stochastic gradient descent optimisation strategy is employed [19]:

$$\boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t + a_t \tilde{\boldsymbol{g}}_t, \quad (6.6)$$

with  $\tilde{\boldsymbol{g}}_t$  an approximation of the exact derivative of the similarity measure evaluated at iteration  $t$ :  $\frac{\partial \mathcal{S}}{\partial \boldsymbol{\mu}}(\boldsymbol{\mu}_t)$ . The derivative is approximated by computing it over a subset of  $N$  points  $\boldsymbol{x} \in \Omega_F$  of the total number of points, randomly chosen in every iteration. This was shown to substantially accelerate the registration, while retaining convergence properties [19]. Samples that are not in the overlap of the fixed and deformed moving image domain are excluded. For the step size  $a_t$  a decaying function of the iteration number is used:  $a_t = a/(A+t)^\tau$ , with  $a > 0$ ,  $A \geq 1$  and  $0 < \tau \leq 1$  user-defined constants [19].

## 6.2.5 Derivative of $\alpha$ -Mutual Information

For the multidimensional case, registration is written as:

$$\hat{\boldsymbol{\mu}} = \arg \min_{\boldsymbol{\mu}} \alpha\text{-}\widehat{\text{MI}}(\boldsymbol{\mu}; \mathcal{Z}_f, \mathcal{Z}_m, \mathcal{Z}_{fm}). \quad (6.7)$$

In order to apply the stochastic gradient descent method given by Equation (6.6), the derivative  $\frac{\partial}{\partial \mu_j} \alpha\text{-}\widehat{\text{MI}}(\boldsymbol{\mu}; \mathcal{Z}_f, \mathcal{Z}_m, \mathcal{Z}_{fm})$  is required. Since computing the derivative using finite differences is extremely slow for a large number of transformation parameters  $M$ , we derive an analytical expression for it. The expression we derive is valid as long as graph topology does not change for small changes of  $\boldsymbol{\mu}$ .

For compact notation, define

$$G_i(\boldsymbol{\mu}) = \Gamma_i^{fm}(\boldsymbol{\mu}) / \sqrt{\Gamma_i^f \Gamma_i^m(\boldsymbol{\mu})} \quad \text{and} \quad (6.8)$$

$$\mathbf{d}_{ip}^{fm}(\boldsymbol{\mu}) = \mathbf{z}^{fm}(\mathbf{x}_i, \mathbf{T}_{\boldsymbol{\mu}}(\mathbf{x}_i)) - \mathbf{z}^{fm}(\mathbf{x}_{ip}, \mathbf{T}_{\boldsymbol{\mu}}(\mathbf{x}_{ip})), \quad (6.9)$$

then the derivative of  $\alpha\text{-}\widehat{\text{MI}}$  equals:

$$\frac{\partial}{\partial \mu_j} \alpha\text{-}\widehat{\text{MI}}(\boldsymbol{\mu}; \mathcal{Z}_f, \mathcal{Z}_m, \Omega_F) = \frac{-2d}{\sum_{i=1}^n G_i(\boldsymbol{\mu})^{2\gamma}} \sum_{i=1}^n G_i(\boldsymbol{\mu})^{2\gamma-1} \frac{\partial}{\partial \mu_j} G_i(\boldsymbol{\mu}). \quad (6.10)$$

The derivative of  $G_i(\boldsymbol{\mu})$  is written as:

$$\frac{\partial}{\partial \mu_j} G_i(\boldsymbol{\mu}) = \frac{\frac{\partial}{\partial \mu_j} \Gamma_i^{fm}(\boldsymbol{\mu}) - \frac{1}{2} \Gamma_i^{fm}(\boldsymbol{\mu}) \Gamma_i^m(\boldsymbol{\mu})^{-1} \frac{\partial}{\partial \mu_j} \Gamma_i^m(\boldsymbol{\mu})}{\sqrt{\Gamma_i^f \Gamma_i^m(\boldsymbol{\mu})}}. \quad (6.11)$$

From (6.11) we expand the derivative of  $\Gamma_i^{fm}(\boldsymbol{\mu})$ :

$$\frac{\partial}{\partial \mu_j} \Gamma_i^{fm}(\boldsymbol{\mu}) = \sum_{p=1}^k \frac{\partial}{\partial \mu_j} \sqrt{\mathbf{d}_{ip}^{fm}(\boldsymbol{\mu})^T \mathbf{d}_{ip}^{fm}(\boldsymbol{\mu})} \quad (6.12)$$

$$= \sum_{p=1}^k \frac{1}{2} \|\mathbf{d}_{ip}^{fm}(\boldsymbol{\mu})\|^{-1} \frac{\partial}{\partial \mu_j} \left( \mathbf{d}_{ip}^{fm}(\boldsymbol{\mu})^T \mathbf{d}_{ip}^{fm}(\boldsymbol{\mu}) \right) \quad (6.13)$$

$$= \sum_{p=1}^k \frac{\mathbf{d}_{ip}^{fm}(\boldsymbol{\mu})^T}{\|\mathbf{d}_{ip}^{fm}(\boldsymbol{\mu})\|} \cdot \frac{\partial}{\partial \mu_j} \mathbf{d}_{ip}^{fm}(\boldsymbol{\mu}), \quad (6.14)$$

where we used the Euclidean distance metric in (6.12). Only the derivative to the moving image features matter, since  $\frac{\partial}{\partial \mu_j} \mathbf{d}_{ip}^{fm}(\boldsymbol{\mu}) = [\mathbf{0}, \frac{\partial}{\partial \mu_j} \mathbf{d}_{ip}^m(\boldsymbol{\mu})]^T$ . So,

$$\begin{aligned} \frac{\partial}{\partial \mu_j} \Gamma_i^{fm}(\boldsymbol{\mu}) &= \sum_{p=1}^k \frac{\mathbf{d}_{ip}^m(\boldsymbol{\mu})^T}{\|\mathbf{d}_{ip}^{fm}(\boldsymbol{\mu})\|} \times \\ &\quad \left[ \frac{\partial \mathbf{z}^m}{\partial \mathbf{x}}(\mathbf{T}_\mu(\mathbf{x}_i)) \cdot \frac{\partial \mathbf{T}_\mu}{\partial \mu_j}(\mathbf{x}_i) - \frac{\partial \mathbf{z}^m}{\partial \mathbf{x}}(\mathbf{T}_\mu(\mathbf{x}_{ip})) \frac{\partial \mathbf{T}_\mu}{\partial \mu_j}(\mathbf{x}_{ip}) \right]. \end{aligned} \quad (6.15)$$

The derivative  $\frac{\partial}{\partial \mu_j} \Gamma_i^m(\boldsymbol{\mu})$  can be derived similarly. Summarising, the derivative of  $\alpha\widehat{\text{MI}}$  can be expressed in terms of the spatial derivative of the moving (feature) images  $\partial \mathbf{z}^m / \partial \mathbf{x}$ , the Jacobian of the transformation  $\partial \mathbf{T}_\mu / \partial \mu_j$ , the differences  $\mathbf{d}_{ip}^{fm}$  and  $\mathbf{d}_{ip}^m$ , and the graph distances  $\Gamma_i^f$ ,  $\Gamma_i^m(\boldsymbol{\mu})$  and  $\Gamma_i^{fm}(\boldsymbol{\mu})$ . Derivation of the analytical derivative is similar to that of Oubel *et al.* [114], who did the analysis for  $\alpha$ -joint entropy. Note that, in contrast with Oubel, we also take the derivative to the nearest neighbours into account in Equation (6.15).

### 6.3 Features

Now that the registration framework for multidimensional mutual information has been defined, features need to be chosen. From the vast amount of choices, features that describe the local structure of images supply supplementary knowledge, which may improve the registration. We choose the set of Cartesian image structure features up to the second order derivatives [119], listed in Table 6.1. Here,  $L$  denotes luminance or intensity,  $\mathbf{g} = (\partial L / \partial \mathbf{x})$  the spatial derivative,  $H$  is the Hessian of  $L$ , and  $tr(\cdot)$  denotes the matrix trace. Note that  $L_i L_i$  is the gradient magnitude, and  $L_{ii}$  the Laplacian. All features are invariant to rotation and translation. Additionally, this set is irreducible: other structure invariants of order up to two can be expressed in terms of features in this set, and the set is minimal. The invariants are computed using Gaussian derivatives at scale  $\sigma$ . Features at scale  $\sigma = 1$  and 2 were included. The additional use of  $\sigma = 4$  did not improve results, see Section 6.5.1. In total we have 15 features: 14 features that describe local structure<sup>1</sup>, and the original intensity data. In Figure 6.2(a)-(f) some examples are given.

<sup>1</sup>Note that also  $L$  can be computed at scale  $\sigma$ , by Gaussian filtering of the original intensity data.

Table 6.1: The irreducible set of second order Cartesian image structure invariants for 3D.

Einstein notation	matrix notation
$L$	$L$
$L_i L_i$	$\mathbf{g}^T \mathbf{g}$
$L_i L_{ij} L_j$	$\mathbf{g}^T H \mathbf{g}$
$L_i L_{ij} L_{jk} L_k$	$\mathbf{g}^T H H \mathbf{g}$
$L_{ii}$	$\text{tr}(H)$
$L_{ij} L_{ji}$	$\text{tr}(H H)$
$L_{ij} L_{jk} L_{ki}$	$\text{tr}(H H H)$

Note that all features are derived from the original image data, and can be computed before registration.

Since the computational complexity of (6.5) and (6.10) increases with the number of features  $d$ , and since not all features may have an equal contribution to the registration quality, it could be beneficial to select a subset of the total feature set for inclusion in the registration. We opt to perform a Principal Component Analysis (PCA) on the complete feature set, and select only the first  $P \leq d$  principal components for inclusion in the registration.

The features are normalised to have zero mean and unit variance. The PCA is performed on the concatenation of the fixed and moving images, resulting in the  $(N_F + N_M) \times d$  matrix  $[\mathbf{z}^f(\mathbf{x}_1), \dots, \mathbf{z}^f(\mathbf{x}_{N_F}), \mathbf{z}^m(\mathbf{y}_1), \dots, \mathbf{z}^m(\mathbf{y}_{N_M})]^T$ , with  $N_F = |\Omega_F|$  and  $N_M = |\Omega_M|$ . By performing the PCA on the concatenation, rather than on the separate images  $I_F$  and  $I_M$ , an identical linear combination is used for the feature images of the fixed and the moving set. The PCA algorithm returns the principal components (a linear combination of the input), together with the explained variance  $\lambda_i^2, i = 1, \dots, d$ , of each component. The percentage of total explained variance is given by  $\sum_{i=1}^P \lambda_i^2 / \sum_{i=1}^d \lambda_i^2$ . In Figure 6.2(g)-(l) the first 6 principal components, derived from the total feature set of the image in Figure 6.2(a), are shown.

## 6.4 Experiments

The graph-based  $\alpha$ -mutual information measure was implemented in the registration package `elastix`, developed by the authors. It is publicly available from <http://elastix.isi.uu.nl>. This package is largely based on the Insight Toolkit (ITK) [33]. The implementation of the  $k$ NN graph is based on the Approximate Nearest Neighbour (ANN) library, freely available from <http://www.cs.umd.edu/~mount/ANN/>.




Figure 6.2: Examples of the image structure features. (a) original image, (b) - (f) are the features that describe local structure ( $\sigma = 1$ ), with (b)  $L_i L_i$ , (c)  $L_i L_{ij} L_j$ , (d)  $L_{ii}$ , (e)  $L_{ij} L_{ji}$ , (f)  $L_{ij} L_{jk} L_{ki}$ . The first 6 principal components of the PCA are shown in (g)-(l). For this particular example the total explained variance for increasing  $P$  from 1 to 6 was 57%, 79%, 88%, 93%, 96%, and 98%. The images were windowed for optimal display.

### 6.4.1 Data

The MR data were acquired with a Philips 1.5T scanner (Gyrosan NT Intera; Philips Medical Systems, Best, The Netherlands), using a T2-weighted sequence taken in the transversal direction. Nineteen patients were scanned. Each patient was scanned 5 times, one scan each week, except for two patients who were scanned 4 times.

The image dimensions were  $512 \times 512 \times 30$  voxels of size  $0.625 \times 0.625 \times 4.5$  mm. The data were cropped before registration to a size of  $210 \times 250 \times 30$  voxels such that the data included the relevant structures. Manual segmentations of the GTV, CTV, bladder and rectum were available for each image. They were created by a radiation oncologist and approved by a radiologist. The segmentations were made with the clinical purpose in mind, and not for evaluation of registration quality specifically.

The data were divided into two sets. The first set contains the images at week 1 and 2 (19 image pairs), and is used for selection of the parameters. The second set consists of the remaining 55 images: 19 pairs of weeks 3 and 4, and 17 pairs of weeks 4 and 5; a total of 36 image pairs. The second set is used for comparing MI and  $\alpha$ -MI with parameters tuned on the first set. With this division evaluation of the method could be performed on all available patients.

### 6.4.2 Registration Settings

The registration parameters were chosen by trial-and-error on the first data set. Images at one week were registered to that of the next week. A rigid registration based on the mutual information of intensities only, was performed prior to nonrigid registration, to get a rough alignment. For the nonrigid registration, the transformation was parameterised with B-splines [4]. We employ a multiresolution scheme with three resolution levels on all feature images. Gaussian smoothing was applied, but no downsampling. Scales  $\sigma = 4.0, 2.0$  and  $1.0$  voxels were used in the  $x$  and  $y$  directions. For the  $z$  direction,  $\sigma = 2.0, 1.0$  and  $0.5$  voxel was used, because of the voxel anisotropy. A multigrid approach was used with a spacing of 80, 40 and 20 mm between the B-spline control points for the first, second, and final resolution, respectively. This yielded  $M = 3300$  parameters in the final resolution. For the optimisation procedure,  $A = 50$ ,  $\tau = 0.6$ , and  $a = 2000$  were set. During the parameter selection stage, 300 iterations were used. MI and  $\alpha$ -MI were compared on the second data set with 600 iterations. The number of samples, randomly selected in every iteration was set to  $N = 5000$ . To be comparable to MI,  $\alpha$  was set to 0.99. For the  $k$ NN graph implementation the ANN package was used. We selected  $k$ D trees, a standard splitting rule, and a bucket size of 50. The parameter  $\epsilon$  that defines the amount of error that is acceptable when computing the nearest neighbours was set to 10. The  $k = 20$  nearest neighbours were used to compute the derivative of  $\alpha$ - $\widehat{\text{MI}}$ .

### 6.4.3 Evaluation Measure

To evaluate the registration quality, manual segmentations of the CTV, bladder and rectum were used. Automatic segmentations were generated by transforming the manual segmentation of the moving image to the fixed image domain, using the transformation



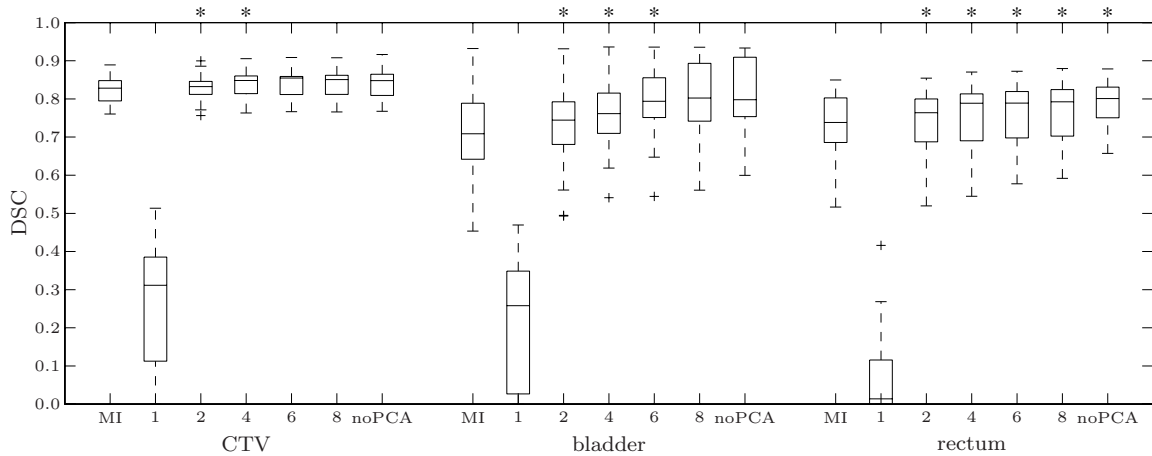


Figure 6.3: The effect of feature selection and the number of principal components  $P$ . For each anatomical structure, the leftmost column shows the result for MI, the right-most for  $\alpha$ -MI without PCA ( $d = 15$ ), and the columns in between for  $\alpha$ -MI with PCA. The number below the graph refers to  $P$ . A star indicates a statistical significant difference of the median overlap compared to the previous column, starting from  $P = 2$ .

$T_{\hat{\mu}}$  found by the registration. This transformed segmentation  $A$  was compared to the manual segmentation  $B$  of the fixed image, using the Dice similarity coefficient (DSC) as a measure of overlap [40]:

$$\text{DSC}(A, B) = \frac{2|A \cap B|}{|A| + |B|}, \quad (6.16)$$

where  $|\cdot|$  denotes the number of voxels within the segmentation. A value of 0 indicates that  $A$  and  $B$  are disjoint, 1 indicates perfect agreement. The DSCs are presented by box-and-whisker plots.

Paired, two-sided Wilcoxon tests were performed on the overlap values of two experiments. A value of  $p < 0.05$  was considered to indicate a statistically significant difference.

## 6.5 Results

### 6.5.1 Results: Data Set 1 (Parameter Selection)

Several aspects of the proposed algorithm were investigated. All experiments described in this section were performed on the first data set.

The selection of the scales of the features was addressed in a first experiment. Compared to using only  $\sigma = 1$ , including features at scale  $\sigma = 1$  and 2 gave slightly better results. The effect was most noticeable for the bladder, where the median overlap increased from 0.76 to 0.79. The additional use of  $\sigma = 4$  did not improve the results. Therefore, for  $\alpha$ -MI, scales  $\sigma = 1$  and 2 were chosen.

In a second experiment, the effect of the number of principal components  $P$  that is used during the registration, was investigated. The results are shown in Figure 6.3. If

Table 6.2: Runtime for the various registration experiments. The experiments were performed on a standard pc (AMD Opteron 250 running at 2.4 GHz.), using 300 iterations.

method	time (min.)
MI	1.2
$\alpha$ -MI $P = 1$	32
$P = 2$	42
$P = 4$	57
$P = 6$	73
$P = 8$	90
$d = 15$	162

only one principal component was taken into account, most of the registrations failed; the result was worse than MI. Results similar to  $P = 1$  were obtained when using the original intensity images as the single feature in  $\alpha$ -MI. A possible explanation for this behaviour can be found by closer inspection of the definition of  $\alpha$ -MI in (6.5). The expression contains the division by  $\Gamma_i^f$ . In the case  $d = 1$  this can be a very small number when  $k$  out of the  $N$  samples have similar intensity. Division by a small number can make the estimation of  $\alpha$ -MI unstable. For  $d > 1$  the probability that samples have a small distance to their nearest neighbours is in general reduced, except in the unlikely event that the second feature is constant. This observation coincides with our experience that the problem is reduced for larger  $k$ . The median total explained variance resulting from the PCA analysis was 0.57, 0.80, 0.93, 0.98, and 0.99 for  $P = 1, 2, 4, 6,$  and  $8$ , respectively. Increasing  $P$  improved the results up to approximately  $P = 6$  or  $P = 8$ .

In a third experiment the influence of the use of PCA was examined.  $\alpha$ -MI was performed without the feature space reduction, i.e. with all local image structure features at scale 1 and 2 and the original intensity data (in total 15 features). For each anatomical structure, the right-most column of Figure 6.3 shows the result for this experiment. Inclusion of all features did not improve registration performance considerably, except for the rectum where the difference with PCA ( $P = 8$ ) is significant. The best trade-off of computation time and registration accuracy was with the use of PCA with  $P = 6$ . With this setting the registration time was 73 minutes, a reduction with a factor of 2.2 compared to using the full feature set, see Table 6.2.

### 6.5.2 Results: Data Set 2 (Evaluation)

MI is compared to  $\alpha$ -MI (PCA,  $P = 6$ ) on data set 2. For each of the two similarity measures, 36 registrations are performed. The results are shown in Figure 6.4. Not much difference between the two methods is seen at the CTV. At the bladder and the rectum, however, the median overlap increases significantly from 0.75 to 0.81 ( $p = 2 \cdot 10^{-5}$ ) and from 0.76 to 0.77 ( $p = 9 \cdot 10^{-5}$ ), respectively. Also note the increase of the first quartile, meaning that if the technique were to be used in the clinic, less manual correction of the treatment plans is needed when using  $\alpha$ -MI.

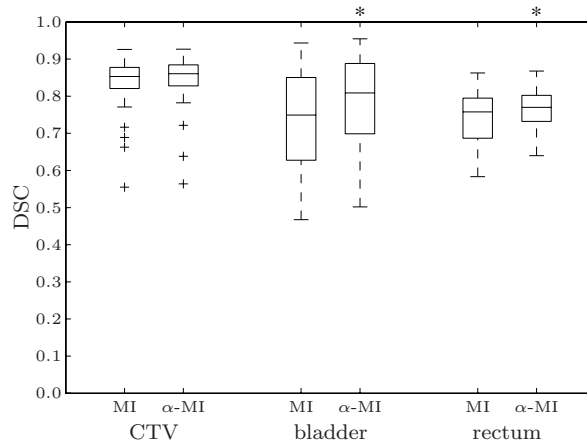


Figure 6.4: MI versus  $\alpha$ -MI. For each anatomical structure, the leftmost column shows the result for MI. The second column shows the result for  $\alpha$ -MI ( $P = 6$ ). A star indicates a statistical significant difference of the median overlaps of the two methods.

In Figure 6.5 a typical example of the result of registration is given, for MI (a) and  $\alpha$ -MI (b). Much less deformation is expected at fatty tissue and near bony anatomy, see for example the bottom and right side of Figure 6.5 (a) and (b). The examples show accurate registrations at those positions, both for MI and  $\alpha$ -MI. A large difference in bladder filling is hard to recover for MI. Although the bladders are not perfectly aligned by  $\alpha$ -MI, the result is better. Registration problems sometimes also occur at the GTV, where the tissue changes due to irradiation. Figures 6.5(c) and (d) show an example of such a tissue change, which is difficult to handle for the registration.

## 6.6 Discussion and Conclusion

In this chapter we have introduced a registration method for cervical MR images capable of taking into account multiple image features. Three aspects were addressed to be able to use the method in practice, also for nonrigid registration. Firstly, a graph-based implementation of  $\alpha$ -mutual information was chosen, instead of a histogram-based implementation. Histogram-based approaches are currently not able to cope with a larger number of features, although efforts to do so have been undertaken [120–122]. The graph-based implementation makes it possible to take into account an enlarged feature space. Secondly, an analytical derivative was constructed, which enables the use of a fast stochastic gradient descent optimisation routine. A finite difference gradient descent approach would require  $2M$  evaluations of the cost function per iteration, which results in a runtime of 153 h., compared to 162 min. when the analytical derivative is employed. Finally, the feature space was reduced by means of a PCA algorithm. This reduces the number of features from 15 to 6, which leads to another reduction in registration time by a factor of  $\approx 2$ . This work is the first to report feasible registration times for multi-dimensional mutual information with the number of parameters of the transformation in the thousands ( $M = 3300$ ).




Figure 6.5: Example results. The bladder filling of the fixed and moving image is very different; a large deformation is required at that position. The fixed image is combined with the deformed moving image, using a checkerboard pattern. (a) The result of MI clearly is not well aligned at the bladder. (b)  $\alpha$ -MI performs much better for this large deformation. (c) and (d) show the GTV (delineated) at two subsequent weeks. Note the change of the tissue around the GTV.

It was shown on clinical data of patients with cervical cancer that the proposed method outperforms a standard approach based on the mutual information of intensity only. For use in the clinic, further improvements are still required. One option is the use of a localised version of mutual information [9]. On MR images acquired for the radiation therapy of the prostate this was shown to improve the registration, at no additional computational cost. Another possibility with  $\alpha$ -MI is the inclusion of other imaging data, acquired at the same time point, such as an MR image acquired in another scanning direction, or acquired with a different protocol. For the purpose of updating treatment plans, it is possible to use a multi-atlas matching approach [8], once two or more previous scans are available. Rohlfing *et al.* [8] reported improved results for this technique, compared to using a single atlas image.

Propagation of the GTV segmentation is probably not possible with a registration approach only. Even if anatomical correspondence is found, tumourous tissue may disappear in time, due to successful treatment. In this case it might be necessary to employ

a tissue classification technique, after the registration.

Although registration using multiple image features is feasible with the proposed method, in our implementation it still requires much more time than the histogram-based approach. For clinical use the computational burden of  $\alpha$ -MI probably needs to be decreased, depending on the specific application. From Equation (6.5) it is clear that the major part of the computation can be parallelised by distributing the summation over multiple processors. Another interesting approach is to use a Minimal Spanning Tree (MST) instead of a  $k$ NN graph. With the MST,  $\alpha$ -MI is computed as a sum over  $N - 1$  edges; with the  $k$ NN graph it is a sum over  $kN$  edges, see (6.5). During the experiments we noticed that most of the computation time is spent in calculating (6.15). The construction of the graph and searching for the  $k$  nearest neighbours appeared to be of minor influence. Therefore, although it takes more time to construct the MST, it is probably faster to compute the analytical derivative with it. The influence of this change on the rate of convergence of the registration remains to be investigated.

In conclusion, compared to a standard approach, the proposed method accomplishes improved nonrigid registration for a challenging registration problem, by inclusion of multiple features in the registration cost function.

## Acknowledgements

We gratefully acknowledge L. van de Bunt, M.D. and G.A.P. de Kort, M.D. for providing the manual segmentations.



# Summary 7

## 7.1 Summary

IMAGE MATCHING is important for the comparison of medical images. Comparison is of clinical relevance for the analysis of differences due to changes in the health of a patient. For example, when a disease is imaged at two time points, then one wants to know if it is stable, has regressed, or progressed. Another example is when a patient is imaged before and after treatment. In this case one wants to know to what extent treatment was successful. The type of clinical action that is taken is dependent on the outcome of this analysis.

This thesis addresses the analysis of change in medical images by means of image registration. We distinguish three contributions of the work described in this thesis, given in the next three sections.

### 7.1.1 `elastix`

Together with Stefan Klein software was developed that aids the research on image registration. The software package is called `elastix` and is described in **Chapter 2**. An image registration algorithm consists of several parts, such as the cost function, the optimisation routine, and the transformation. For each part many choices are available, but it is beforehand not known which selection is optimal for the problem at hand. `elastix` contains a collection of the algorithms available in the literature, and allows for easy selection, testing and comparison of the different methods. The command-line interface enables automated processing of large numbers of data sets, by means of scripting. `elastix` is based on the Insight Toolkit and is open source (distributed under the BSD license). The software was used for all experiments described in this thesis, and is increasingly utilised, both within and outside the Image Sciences Institute.

### 7.1.2 Tissue-dependent Image Registration and Change Detection

Popular nonrigid image registration algorithms do not take the rigidity of different tissue types into account. This can lead to undesired effects in three situations:

- a. The image contains various tissue types, each with its own mechanical stiffness.

- b. Structural changes over time (e.g. tumour growth) need to be visualised.
- c. The visibility of structures varies between acquisitions due to the use of contrast material.

In all three situations the objects of interest should be considered as locally rigid by the registration algorithm, to allow change analysis by image subtraction.

In **Chapter 3** this problem is tackled by locally adapting the deformation field at structures that must be kept rigid, using a tissue-dependent filtering technique. This adaptive filtering of the deformation field results in locally linear transformations without scaling or shearing. The degree of filtering is related to tissue stiffness: more filtering is applied at stiff tissue locations, less at parts of the image containing nonrigid tissue.

Evaluation of the proposed tissue-dependent filtering was performed on 3D CT data of the thorax and on 2D Digital Subtraction Angiography (DSA) images. It was shown that the adaptive filter leads to improved registration results: tumour volumes and vessel widths are preserved rather than affected. On average 99% of the tumour volume is preserved, compared to 95% with a standard algorithm that does not take into account rigidity of the tumour. For vessels 98% of the diameter is retained, compared to 85% of a standard algorithm. The filter has as a downside that when it is applied strongly it does not allow for rotations of rigid objects.

A rigidity penalty term does not have this downside. The construction of the term is described in **Chapter 4**. It is derived from the definition of a rigid transformation, and results in three distinct terms, related to affine-ness of the transformation, and orthonormality and properness of the rotation matrix. The three conditions are combined in one term. Locality is introduced by a weight function, dependent on the spatial position. Local rigidity is included in the registration by including a penalty term in the cost function. A fast algorithm is obtained by exploiting the B-spline parameterisation of the transformation.

Evaluation was again performed on the 3D CT data of the thorax and the 2D angio data. The tumour volume decreased on average to 78% of the original volume for a standard algorithm; for the algorithm that included the rigidity penalty term 98% of the volume was retained. For the vessels the diameter preservation improved from 84% to 99%. From these results we concluded that the rigidity penalty term is capable of keeping user-defined structures locally rigid, while performing nonrigid registration.

Situation b. from the list above is analysed further in a clinical setting in **Chapter 5**. To improve patient prognosis it is important to determine at an early stage whether a lung nodule is benign or malignant. Changes in size or density are two important criteria for malignancy. Detection of growth or density change, however, is frequently difficult for a specific type of nodules, ground-glass opacities, because of their appearance on CT. In current clinical practice the nodules are compared visually, with the aid of manual diameter measurements. We studied the impact of the additional use of a subtraction image on the detection of changes for these nodules. The subtraction image was generated by the registration algorithm described in Chapter 4.



To this end we let four observers rate size and density change for approximately 60 pairs of images containing a ground-glass opacity, once with and once without the subtraction image. Especially for nodules that had changed only in a subtle fashion, the subtraction image improved observer performance substantially compared to the current clinical practice. The agreement with the standard of reference, measured with a kappa statistic, increased from  $\kappa_w = 0.26$  to 0.57 for size change detection, and from 0.19 to 0.47 for density change detection, for these subtle cases. Also the inter-observer agreement improved, indicating that the subtraction image supplies a more reproducible method.

### 7.1.3 Improved Registration Performance

In **Chapter 6**, the registration of cervical data was addressed using mutual information of not only image intensity, but also features that describe local image structure. Several technical aspects of the proposed cost function were addressed to obtain a feasible registration time.

It was shown on challenging data (challenging due to image quality and large deformations) from the Department of Radiotherapy (T2-weighted MR data of the cervix) that, compared to a standard algorithm based on intensity only, the proposed algorithm improves performance significantly. For structures that have very different size and shape between the two images, the performance increase was most noticeable: the median overlap of the bladder after registration increased significantly from 0.75 to 0.81.

## 7.2 Clinical Relevance

The work described in this thesis enables nonrigid registration while keeping user-defined structures locally rigid. As shown in Chapter 5, this can be used to visualise changes of ground-glass opacities. The method enables early detection of change, which improves the patient prognosis and may help to save lives.

Research is needed to further automate the algorithm, since currently a rough manual segmentation of the region that is to be kept rigid (the nodule) is required. Since the segmentation only needs to be rough, it is probably possible to automate this part, using standard segmentation algorithms. Alternatively, one may focus research attention on completely automatic and accurate segmentation of these ground-glass nodules, and perform volumetry to detect changes. It would be interesting to see how that method would compare to relying on the human eye by means of a subtraction image. For solid nodules, tools for (semi)automatic volumetry already exist, since these nodules are easier to segment. For these nodules, a comparison of both methods could be started immediately.

The radiotherapy department of the University Medical Center Utrecht is developing an integrated MRI-accelerator system. This system enables imaging of a target region (e.g. the cervix), while at the same time irradiating it. A potential advantage of such a system compared to the current methodology is that more accurate dose targeting is possible. This enables the use of a higher dose for the target region, and the sparing

of tissue near a tumour, thereby reducing patient complications. However, for the dose targeting the relevant organs need to be segmented. Manual segmentation is not an option, because of time constraints. Automation can be done by means of nonrigid image registration, but currently these techniques lack accuracy. In Chapter 6 a step is taken towards clinically acceptable registration accuracy.

In conclusion, although nonrigid image registration is not yet widely used in the clinic, it can be of added value for the analysis of changes in medical images.

## References

- [1] “PET-CT,” June 2008. [Online]. Available: [www.medical.philips.com/main/products/pet/clinicalimages/geminigxl/10/](http://www.medical.philips.com/main/products/pet/clinicalimages/geminigxl/10/)
- [2] D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank, “PET-CT image registration in the chest using free-form deformations,” *IEEE Transactions on Medical Imaging*, vol. 22, no. 1, pp. 120 – 128, 2003.
- [3] M. Staring, S. Klein, and J. P. W. Pluim, “A rigidity penalty term for nonrigid registration,” *Medical Physics*, vol. 34, no. 11, pp. 4098 – 4108, 2007.
- [4] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, “Nonrigid registration using free-form deformations: Application to breast MR images,” *IEEE Transactions on Medical Imaging*, vol. 18, no. 8, pp. 712 – 721, 1999.
- [5] T. Rohlfing, C. R. Maurer Jr., D. A. Bluemke, and M. A. Jacobs, “Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint,” *IEEE Transactions on Medical Imaging*, vol. 22, no. 6, pp. 730 – 741, 2003.
- [6] M. Foskey, B. Davis, L. Goyal, S. Chang, E. Chaney, N. Strehl, S. Tomei, J. Rosenman, and S. Joshi, “Large deformation three-dimensional image registration in image-guided radiation therapy,” *Physics in Medicine and Biology*, vol. 50, pp. 5869–5892, 2005.
- [7] X. Pennec, P. Cachier, and N. Ayache, “Tracking brain deformations in time sequences of 3D US images,” *Pattern Recognition Letters*, vol. 24, no. 4-5, pp. 801–813, 2003.
- [8] T. Rohlfing, R. Brandt, R. Menzel, and C. R. Maurer Jr., “Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains,” *NeuroImage*, vol. 21, no. 4, pp. 1428–1442, 2004.
- [9] S. Klein, U. A. van der Heide, I. M. Lips, M. van Vulpen, M. Staring, and J. P. W. Pluim, “Automatic segmentation of the prostate in 3D MR images by atlas matching using localized mutual information,” *Medical Physics*, vol. 35, no. 4, pp. 1407 – 1417, April 2008.
- [10] T. A. Sundaram and J. C. Gee, “Towards a model of lung biomechanics: pulmonary kinematics via registration of serial lung images,” *Medical Image Analysis*, vol. 9, no. 6, pp. 524–537, 2005.

- [11] H. A. Vrooman, C. A. Cocosco, F. van der Lijn, R. Stokking, M. A. Ikram, M. W. Vernooij, M. M. B. Breteler, and W. J. Niessen, “Multi-spectral brain tissue segmentation using automatically trained  $k$ -nearest-neighbor classification,” *NeuroImage*, vol. 37, no. 1, pp. 71–81, 2007.
- [12] P. Anbeek, K. L. Vincken, M. J. P. van Osch, R. H. C. Bisschops, and J. van der Grond, “Automatic segmentation of different-sized white matter lesions by voxel probability estimation,” *Medical Image Analysis*, vol. 8, no. 3, pp. 205–215, 2004.
- [13] L. G. Brown, “A survey of image registration techniques,” *ACM Computing Surveys*, vol. 24, no. 4, pp. 325–376, 1992.
- [14] J. B. A. Maintz and M. A. Viergever, “A survey of medical image registration,” *Medical Image Analysis*, vol. 2, no. 1, pp. 1 – 36, 1998.
- [15] H. Lester and S. R. Arridge, “A survey of hierarchical non-linear medical image registration,” *Pattern Recognition*, vol. 32, no. 1, pp. 129 – 149, 1999.
- [16] D. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, “Medical image registration,” *Physics in Medicine and Biology*, vol. 46, no. 3, pp. R1 – R45, 2001.
- [17] J. Modersitzki, *Numerical Methods for Image Registration*. Oxford University Press, 2004.
- [18] F. Maes, D. Vandermeulen, and P. Suetens, “Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information,” *Medical Image Analysis*, vol. 3, no. 4, pp. 373–386, 1999.
- [19] S. Klein, M. Staring, and J. P. W. Pluim, “Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines,” *IEEE Transactions on Image Processing*, vol. 16, no. 12, pp. 2879 – 2890, December 2007.
- [20] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, “Interpolation artefacts in mutual information-based image registration,” *Computer Vision and Image Understanding*, vol. 77, no. 2, pp. 211–232, 2000.
- [21] P. Viola and W. M. Wells III, “Alignment by maximization of mutual information,” *International Journal of Computer Vision*, vol. 24, no. 2, pp. 137 – 154, 1997.
- [22] J. West, J. Fitzpatrick, M. Wang, B. Dawant, C. Maurer Jr., R. Kessler, R. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D. Vandermeulen, P. van den Elsen, S. Napel, T. Sumanaweera, B. Harkness, P. Hemler, D. Hill, D. Hawkes, C. Studholme, J. Maintz, M. Viergever, G. Malandain, X. Pennec, M. Noz, G. Maguire Jr., M. Pollack, C. Pelizzari, R. Robb, D. Hanson, and R. Woods, “Comparison and evaluation of retrospective intermodality brain image registration techniques,” *Journal of Computer Assisted Tomography*, vol. 21, no. 4, pp. 554 – 566, 1997.

- [23] R. Bajcsy and S. Kovačič, “Multiresolution elastic matching,” *Computer Vision, Graphics and Image Processing*, vol. 46, no. 1, pp. 1 – 21, 1989.
- [24] B. Fischer and J. Modersitzki, “A unified approach to fast image registration and a new curvature based registration technique,” *Linear Algebra Applications*, vol. 380, pp. 107 – 124, 2004.
- [25] P. Thévenaz, U. E. Ruttimann, and M. Unser, “A pyramid approach to subpixel registration based on intensity,” *IEEE Transactions on Image Processing*, vol. 7, no. 1, pp. 27 – 41, 1998.
- [26] J. Kybic and M. Unser, “Fast parametric elastic image registration,” *IEEE Transactions on Image Processing*, vol. 12, no. 11, pp. 1427 – 1442, November 2003.
- [27] C. Studholme, D. L. G. Hill, and D. J. Hawkes, “Automated 3-D registration of MR and CT images of the head,” *Medical Image Analysis*, vol. 1, no. 2, pp. 163–175, 1996.
- [28] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, “A comparison of similarity measures for use in 2D-3D medical image registration,” *IEEE Transactions on Medical Imaging*, vol. 17, no. 4, pp. 586–595, 1998.
- [29] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, “Multimodality image registration by maximization of mutual information,” *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 187 – 198, 1997.
- [30] P. Thévenaz and M. Unser, “Optimization of mutual information for multiresolution image registration,” *IEEE Transactions on Image Processing*, vol. 9, no. 12, pp. 2083 – 2099, 2000.
- [31] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever, “Mutual-information-based registration of medical images: a survey,” *IEEE Transactions on Medical Imaging*, vol. 22, no. 8, pp. 986 – 1004, 2003.
- [32] C. Studholme, D. L. G. Hill, and D. J. Hawkes, “An overlap invariant entropy measure of 3D medical image alignment,” *Pattern Recognition*, vol. 32, no. 1, pp. 71–86, 1999.
- [33] L. Ibáñez, W. Schroeder, L. Ng, and J. Cates, *The ITK Software Guide*, 2nd ed., Kitware, Inc. ISBN 1-930934-15-7, 2005.
- [34] G. K. Rohde, A. Aldroubi, and B. M. Dawant, “The adaptive bases algorithm for intensity-based nonrigid image registration,” *IEEE Transactions on Medical Imaging*, vol. 22, no. 11, pp. 1470–1479, 2003.
- [35] V. Arsigny, X. Pennec, and N. Ayache, “Polyrigid and polyaffine transformations: A novel geometrical tool to deal with non-rigid deformations application to the registration of histological slices,” *Medical Image Analysis*, vol. 9, no. 6, pp. 507 – 523, 2005.

- [36] M. Unser, “Splines: A perfect fit for signal and image processing,” *IEEE Signal Processing Magazine*, vol. 16, no. 6, pp. 22 – 38, 1999.
- [37] B. Likar and F. Pernuš, “A hierarchical approach to elastic registration based on mutual information,” *Image and Vision Computing*, vol. 19, no. 1-2, pp. 33–44, 2001.
- [38] P. Thévenaz, M. Bierlaire, and M. Unser, “Halton sampling for image registration based on mutual information,” *Sampling Theory in Signal and Image Processing*, vol. (in press), 2008.
- [39] P. Thévenaz, T. Blu, and M. Unser, “Interpolation revisited,” *IEEE Transactions on Medical Imaging*, vol. 19, no. 7, pp. 739 – 758, July 2000.
- [40] L. R. Dice, “Measures of the amount of ecologic association between species,” *Ecology*, vol. 26, no. 3, pp. 297 – 302, 1945.
- [41] C. A. van Iersel, H. J. de Koning, G. Draisma, W. P. Mali, E. T. Scholten, K. Nackaerts, M. Prokop, J. D. F. Habbema, M. Oudkerk, and R. J. van Klaveren, “Risk-based selection from the general population in a screening trial: Selection criteria, recruitment and power for the Dutch-Belgian randomised lung cancer multi-slice CT screening trial (NELSON),” *International Journal of Cancer*, vol. 120, no. 4, pp. 868 – 874, 2007.
- [42] K. Murphy, B. van Ginneken, S. Klein, M. Staring, and J. P. W. Pluim, “Semi-automatic reference standard construction for quantitative evaluation of lung CT registration,” in *MICCAI*, 2008, submitted.
- [43] M. Staring, S. Klein, and J. P. W. Pluim, “Nonrigid registration with tissue-dependent filtering of the deformation field,” *Physics in Medicine and Biology*, vol. 52, no. 23, pp. 6879 – 6892, 2007.
- [44] E. M. van Rikxoort, Y. Arzhaeva, and B. van Ginneken, “A multi-atlas approach to automatic segmentation of the caudate nucleus in MR brain images,” in *3D Segmentation In The Clinic: A Grand Challenge*, T. Heimann, M. Styner, and B. van Ginneken, Eds., 2007, pp. 29–36, <http://mbi.dkfz-heidelberg.de/grand-challenge2007>.
- [45] M. Staring, S. Klein, and J. P. W. Pluim, “Nonrigid registration with adaptive, content-based filtering of the deformation field,” in *SPIE Medical Imaging: Image Processing*, ser. Proceedings of SPIE, vol. 5747. SPIE Press, 2005, pp. 212 – 221.
- [46] C. Tanner, J. A. Schnabel, D. Chung, M. J. Clarkson, D. Rueckert, D. L. G. Hill, and D. J. Hawkes, “Volume and shape preservation of enhancing lesions when applying nonrigid registration to a time series of contrast enhancing MR breast images,” in *MICCAI*, ser. Lecture Notes In Computer Science, vol. 1935, 2000, pp. 327 – 337.

- [47] T. Rohlfing and C. R. Maurer Jr., “Intensity-based nonrigid registration using adaptive multilevel free-form deformation with an incompressibility constraint,” in *MICCAI*, ser. Lecture Notes in Computer Science, vol. 2208, 2001, pp. 111 – 119.
- [48] D. Loeckx, F. Maes, D. Vandermeulen, and P. Suetens, “Nonrigid image registration using free-form deformations with a local rigidity constraint,” in *MICCAI*, ser. Lecture Notes in Computer Science, vol. 3216, 2004, pp. 639 – 646.
- [49] J. A. Little, D. L. G. Hill, and D. J. Hawkes, “Deformations incorporating rigid structures,” *Computer Vision and Image Understanding*, vol. 66, no. 2, pp. 223 – 232, 1997.
- [50] P. J. Edwards, D. L. G. Hill, J. A. Little, and D. J. Hawkes, “A three-component deformation model for image-guided surgery,” *Medical Image Analysis*, vol. 2, no. 4, pp. 355 – 367, 1998.
- [51] K. K. Brock, M. B. Sharpe, L. A. Dawson, S. M. Kim, and D. A. Jaffray, “Accuracy of finite element model-based multi-organ deformable image registration,” *Medical Physics*, vol. 32, no. 6, pp. 1647 – 1659, 2005.
- [52] H. Lester, S. R. Arridge, K. M. Jansons, L. Lemieux, J. V. Hajnal, and A. Oatridge, “Non-linear registration with the variable viscosity fluid algorithm,” in *IPMI*, ser. Lecture Notes In Computer Science, vol. 1613. Springer-Verlag, 1999, pp. 238 – 251.
- [53] J. P. Thirion, “Image matching as a diffusion process: an analogy with Maxwell’s demons,” *Medical Image Analysis*, vol. 2, no. 3, pp. 243 – 260, 1998.
- [54] P. Cachier, X. Pennec, and N. Ayache, “Fast non-rigid matching by gradient descent: Study and improvements of the demons algorithm,” INRIA, Research Report 3706, 1999.
- [55] R. Stefanescu, X. Pennec, and N. Ayache, “Grid powered nonlinear image registration with locally adaptive regularization,” *Medical Image Analysis*, vol. 8, no. 3, pp. 325 – 342, 2004.
- [56] J. C. Spall, “Implementation of the simultaneous perturbation method for stochastic optimization,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, no. 3, pp. 817 – 823, 1998.
- [57] D. Rey, G. Subsol, H. Delingette, and N. Ayache, “Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis,” *Medical Image Analysis*, vol. 6, no. 2, pp. 163 – 179, 2002.
- [58] S. Hu, E. A. Hoffman, and J. M. Reinhardt, “Automatic lung segmentation for accurate quantitation of volumetric X-Ray CT images,” *IEEE Transactions on Medical Imaging*, vol. 20, no. 6, pp. 490 – 498, 2001.

- [59] I. C. Sluimer, M. Prokop, and B. van Ginneken, "Towards automated segmentation of the pathological lung in CT," *IEEE Transactions on Medical Imaging*, vol. 24, no. 8, pp. 1025 – 1038, 2005.
- [60] C. Studholme, C. Drapaca, B. Iordanova, and V. Cardenas, "Deformation-based mapping of volume change from serial brain MRI in the presence of local tissue contrast change," *IEEE Transactions on Medical Imaging*, vol. 25, no. 5, pp. 626–639, 2006.
- [61] J. Ashburner and K. Friston, "Why voxel-based morphometry should be used," *NeuroImage*, vol. 14, no. 6, pp. 1238 – 1243, 2001.
- [62] F. Bookstein, "'Voxel-based morphometry' should not be used with imperfectly registered images," *NeuroImage*, vol. 14, no. 6, pp. 1454 – 1462, 2001.
- [63] G. E. Christensen and H. J. Johnson, "Consistent image registration," *IEEE Transactions on Medical Imaging*, vol. 20, no. 7, pp. 568 – 582, 2001.
- [64] E. Haber and J. Modersitzki, "Numerical methods for volume preserving image registration," *Inverse Problems*, vol. 20, no. 5, pp. 1621 – 1638, 2004.
- [65] D. Ruan, J. A. Fessler, M. Roberson, J. Balter, and M. Kesler, "Nonrigid registration using regularization that accomodates local tissue rigidity," in *SPIE Medical Imaging: Image Processing*, ser. Proceedings of SPIE, vol. 6144. SPIE Press, 2006, pp. 346 – 354.
- [66] M. Staring, S. Klein, and J. P. W. Pluim, "Nonrigid registration using a rigidity constraint," in *SPIE Medical Imaging: Image Processing*, ser. Proceedings of SPIE, vol. 6144. SPIE Press, 2006, pp. 355 – 364.
- [67] S. Klein, M. Staring, and J. P. W. Pluim, "Comparison of gradient approximation techniques for optimisation of mutual information in nonrigid registration," in *SPIE Medical Imaging: Image Processing*, ser. Proceedings of SPIE, vol. 5747. SPIE Press, 2005, pp. 192 – 203.
- [68] C. Henschke, D. McCauley, D. Yankelevitz, D. Naidich, G. McGuinness, O. Miettinen, D. Libby, M. Pasmantier, J. Koizumi, N. Altorki, and J. Smith, "Early lung cancer action project: overall design and findings from baseline screening," *Lancet*, vol. 354, pp. 99 – 105, 1999.
- [69] C. Henschke, D. Yankelevitz, D. Libby, M. Pasmantier, J. Smith, and O. Miettinen, "Survival of patients with stage i lung cancer detected on ct screening," *New England Journal of Medicine*, vol. 355, pp. 1763 – 1771, 2006.
- [70] "National lung screening trial," Available at [www.cancer.gov/nlst](http://www.cancer.gov/nlst), National Cancer Institute, accessed December 27, 2007.
- [71] K. Marten and C. Engelke, "Computer-aided detection and automated ct volumetry of pulmonary nodules," *Eur Radiol*, vol. 17, pp. 888 – 901, 2007.



- [72] M. Nakata, S. Sawada, H. Saeki, S. Takashima, H. Mogami, N. Teramoto, and K. Eguchi, "Prospective study of thoracoscopic limited resection for ground-glass opacity selected by computed tomography," *The Annals of Thoracic Surgery*, vol. 75, pp. 1601 – 1605, 2003.
- [73] R. Nakajima, T. Yokose, R. Kakinuma, K. Nagai, Y. Nishiwaki, and A. Ochiai, "Localized pure ground-glass opacity on high-resolution ct: histologic characteristics," *Journal of Computer Assisted Tomography*, vol. 26, pp. 323 – 329, 2002.
- [74] K. Suzuki, H. Asamura, M. Kusumoto, H. Kondo, and R. Tsuchiya, "'early" peripheral lung cancer: prognostic significance of ground-glass opacity on thin-section computed tomographic scan," *The Annals of Thoracic Surgery*, vol. 74, pp. 1635 – 1639, 2002.
- [75] C. Henschke, D. Yankelevitz, R. Mirtcheva, G. McGuinness, D. McCauley, and O. Miettinen, "CT screening for lung cancer: frequency and significance of part-solid and nonsolid nodules," *American Journal of Roentgenology*, vol. 178, pp. 1053 – 1057, 2002.
- [76] D. Arenberg, "Bronchioloalveolar lung cancer: ACCP evidence-based clinical practice guidelines (2nd edition)," *Chest*, vol. 132, pp. 306 – 313, September 2007.
- [77] Z. Yang, S. Sone, S. Takashima, F. Li, T. Honda, Y. Maruyama, M. Hasegawa, and S. Kawakami, "High-resolution CT analysis of small peripheral lung adenocarcinomas revealed on screening helical CT," *American Journal of Roentgenology*, vol. 176, no. 6, pp. 1399 – 1407, June 2001.
- [78] C. Henschke, D. Yankelevitz, J. Smith, and O. Miettinen, "Screening for lung cancer: the early lung cancer action approach," *Lung Cancer*, vol. 35, pp. 143 – 148, 2002.
- [79] R. Kakinuma, H. Ohmatsu, M. Kaneko, M. Kusumoto, J. Yoshida, K. Nagai, Y. Nishiwaki, T. Kobayashi, R. Tsuchiya, H. Nishiyama, E. Matsui, K. Eguchi, and N. Moriyama, "Progression of focal pure ground-glass opacity detected by low-dose helical computed tomography screening for lung cancer," *Journal of Computer Assisted Tomography*, vol. 28, pp. 17 – 23, 2004.
- [80] H. Nakamura, H. S. H, A. Ogata, T. Saijo, S. Okada, and H. Kato, "Lung cancer patients showing pure ground-glass opacity on computed tomography are good candidates for wedge resection," *Lung Cancer*, vol. 44, pp. 61 – 68, 2004.
- [81] H. Gietema, C. Schaefer-Prokop, W. Mali, G. Groenewegen, and M. Prokop, "Pulmonary nodules: Interscan variability of semiautomated volume measurements with multisection CT - influence of inspiration level, nodule size, and segmentation performance," *Radiology*, vol. 245, pp. 888 – 894, 2007.

- [82] D. Wormanns, G. Kohl, E. Klotz, A. Marheine, F. Beyer, W. Heindel, and S. Diederich, "Volumetric measurements of pulmonary nodules at multi-row detector ct: in vivo reproducibility," *European Radiology*, vol. 14, no. 1, pp. 86 – 92, January 2004.
- [83] W. Kostis, D. Yankelevitz, A. Reeves, S. Fluture, and C. Henschke, "Small pulmonary nodules: reproducibility of three-dimensional volumetric measurement and estimation of time to follow-up CT," *Radiology*, vol. 231, pp. 446 – 452, 2004.
- [84] L. Goodman, M. Gulsun, L. Washington, P. Nagy, and K. Piacsek, "Inherent variability of CT lung nodule measurements in vivo using semiautomated volumetric measurements," *American Journal of Roentgenology*, vol. 186, pp. 989 – 994, 2006.
- [85] J. Zhou, S. Chang, D. Metaxas, B. Zhao, L. Schwartz, and M. Ginsberg, "Automatic detection and segmentation of ground glass opacity nodules," in *MICCAI*, ser. Lecture Notes In Computer Science, vol. 4190, 2006, pp. 784 – 791.
- [86] D. Xu, H. Gietema, H. de Koning, R. Vernhout, K. Nackaerts, M. Prokop, C. Weenink, J. Lammers, H. Groen, and M. Oudkerk, "Nodule management protocol of the NELSON randomised lung cancer screening trial," *Lung Cancer*, vol. 54, no. 2, pp. 177 – 184, November 2006.
- [87] J. Cohen, "Weighted kappa: Nominal scale agreement with provision for scaled disagreement or partial credit," *Psychological Bulletin*, vol. 70, pp. 213 – 220, 1968.
- [88] D. G. Altman, *Practical Statistics for Medical Research*. London: Chapman & Hall / CRC, 1991.
- [89] R. J. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, no. 1, pp. 159 – 174, 1977.
- [90] J. J. Erasmus, G. W. Gladish, L. Broemeling, B. S. Sabloff, M. T. Truong, R. S. Herbst, and R. F. Munden, "Interobserver and intraobserver variability in measurement of nonsmall-cell carcinoma lung lesions: Implications for assessment of tumor response," *Journal of Clinical Oncology*, vol. 21, no. 13, pp. 2574 – 2582, 2003.
- [91] M. Revel, A. Bissery, M. Bienvenu, L. Aycard, C. Lefort, and G. Frija, "Are two-dimensional CT measurements of small non-calcified pulmonary nodules reliable?" *Radiology*, vol. 231, pp. 453 – 458, 2004.
- [92] S. Jennings, H. Winer-Muram, R. Tarver, and M. Farber, "Lung tumor growth: assessment with CT - comparison of diameter and cross-sectional area with volume measurements," *Radiology*, vol. 231, pp. 866 – 871, 2004.
- [93] H. Saito, Y. Minamiya, H. Kawai, T. Nakagawa, M. Ito, Y. Hosono, S. Motoyama, M. Hashimoto, K. Ishiyama, and J. Ogawa, "Usefulness of circumference difference for estimating the likelihood of malignancy in small solitary pulmonary nodules on CT," *Lung Cancer*, vol. 58, pp. 348 – 354, 2007.

- [94] H. Sumikawa, T. Johkoh, T. Nagareda, J. Sekiguchi, K. Matsuo, Y. Fujita, J. Nat-sag, A. Inoue, N. Mihara, O. Honda, N. Tomiyama, M. Minami, M. Okumura, and H. Nakamura, "Pulmonary adenocarcinomas with ground-glass attenuation on thin-section CT: Quantification by three-dimensional image analyzing method," *European Journal of Radiology*, vol. 65, no. 1, pp. 104 – 111., 2008.
- [95] D. M. Parkin, F. Bray, J. Ferlay, and P. Pisani, "Global cancer statistics, 2002," *CA: A Cancer Journal for Clinicians*, vol. 55, no. 2, pp. 74 – 108, 2005.
- [96] R. Pötter, C. Haie-Meder, E. V. Limbergen, I. Barillot, M. D. Brabandere, J. Di-mopoulos, I. Dumas, B. Erickson, S. Lang, A. Nulens, P. Petrow, J. Rownd, and C. Kirisits, "Recommendations from gynaecological (GYN) GEC ESTRO working group (II): concepts and terms in 3D image-based treatment planning in cervix cancer brachytherapy-3D dose volume parameters and aspects of 3D image-based anatomy, radiation physics, radiobiology," *Radiotherapy and Oncology*, vol. 78, no. 1, pp. 67 – 77, 2006.
- [97] R. Pötter, J. Dimopoulos, P. Georg, S. Lang, C. Waldhäusl, N. Wachter-Gerstner, H. Weitmann, A. Reinthaller, T. H. Knocke, S. Wachter, and C. Kirisits, "Clinical impact of MRI assisted dose volume adaptation and dose escalation in brachyther-apy of locally advanced cervix cancer," *Radiotherapy and Oncology*, vol. 83, no. 2, pp. 148 – 155, 2007.
- [98] L. van de Bunt, U. A. van der Heide, M. Ketelaars, G. A. de Kort, and I. M. Jürgenliemk-Schulz, "Conventional, conformal, and intensity-modulated radiation therapy treatment planning of external beam radiotherapy for cervical cancer: The impact of tumor regression," *International Journal of Radiation Oncology Biology Physics*, vol. 64, no. 1, pp. 189 – 196, 2006.
- [99] L. van de Bunt, I. M. Jürgenliemk-Schulz, G. A. de Kort, J. M. Roesink, R. J. Tersteeg, and U. A. van der Heide, "Motion and deformation of the target volumes during IMRT for cervical cancer: What margins do we need?" *Radiotherapy and Oncology*, 2008, in press.
- [100] M. R. Kaus, K. K. Brock, V. Pekar, L. A. Dawson, A. M. Nichol, and D. A. Jaffray, "Assessment of a model-based deformable image registration approach for radiation therapy planning," *International Journal of Radiation Oncology Biology Physics*, vol. 68, no. 2, pp. 572 – 580, 2007.
- [101] K. K. Brock, L. A. Dawson, M. B. Sharpe, D. J. Moseley, and D. A. Jaffray, "Feasi-bility of a novel deformable image registration technique to facilitate classification, targeting, and monitoring of tumor and normal tissue," *International Journal of Radiation Oncology Biology Physics*, vol. 64, no. 4, pp. 1245 – 1254, 2006.
- [102] K. S. C. Chao, S. Bhide, H. Chen, J. Asper, S. Bush, G. Franklin, V. Kavadi, V. Liengswangwong, W. Gordon, A. Raben, J. Strasser, C. Koprowski, S. Frank, G. Chronowski, A. Ahamad, R. Malyapa, L. Zhang, and L. Dong, "Reduce in

- variation and improve efficiency of target volume delineation by a computer-assisted system using a deformable image registration approach,” *International Journal of Radiation Oncology Biology Physics*, vol. 68, no. 5, pp. 1512 – 1521, 2007.
- [103] K. K. Brock, D. J. Moseley, J. Stewart, K. Lim, P. Chan, R. Dinniwell, Y. Cho, A. Fyles, M. Milosevic, and D. A. Jaffray, “Deformable modeling of physiological motion and treatment response of cervical cancer patients,” in *International Conference on the Use of Computers in Radiation Therapy*, 2007.
- [104] C. López and S. Chakravarti, “Imaging of cervical cancer,” *Imaging*, vol. 18, no. 1, pp. 10 – 19, 2006.
- [105] D. G. Mitchell, B. Snyder, F. Coakley, C. Reinhold, G. Thomas, M. Amendola, L. H. Schwartz, P. Woodward, H. Pannu, and H. Hricak, “Early invasive cervical cancer: Tumor delineation by magnetic resonance imaging, computed tomography, and clinical examination, verified by pathologic results in the ACRIN 6651/GOG 183 intergroup study,” *Journal Of Clinical Oncology*, vol. 24, no. 36, pp. 5687 – 5693, 2006.
- [106] J. Dimopoulos, G. Schard, D. Berger, S. Lang, G. Goldner, T. Helbich, and R. Pötter, “Systematic evaluation of MRI findings in different stages of treatment of cervical cancer: Potential of MRI on delineation of target, pathoanatomic structures, and organs at risk,” *International Journal of Radiation Oncology Biology Physics*, vol. 64, no. 5, pp. 1380 – 1388, 2006.
- [107] A. Hero, B. Ma, O. Michel, and J. Gorman, “Applications of entropic spanning graphs,” *IEEE Signal Processing Magazine*, vol. 19, no. 5, pp. 85 – 95, September 2002.
- [108] H. Neemuchwala and A. Hero, *Entropic Graphs for Registration*, 2005, pp. 185 – 235.
- [109] H. Neemuchwala, A. Hero, and P. Carson, “Image matching using alpha-entropy measures and entropic graphs,” *Signal Processing*, vol. 85, no. 2, pp. 277 – 296, February 2005.
- [110] H. F. Neemuchwala and A. O. Hero, “Image registration in high dimensional feature space,” in *SPIE Electronic Imaging: Computational Imaging III*, ser. Proceedings of SPIE, C. A. Bouman and E. L. Miller, Eds., vol. 5674. San Jose, California, USA: SPIE Press, March 2005, pp. 99 – 113.
- [111] E. Oubel, A. F. Frangi, and A. Hero, “Complex wavelets for registration of tagged MRI sequences,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2006, pp. 622 – 625.
- [112] B. Ma, R. Narayanan, H. Park, A. O. Hero, P. H. Bland, and C. R. Meyer, “Comparing pairwise and simultaneous joint registrations of decorrelating interval exams

- using entropic graphs,” in *Information Processing in Medical Imaging*, ser. Lecture Notes in Computer Science, vol. 4584, 2007, pp. 270 – 282.
- [113] M. R. Sabuncu and P. J. Ramadge, “Using spanning graphs for efficient image registration,” *IEEE Transactions on Image Processing*, vol. 17, no. 5, pp. 788 – 797, 2008.
- [114] E. Oubel, M. de Craene, M. Gazzola, A. Hero, and A. F. Frangi, “Multiview registration of cardiac tagging MRI sequences,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, Arlington, VA, USA, 2007, pp. 388 – 391.
- [115] A. Collignon, “Multi-modality medical image registration by maximization of mutual information,” Ph.D. dissertation, Catholic University of Leuven, Leuven, Belgium, 1998.
- [116] J. Beardwood, J. H. Halton, and J. M. Hammersley, “The shortest path through many points,” in *Proceedings Cambridge Philosophical Society*, vol. 55, 1959, pp. 299 – 327.
- [117] A. Hero and O. Michel, “Asymptotic theory of greedy approximations to minimal  $k$ -point random graphs,” *IEEE Transactions on Information Theory*, vol. 45, no. 6, pp. 1921 – 1939, 1999.
- [118] A. Rényi, “On measures of entropy and information,” in *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability*, vol. 1, Berkeley, CA: University of California Press, 1961, pp. 547 – 561.
- [119] L. Florack, “The syntactical structure of scalar images,” Ph.D. dissertation, University of Utrecht, 1993.
- [120] D. Rueckert, M. J. Clarkson, D. L. G. Hill, and D. J. Hawkes, “Non-rigid registration using higher-order mutual information,” in *SPIE Medical Imaging: Image Processing*, ser. Proceedings of SPIE, M. Sonka and J. M. Fitzpatrick, Eds., vol. 3979. San Diego, CA, USA: SPIE Press, 2000, pp. 438 – 447.
- [121] D. B. Russakoff, C. Tomasi, T. Rohlfing, and C. R. Maurer Jr., “Image similarity using mutual information of regions,” in *European Conference on Computer Vision (ECCV)*, ser. Lecture Notes in Computer Science, T. Pajdla and J. Matas, Eds., vol. 3023. Prague, Czech Republic: Springer Verlag, 2004, pp. 596 – 607.
- [122] D. Tomaževič, B. Likar, and F. Pernuš, “Multi-feature mutual information,” in *SPIE Medical Imaging: Image Processing*, ser. Proceedings of SPIE, J. M. Fitzpatrick and M. Sonka, Eds., vol. 5370. San Diego, CA, USA: SPIE, 2004, pp. 143 – 154.



# Samenvatting

**B**EELDREGISTRATIE is belangrijk voor het vergelijken van medische beelden. Vergelijken is klinisch relevant voor de analyse van verschillen die ontstaan door veranderingen in de gezondheid van een patiënt. Bijvoorbeeld, wanneer een ziekte op twee tijdpunten in beeld wordt gebracht, wil men weten of het stabiel is, of dat er regressie of progressie is opgetreden. Een ander voorbeeld is het scannen van een patiënt voor en na behandeling. In dat geval wil men weten in welke mate de behandeling succesvol was. De klinische actie die genomen wordt, is afhankelijk van de uitkomst van de analyse van de verschillen.

Dit proefschrift behandelt beeldregistratietechnieken ten behoeve van de analyse van veranderingen in medische beelden. We onderscheiden drie bijdrages van het werk beschreven in dit proefschrift, die behandeld worden in de volgende drie paragrafen. Tot slot bespreken we nog de klinische relevantie van het werk.

## `elastix`

Samen met Stefan Klein is er software ontwikkeld dat het onderzoek naar beeldregistratie ondersteunt. Het software pakket heet `elastix` en is beschreven in **Hoofdstuk 2**. Een algoritme voor beeldregistratie bestaat uit verschillende delen, zoals de kostenfunctie, de optimalisatieroutine en de transformatie. Voor ieder deel zijn er veel keuzes beschikbaar, maar het is van tevoren niet bekend welke keuze optimaal is voor een bepaald probleem. `elastix` is een verzameling van algoritmes die beschikbaar zijn in de literatuur. Het vergemakkelijkt de selectie, het testen en het vergelijken van de verschillende methoden. De command-line interface maakt het mogelijk om grote hoeveelheden data te verwerken, door gebruik te maken van scripts. `elastix` is gebaseerd op de Insight Toolkit en is open source (wordt verspreid onder the BSD licentie). De software werd gebruikt voor alle experimenten die beschreven zijn in dit proefschrift en wordt steeds meer gebruikt, zowel binnen als buiten het Image Sciences Institute.

## Weefselafhankelijke Beeldregistratie en de Detectie van Veranderingen

De populaire, vaak gebruikte niet-rigide beeldregistratiealgoritmes houden geen rekening met de stijfheid van verschillende weefseltypes. Dit kan tot ongewenste effecten leiden in drie situaties:

- a. Het beeld bevat verscheidene weefseltypes, ieder met zijn eigen mechanische stijfheid.
- b. Veranderingen van structuren over de tijd moeten gevisualiseerd worden.

- c. Tussen acquisities verandert de zichtbaarheid van structuren doordat er contrast-materiaal gebruikt wordt.

In alle drie de situaties moeten de objecten waar men geïntereeserd in is door het registratiealgoritme beschouwd worden als lokaal rigide, om zo de analyse van veranderingen mogelijk te maken door middel van beeldsubtractie.

In **Hoofdstuk 3** wordt dit probleem aangepakt door het deformatieveld lokaal aan te passen bij structuren die rigide moeten blijven. Dit wordt bereikt met een weefselafhankelijke filtertechniek. Dit adaptieve filteren van het deformatieveld resulteert in lokaal lineaire transformaties zonder schaling. De mate van filtering is gerelateerd aan de stijfheid van het weefsel: er wordt meer gefilterd als het weefsel stijf is, minder in delen van het beeld die zacht weefsel bevatten.

Evaluatie van de voorgestelde weefselafhankelijke filtering werd gedaan op 3D CT data van de borstkas en op 2D Digitale Subtractie Angiografie (DSA) beelden. Er werd aangetoond dat het adaptieve filter tot verbeterde registratieresultaten leidt: tumorvolume en de wijdte van bloedvaten bleven behouden. Gemiddeld bleef 99% van het tumorvolume behouden, vergeleken met 95% met een standaard methode die geen rekening houdt met de rigiditeit van de tumor. Voor bloedvaten bleef 98% van de diameter behouden, vergeleken met 85% met het standaard algoritme. Het filter heeft als nadeel dat wanneer het te extreem wordt toegepast, rotaties van rigide objecten niet worden toegestaan.

Een rigiditeitstrafterm heeft dit nadeel niet. De constructie van deze term wordt beschreven in **Hoofdstuk 4**. De term is afgeleid van de definitie van een rigide transformatie en bestaat uit drie verschillende termen die gerelateerd zijn aan het affien zijn van de transformatie en de orthonormaliteit en “properness” van de rotatiematrix. De drie voorwaarden zijn gecombineerd in één term. Plaatsafhankelijkheid wordt geïntroduceerd via een gewichtsfunctie, die afhankelijk is van de spatiële locatie. Lokale rigiditeit wordt opgenomen in de registratie door de strafterm toe te voegen aan de kostenfunctie. Een snel algoritme wordt verkregen door de B-spline parameterisatie van de transformatie uit te buiten.

Evaluatie werd weer gedaan op de 3D CT data van de borstkas en de 2D angio data. Het tumorvolume nam gemiddeld af tot 78% van het originele volume met een standaard algoritme; voor het algoritme met de rigiditeitstrafterm bleef 98% van het volume behouden. Voor de bloedvaten verbeterde het diameterbehoud van 84% tot 99%. Uit deze resultaten concludeerden we dat de rigiditeitstrafterm in staat is om door de gebruiker bepaalde structuren lokaal rigide te houden tijdens het uitvoeren van een niet-rigide registratie.

Situatie b. van de lijst hierboven wordt in een klinisch kader verder geanalyseerd in **Hoofdstuk 5**. Om de prognose van de patiënt te verbeteren is het belangrijk in een vroeg stadium te weten of een long nodule goed- of kwaadaardig is. Veranderingen in grootte en dichtheid zijn twee belangrijke criteria voor kwaadaardigheid. Detectie van veranderingen in grootte of dichtheid is echter vaak moeilijk voor een specifiek type nodule, matglas nodules, door hun verschijningsvorm op CT. In de huidige klinische praktijk worden nodules



visueel vergeleken, met behulp van handmatige diametermetingen. We bestudeerden het effect van het toevoegen van een verschilbeeld aan de detectie van veranderingen van dit type nodules. Het verschilbeeld was gemaakt met het registratiealgoritme beschreven in Hoofdstuk 4.

Daartoe zijn vier waarnemers gevraagd de verandering van grootte en dichtheid te scoren voor ongeveer 60 paren van beelden die een matglas nodule bevatten, een keer met en een keer zonder het verschilbeeld. In het bijzonder voor nodules die subtiel veranderd waren, verbeterde het verschilbeeld de prestatie van de waarnemers substantieel vergeleken met de huidige klinische praktijk. Voor deze subtiele gevallen verbeterde de overeenkomst met de gouden standaard, gemeten met de kappa statistiek, van  $\kappa_w = 0.26$  tot 0.57 voor de detectie van veranderingen in grootte en van 0.19 tot 0.47 voor de detectie van veranderingen in dichtheid. Ook de overeenkomst tussen waarnemers verbeterde, wat er op wijst dat het verschilbeeld een beter reproduceerbare methode oplevert.

## Verbeterde Registratie

In Hoofdstuk 6 behandelden we de registratie van cervix data met behulp van ‘wederzijdse informatie’ van niet alleen de intensiteit, maar ook van kenmerken die de lokale structuur van een beeld beschrijven. Verschillende technische aspecten van de voorgestelde kostenfunctie werden behandeld om een acceptabele registratietijd te verkrijgen.

We hebben op uitdagende data (uitdagend door de beeldkwaliteit en de grote deformaties) van de afdeling Radiotherapie laten zien dat, vergeleken met een standaard algoritme gebaseerd op intensiteit alleen, het voorgestelde algoritme het significant beter doet. De verbetering was het meest duidelijk voor structuren die verschillend van grootte en vorm waren tussen de twee beelden: de mediaan van de overlap van de blaas na registratie nam toe van 0.75 tot 0.81.

## Klinische Relevantie

Het werk beschreven in dit proefschrift maakt het mogelijk om niet-rigide registratie te doen, terwijl door de gebruiker gedefinieerde structuren lokaal rigide gehouden worden. Zoals Hoofdstuk 5 laat zien, kan dit gebruikt worden om veranderingen van matglas nodules te visualiseren. De methode maakt het mogelijk om veranderingen in een vroeg stadium op te sporen, wat de prognose voor de patiënt verbetert en zou kunnen helpen om levens te redden.

Onderzoek is nodig om het algoritme verder te automatiseren, omdat nu een manuele segmentatie van de regio die rigide gehouden moet worden (de nodule) nodig is. Omdat alleen een ruwe segmentatie nodig is, is het waarschijnlijk wel mogelijk dit te automatiseren met behulp van standaard segmentatie algoritmes. Als alternatief zou het onderzoek gericht kunnen worden op compleet automatische en accurate segmentatie van deze matglas nodules, om vervolgens de volumes te gebruiken om veranderingen te detecteren. Het zou interessant zijn om te zien hoe deze methode zich verhoudt tot de subtractiemethode die vertrouwt op het menselijk oog. Voor sterk verdichte (solid) nodules bestaan er al technieken voor (semi-)automatische volumemetingen, omdat deze nodules makkelijker te

segmenteren zijn. Voor dit type nodules zou een vergelijkende studie van beide methodes direct kunnen starten.

De afdeling Radiotherapie van het Universitair Medisch Centrum Utrecht ontwikkelt een geïntegreerde MRI-versneller. Dit systeem maakt het mogelijk om een doelgebied (b.v. de cervix) te scannen en het tegelijkertijd te bestralen. Een potentieel voordeel van een dergelijk systeem vergeleken met de huidige methodologie is dat het richten van de dosis meer accuraat is. Dit maakt het mogelijk om een hogere dosis te gebruiken voor het doelgebied en om weefsel in de nabijheid van de tumor te sparen, wat de complicaties voor de patiënt vermindert. Voor het richten van de dosis zijn echter segmentaties van de relevante organen nodig. Manuele segmentatie is geen optie, door tijdsrestricties. Automatiseren kan door middel van niet-rigide beeldregistratie, maar deze technieken zijn nog niet accuraat genoeg. In Hoofdstuk 6 wordt er een stap genomen in de richting van klinisch acceptabele accuraatheid van registratie.

Concluderend, hoewel niet-rigide beeldregistratie nog niet wijd verspreid is in de kliniek, kan het van toegevoegde waarde zijn voor de analyse van veranderingen in medische beelden.

# Publications

## Publications in international journals

- I. Išgum, **M. Staring**, A. Rutten, M. Prokop, M.A. Viergever, and B. van Ginneken, “Multi-atlas-based segmentation with local decision fusion - application to cardiac and aortic segmentation in CT scans”, *submitted*.
- S. Klein, **M. Staring**, P. Andersson, and J.P.W. Pluim, “Preconditioned Stochastic Gradient Descent Optimisation for Monomodal Image Registration”, *submitted*.
- S. Klein, J.P.W. Pluim, **M. Staring**, and M.A. Viergever, “Adaptive Stochastic Gradient Descent Optimisation for Image Registration”, *International Journal of Computer Vision*, *in press*.
- **M. Staring**, U.A. van der Heide, S. Klein, M.A. Viergever, and J.P.W. Pluim, “Registration of Cervical MRI Using Multifeature Mutual Information”, *submitted*.
- S. Klein, **M. Staring**, K. Murphy, M.A. Viergever, and J.P.W. Pluim, “**elastix**: a toolbox for intensity-based medical image registration”, *submitted*.
- **M. Staring**, J.P.W. Pluim, B. de Hoop, S. Klein, B. van Ginneken, H. Gietema, G. Nossent, C. Schaefer-Prokop, S. van de Vorst, and M. Prokop, “Influence of Image Subtraction on the Assessment of Volume and Density Change in Ground-Glass Opacities in Chest CT”, *submitted*.
- S. Klein, U.A. van der Heide, I.M. Lips, M. van Vulpen, **M. Staring**, and J.P.W. Pluim, “Automatic Segmentation of the Prostate in 3D MR Images by Atlas Matching using Localized Mutual Information”, *Medical Physics*, vol. 35, no. 4, pp. 1407-1417, April 2008.
- S. Klein, **M. Staring**, and J.P.W. Pluim, ” “Evaluation of Optimization Methods for Nonrigid Medical Image Registration using Mutual Information and B-splines”, *IEEE Transactions on Image Processing*, vol. 16, no. 12, pp. 2879 - 2890, December 2007.
- **M. Staring**, S. Klein, and J.P.W. Pluim, “A Rigidity Penalty Term for Nonrigid Registration”, *Medical Physics*. vol. 34, no. 11, pp. 4098 - 4108, November 2007.
- **M. Staring**, S. Klein, and J.P.W. Pluim, “Nonrigid Registration with Tissue-Dependent Filtering of the Deformation Field”, *Physics in Medicine and Biology*, vol. 52, no. 23, pp. 6879 - 6892, December 2007.

## Conference contributions

- K. Murphy, B. van Ginneken, S. Klein, **M. Staring**, and J.P.W. Pluim, “Semi-Automatic Reference Standard Construction for Quantitative Evaluation of Lung CT Registration”, in: Medical Image Computing and Computer-Assisted Intervention, 2008, *accepted*.
- E.M. van Rikxoort, I. Išgum, **M. Staring**, S. Klein, and B. van Ginneken, “Adaptive Local Multi-Atlas Segmentation: Application to Heart Segmentation in Chest CT Scans”, in: SPIE Medical Imaging: Image Processing, *in press*.
- S. Klein, U.A. van der Heide, B.W. Raaymakers, A.N.T.J. Kotte, **M. Staring**, and J.P.W. Pluim, “Segmentation of the prostate in MR images by atlas matching”, in: International Symposium on Biomedical Imaging, 2007, pp. 1300-1303.
- S. Klein, U. A. van der Heide, **M. Staring**, A. N. T. J. Kotte, B. W. Raaymakers, and J.P.W. Pluim, “Segmentation of the Prostate in MR images by Atlas Matching using Localised Mutual Information”, in: XVth International Conference on the use of Computers in Radiation Therapy, Editor(s): D. A. Jaffray, M. Sharpe, J. van Dyk, J. P. Bissonnette, 2007, vol. 2, pp. 585-589.
- **M. Staring**, S. Klein, and J.P.W. Pluim, “Nonrigid Registration Using a Rigidity Constraint”, in: SPIE Medical Imaging: Image Processing, Editor(s): J.M. Reinhardt, J.P.W. Pluim, SPIE press, 2006, vol. 6144, Proceedings of SPIE, pp. 614413-1 - 614413-10.
- S. Klein, **M. Staring**, and J.P.W. Pluim, “A Comparison of Acceleration Techniques for Nonrigid Medical Image Registration”, in: International Workshop on Biomedical Image Registration, Editor(s): J.P.W. Pluim, B. Likar, F.A. Gerritsen, Springer-Verlag Berlin Heidelberg, 2006, vol. 4057, Lecture Notes in Computer Science, pp. 151-159.
- **M. Staring**, S. Klein, and J.P.W. Pluim, “Evaluation of a Rigidity Penalty Term for Nonrigid Registration”, in: Workshop on Image Registration in Deformable Environments, Editor(s): A. Bartoli, N. Navab, V. Lepetit, 2006, pp. 41 - 50.
- H.J. Noordmans, R. de Roode, **M. Staring**, and R. Verdaasdonk, “Registration and analysis of in-vivo multispectral images for correction of motion and comparison in time”, in: SPIE Photonics West: Multimodal Biomedical Imaging, Editor(s): F.S. Azar, D.N. Metaxas, 2006, vol. 6081, Proceedings of SPIE, pp. 35 - 43.
- **M. Staring**, S. Klein, and J.P.W. Pluim, “Nonrigid Registration with Adaptive, Content-Based Filtering of the Deformation Field”, in: SPIE Medical Imaging: Image Processing, Editor(s): J.M. Fitzpatrick, J.M. Reinhardt, SPIE press, 2005, vol. 5747, Proceedings of SPIE, pp. 212 - 221.

- S. Klein, **M. Staring**, and J.P.W. Pluim, “Comparison of gradient approximation techniques for optimisation of mutual information in nonrigid registration”, in: SPIE Medical Imaging: Image Processing, Editor(s): J.M. Fitzpatrick, J.M. Reinhardt, SPIE press, 2005, vol. 5747, Proceedings of SPIE, pp. 192-203.
- J.C. Oostveen, A.A.C. Kalker, and **M. Staring**, “Adaptive quantization watermarking”, in: Proceedings of SPIE: Security, Steganography, and Watermarking of Multimedia Contents VI, 2004, vol. 5306.
- **M. Staring**, J.C. Oostveen, and A.A.C. Kalker, “Optimal distortion compensation for quantization watermarking”, in: International Conference on Image Processing, 2003, vol. 2, pp. 727 - 730.

## Abstracts

- **M. Staring**, J.P.W. Pluim, B. de Hoop, S. Klein, and M. Prokop, “Subtraction imaging for improved detection of change in ground glass nodules in chest computed tomography”, European Congress of Radiology.



# Dankwoord

Promoveren doe je niet alleen.

Ik ben veel mensen dank verschuldigd voor hun hulp, in allerlei vorm, voor de afgelopen jaren. Allereerst wil ik Max bedanken, omdat hij genoeg vertrouwen in me had om me aan te nemen. Tijdens de promotie was je altijd op een goede manier gericht op het halen van allerlei doelen (als promovendus dus vooral publiceren). Dat hielp mij om de grote lijn vast te houden. Ik heb je constructieve houding erg gewaardeerd.

Josien, onze wekelijkse ontmoetingen begonnen stevast met allerlei gepraat over dingen die niets met promoveren te maken hebben. Als we dan uiteindelijk over het lopende onderzoek begonnen, begreep je altijd snel wat ik erg brak probeerde uit te leggen. En wat er nog mis mee was helaas. Je had ook altijd wel een verhaal uit je eigen promotietijd klaar, wat stimuleerde of waarmee allerlei publicatieleed of promotiestress toch wat minder erg werd. Bedankt ook dat je altijd tijd voor me had, en dat versies van papers lekker snel weer terug kwamen.

Stefan, we zijn op hetzelfde moment op hetzelfde project begonnen en hebben de hele tijd een kantoor gedeeld. En niet alleen op het werk tijd doorgebracht, maar ook daarbuiten. De onderwerpen die we bespraken, gingen over van alles: wetenschap, samenleving, religie, en tal van andere dingen. Ook jij bent van ontzettend veel waarde geweest voor mijn boekje. Dank je wel dat ik bijna elke dag je brein wel mocht gebruiken voor mijn eigen doeleinden. Ik denk met veel plezier terug aan de etentjes, congressen, woestijnreizen, en biertjes. En ik ben weer een vriend rijker.

Bartjan, Hester, George, Cornelia, Saskia en Mathias wil ik bedanken voor hun bijdrage aan Hoofdstuk 5 van mijn boekje. Ontzettend bedankt voor de medische kennis die jullie met me deelden, voor het enthousiasme waarmee jullie met me mee dachten en voor me wilden ‘scoren’ en voor het verbeteren van het paper; een klinisch verhaal is toch een heel ander verhaal. Mathias, jij ook bedankt voor het meedenken toen ik op zoek was naar een toepassing van ontwikkelde technieken. Terwijl ik niet eens jouw aio was. Verder wil ik Uulke bedanken voor zijn bijdrage aan Hoofdstuk 6. Je bent altijd ontzettend snel met het regelen van data, het intekenen van structuren zodat ik een ‘waarheid’ tot mijn beschikking had, en erg geïnteresseerd in mogelijkheden tot automatisering. Dat was echt prima samenwerken!

Essentieel voor een succesvolle promotie zijn allerlei randvoorwaarden. De administratieve zaken werden uitstekend verzorgd door het secretariaat: Marjan, Jacqueline en Renée. En voor een praatje waren jullie ook altijd te porren. Gerard, bedankt voor het systeembeheer. Als ons netwerk en de numbercruchers niet zo goed geregeld waren, zou ik niet op deze manier hebben kunnen promoveren in dit vakgebied. En dat is jouw verdienste.

Volkan, ik heb het leuk gevonden jou te begeleiden tijdens je afstuderen. Dat was voor mij de eerste keer, en ook ik heb er veel van opgestoken. Ik hoop dat je inmiddels

een leuke baan hebt gevonden.

Als je wilt promoveren en het ook nog leuk wilt hebben, ga dan naar het ISI! Ik heb de collegialiteit en de relaxte, gezellige sfeer erg gewaardeerd. Als het effe niet zo lekker loopt me je promotie, dan zijn de koffiepauzes (bedankt voor al die potten koffie, Martijn) in ieder geval de moeite waard om je bed uit te komen. Ik heb een erg goede tijd gehad in het oude R-lab (weet je nog, toen Rashindra, Hans en Paul nog aio waren) en later op de aio-gang. Adriënne, Patrick, Martijn, Peter, Sara, Stefan, weer Rashindra, Sandra en later Jeroen, Gerrit, Casper, Joffrey en Maartje, ontzettend bedankt voor de gezellige tijd, de koffie, de grappen, het gepraat over alles en niets, en de koffie natuurlijk. Cynthia, Zeger, Kajo, Keelin, Sascha, Martijn, Stefan: AMOR dinners were always vèry pleasant. Josien kan dan ook erg goed koken, al is het dan vegetarisch. Ook de rest van de groep bedankt voor superatmosfeer. Een tip voor nieuwe mensen: zet een bank neer in je kamer.

Ook wil ik mijn familie bedanken voor wat ze betekenen in mijn leven. Papa en mama, jullie hebben me altijd gestimuleerd en geholpen bij wat ik ook maar aan het doen was. En nou schop ik het tot doctor. Niet te veel opscheppen hè! Femmie, dank je wel voor je steun, en natuurlijk dat je vorig jaar met me wilde trouwen. We hebben het goed samen. Femmie schatje, ik hald fan dei, foar altyd. Tot slot dank ik mijn Goede God: U heeft mijn leven rijk gezegend, U bent God, U bent mijn Heiland.

MARIUS



# Curriculum Vitae



Marius Staring was born March 14 1977 in Den Helder, the Netherlands. In 1996 he received his VWO diploma from the Openbare Scholengemeenschap Nieuwediep. He studied Applied Mathematics at the university of Twente, where he graduated in December 2002 with a thesis on quantisation based watermarking. This project was carried out at Philips Research Labs (Nat.Lab.) in Eindhoven. Since April 2003 he has been working as a PhD. student at the Image Science Institute, Utrecht, the results of which are described in this thesis. Since May 2008 he is with the Division of Image Processing (LKEB) in Leiden.