MEDICAL IMAGE REGISTRATION GUIDED BY APPLICATION-SPECIFIC GEOMETRY

FLORIS BERENDSEN

MEDICAL IMAGE REGISTRATION GUIDED BY APPLICATION-SPECIFIC GEOMETRY
PhD thesis, Utrecht University, The Netherlands

This thesis was typeset by the author using LaTeX 2$_\varepsilon$
ISBN: 978-94-6295-063-4
Cover design by Floris Berendsen
Printed by Proefschriftmaken.nl || Uitgeverij BOXPress

Financial support for publication of this thesis was kindly provided by Elekta Brachytherapy, Pie Medical Imaging, and the Röntgenstichting Utrecht.

Copyright © 2015 Floris Berendsen.
All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any other information storage and retrieval system, without the written permission from the copyright owner.

MEDICAL IMAGE REGISTRATION GUIDED BY APPLICATION-SPECIFIC GEOMETRY

MEDISCHE BEELDREGISTRATIE OP GELEIDE VAN APPLICATIESPECIFIEKE GEOMETRIE
(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. G.J. van der Zwaan, ingevolge het besluit van het college voor promoties in het openbaar te verdedigen op dinsdag 20 januari 2015 des middags te 4.15 uur

door

FLORIS FREDERIK BERENDSEN
geboren op 28 juli 1983 te Hengelo

Promotoren: prof. dr. J. P. W. Pluim
prof. dr. ir. M. A. Viergever

Contents

1 General Introduction **1**
Applications . . . 2
Registration process . . . 3
Outline . . . 5

2 Free-form image registration regularized by a statistical shape model **7**
2.1 Introduction . . . 9
2.2 Method . . . 10
Data . . . 11
Registration framework . . . 11
Statistical model . . . 13
Penalty . . . 14
2.3 Experiments . . . 15
Parameter selection experiments . . . 16
Evaluation experiments . . . 17
2.4 Results . . . 17
2.5 Discussion . . . 22
2.6 Conclusion . . . 26

3 Registration of Structurally Dissimilar Images **27**
3.1 Introduction . . . 29
3.2 Method . . . 30
3.3 Experimental Setup . . . 33
Datasets . . . 33
Registration Experiments . . . 34
Evaluation measures . . . 36
3.4 Experimental Results . . . 36
3.5 Discussion and Conclusion . . . 37

4 Registration of organs with sliding interfaces and changing topologies 43
4.1 Introduction . . . 45
Related work . . . 45
Aim . . . 46
4.2 Method . . . 46
Transformation Model . . . 46
Penalty . . . 47
Implementation details . . . 49
Surface point sets . . . 49
Neighborhood search . . . 49
4.3 Experiments . . . 50
Data . . . 50
Registration settings . . . 51
Evaluation . . . 51
4.4 Results . . . 52
4.5 Discussion . . . 52
4.6 Conclusion . . . 57

5 Efficient normalized cross-correlation for groupwise registration 61
5.1 Introduction . . . 63
5.2 Transformation model . . . 64
5.3 Image similarity metric . . . 64
5.4 Metric Derivative . . . 66
5.5 Partial overlaps . . . 67
5.6 Discussion and Conclusion . . . 68

6 General Summary and Discussion 71

Bibliography 77

Nederlandse samenvatting 85

Dankwoord 89

List of publications 91

Curriculum vitae 95

CHAPTER 1

General Introduction

Image registration is a fundamental task in medical image processing and analysis. Registration of images is the process of aligning the contents of images. If no external reference frame or calibration is available the images must be aligned based on the image content itself, i.e. the structures visible. The aim of the registration process is to establish a realistic correspondence that maps each point in one image to a point in the other image. The images that are to be registered can be from one patient acquired at different time points, e.g., to follow disease progression and/or response to treatment, or from different sources/modalities to fuse complementary information, or from different patients to analyze a population. Medical images are commonly three-dimensional, such as the well known MRI (magnetic resonance imaging) or CT (computed tomography). Without registration the analysis or visual comparison of three-dimensional scans can be very difficult to perform. In scans the global pose of the patient may have changed and/or local deformations of internal structures may have occurred. Therefore not only linear transformations, such as rigid and affine, are needed to model correspondence, but also non-linear transformations. Extensive surveys on methods for image registration can be found in [12, 21, 28, 30, 39, 53, 57].

In this thesis the focus is on parametric intensity-based registration methods. These currently well-established registration methods have many applications, but their application can be limited in certain situations. In this thesis we extend a conventional parametric intensity-based registration algorithm with new methods that tailor the algorithm to more application specific challenges. In section 1 we describe the mathematical details of the registration framework that is used throughout this thesis.

Applications

Although image registration has many applications, we describe some of the applications specifically addressed in this thesis in more detail. A general application of registration is to carry out image segmentation, e.g. labeling of the anatomical structures of interest in the image. Image segmentation is, for example, essential for the planning of radiotherapy treatments. The planning is often a (clinical) trade-off between a sufficient amount of radiation dose in the tumor region and minimizing the damage to neighboring organs by the radiation beam. In cervical cancer treatment, the so-called organs at risk, including the bladder and rectum, as well as the cervix/uterus need to be segmented in order to make such a planning. These segmentations can be made manually, by delineating the structures in each slice of the 3-D (MR) scan, but this is a time-consuming and costly exercise. Atlas-based segmentation is a technique to automatically create segmentations. In this context the atlas is an image of a patient together with an expert manual labeling and was acquired earlier. By registering the atlas image to the newly obtained target image, the labeling of the atlas can be propagated to the target image as well. For cervical

images, the challenge in registering different patients to each other is that the variation of anatomical shape and the position of the organs among different patients is very large (Chapter 2).

A radiotherapy treatment of cervical cancer typically consists of multiple fractions of radiation dose applied in daily or weekly sessions. With each session, the radiation dose at the tissues accumulates. To investigate the accumulated dose at each anatomical position accurately, the displacement of organs that occurred in between the sessions, must be taken into account. By registering the images of subsequent sessions, the dose distribution (images) of the planning can be propagated. The radiotherapy treatment may consist both of sessions with an external beam and of brachytherapy fractions, i.e. internal treatment by an applicator. Registering the images of a patient that underwent external beam therapy and subsequently brachytherapy is a difficult task, since the treatment applicator occupies a relatively large volume and is present in one of the images and not in the other (Chapter 3).

Organ motion in between scans can not only consist of local bending and stretching, but also of sliding organ surfaces which results in complex non-smooth deformations (Chapter 4). Sliding of surfaces can be observed at the rib-cage-lung boundary during a normal breathing cycle. When a correct mapping is established between the images at the maximum inhale phase and the maximum exhale phase, the lung ventilation properties can be inferred from the deformation field by calculating the local expansion or compression of the volume (i.e. the Jacobian determinant of the deformation field).

For the analysis of a group of patients or a time-series of scans it can be beneficial to register the images to a common reference space without choosing one of the images as a fixed reference. Such a groupwise registration can provide a spatial mean anatomy for the use in atlas creation or can reveal the spatial variations within the group (Chapter 5).

Registration process

All methods proposed in this thesis are applied within the framework of parametric-transformation intensity-based registration. The concepts and definitions within this framework will recur throughout the chapters of this thesis and therefore will be discussed in more detail in this section. Figure 1.1 shows a diagram of the processing concepts (modules) and their interconnections; the solid lines are part of a conventional registration process. The registration process is formulated as a minimization of a cost function to estimate the optimal transformation that aligns the images.

In registration one of the images is usually referred to as the moving (or source or reference) image, while the other is referred to as the fixed (or target) image. Mathematically, an image is defined as a (three-dimensional) volume in space (defined by domains Ω_F and Ω_M for a fixed and moving image, respec-

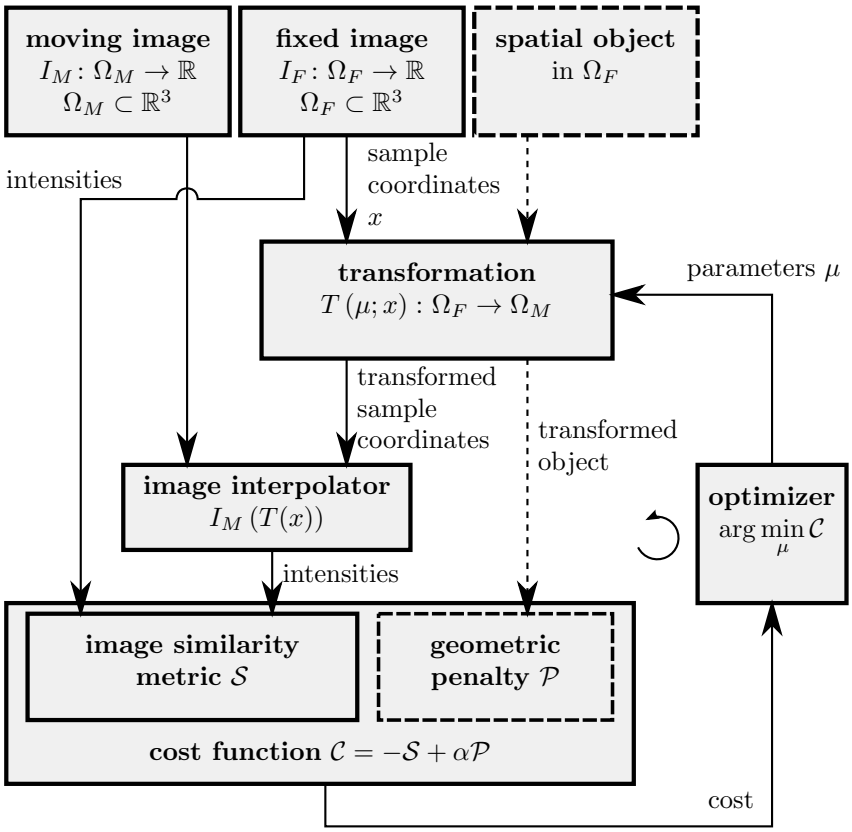

Figure 1.1: Registration framework

tively) containing the real-valued intensities ($\in \mathbb{R}$) of the voxels[1]. The moving image is warped to the fixed domain by a transformation and an interpolation step. First, the sample (voxel) coordinates of the fixed image are transformed to map to the moving image. Subsequently, the moving image is interpolated at each transformed coordinate in order to look up the intensities to create the warped image. The image similarity metric compares the warped image to the fixed image and calculates a similarity score. Various similarity metrics exist, we generally use normalized cross-correlation. The non-linear transformation model ($T(\mu; x)$) we use is a three-dimensional cubic B-spline, which is parame-

[1] A volume element, the 3-D analogy of a 'pixel'

terized by a vector of control points (μ). Initially, the setting of the parameter vector μ will not produce the correct alignment of the fixed and moving image, but by iteratively tuning the parameters the similarity is optimized.

In this thesis, a number of proposed methods (Chapters 2, 3 and 4) incorporate additional penalty terms in the cost function to positively influence the registration process for specific applications. Without specifying each of these in detail, the penalty terms are based on the deformation of spatial objects such as pointsets, triangulated surface meshes, or pointsets with (normal) direction vectors. The coordinates (and vectors) of the spatial object are transformed by the same function as the image coordinates. Conceptually, this may lead to confusion, since the spatial objects, defined in the fixed domain, are transformed to the moving domain, while simultaneously the moving image is warped to the fixed domain. Since the spatial objects we use are (derived from) manual delineations or atlasses, we may choose the atlas image as the fixed image and the image requiring a segmentation as moving image, which is in contrast to the conventions. Conventionally, the fixed image is referred to as target and the moving image as reference, since a segmentation (image) can be warped from moving to fixed. In order to create a segmentation of the moving image, the delineations in the form of surface meshes are transformed from the fixed domain to the moving domain.

All our proposed methods are implemented in the open source registration toolbox `elastix` [35], which provides the framework for intensity-based registration. Its modular design allows to configure the registration method for specific applications. In this thesis multiple geometric penalties are proposed (Chapters 2, 3 and 4), a new transformation model (Chapter 4) and a new similarity metric (Chapter 5).

Outline

The method proposed in Chapter 2 deals with inter-patient registration of patients with cervical cancer. Between patients are large variability in organ shape and position is observed that requires large and complex deformations. To guide registration in finding these deformations a statistical shape, trained on the shape of the segmentations of the population, is incorporated as penalty term in the optimization process.

In Chapter 3 intra-patient registration of the images acquired for external beam radiation therapy and brachytherapy is performed. The missing volume of the applicator, as used in cervical brachytherapy, is modeled as a surface mesh and its volume is minimized during registration.

For registration of images with sliding organs, in Chapter 4 we propose a new transformation model that is accompanied by a geometric penalty term. The proposed method is applied on inhale-exhale CT scans in which the lungs slide along the thoracic cage. Additionally, examples of registration of synthetic images and a registration of a patient with cervical cancer are given.

Chapter 5 proposes a improved similarity metric for groupwise registration. The computation of the groupwise metric that consists of the normalized cross-correlation between all image pairs in the group is reduced to a linear complexity.

In the last Chapter, Chapter 6, a summary and a general discussion is given.

CHAPTER 2

Free-form image registration regularized by a statistical shape model: application to organ segmentation in cervical MR

Abstract

Deformable registration is prone to errors when it involves large and complex deformations, since the procedure can easily end up in a local minimum. To reduce the number of local minima, and thus the risk of misalignment, regularization terms based on prior knowledge can be incorporated in registration. We propose a regularization term that is based on statistical knowledge of the deformations that are to be expected. A statistical model, trained on the shapes of a set of segmentations, is integrated as a penalty term in a free-form registration framework. For the evaluation of our approach, we perform inter-patient registration of MR images, which were acquired for planning of radiation therapy of cervical cancer. The manual delineations of structures such as the bladder and the clinical target volume are available. For both structures, leave-one-patient-out registration experiments were performed. The propagated atlas segmentations were compared to the manual target segmentations by Dice similarity and Hausdorff distance. Compared with registration without the use of statistical knowledge, the segmentations were significantly improved, by 0.1 in Dice similarity and by 8 mm Hausdorff distance on average for both structures.

Published as: Floris F. Berendsen, Uulke A. van der Heide, Thomas R. Langerak, Alexis N. T. J. Kotte, and Josien P. W. Pluim. „Free-form image registration regularized by a statistical shape model: application to organ segmentation in cervical MR." In: *Computer Vision and Image Understanding* 117.9 (Sept. 2013), pp. 1119 –1127

2.1 Introduction

Many medical imaging applications rely on deformable registration [21]. Deformable registration is a complicated task when it involves large and complex deformation. Registration schemes that have to deal with such deformations generally require a transformation model with a very large number of parameters. The downside of this high number of degrees of freedom is that the registration procedure can easily end up in a local minimum. For instance, the registration process may advance to a local minimum when an edge aligns to an edge of a neighboring structure. In case of complex deformations, even coarse-to-fine approaches may not suffice.

Incorporation of prior knowledge of deformation into registration is a way of regularizing the optimization procedure and reducing local minima. This is a well-known approach and it is applied in many situations; a number of these will be discussed. The prior knowledge that is used by Brock et al. [11] comprises physical properties of the tissue, whereas Pennec et al. [47] regularize with statistics on the physical strain tensors within one registration. These approaches are used for intra-subject registrations, but may not be valid for inter-subject registrations, since the transformation of one person to another is not a physical one.

Rueckert et al. [54] use statistical deformation models (SDMs) that model all the control points of the underlying B-spline deformation grid in a statistical way. A similar approach was taken by Loeckx et al. [37], who limit deformations to the space spanned by a trained deformation model. However, this means that a target deformation that deviates from the spanned space cannot be reached by registration, even if the image data provide enough evidence. To avoid this effect, Xue et al. [65] perform registration where statistical regularization is applied in a separate step. An integrated formulation is presented by Albrecht et al. [1], with the model used as a soft constraint in registration. To not restrict the target deformation space spanned by the training data, the span of the model is increased to all the degrees of freedom of the transformation, by the regularization method called shrinkage estimation. However, only registration examples of 2D images are shown.

Rueckert et al. [54], Loeckx et al. [37] and Xue et al. [65] model all displacement vectors in the image. These displacement vectors can only be obtained if a (registration) method is available that establishes correct displacement vectors for the example images. Alternatively, as is done by Albrecht et al. [1], the displacement vectors can be obtained by registering example segmentations. In the latter approach the resulting displacement vectors only carry meaningful information at the boundaries of the segmentation. All other displacement vectors in the image are based on the deformation regularization used, but will still be used for training purposes.

We chose to stay close to the idea of statistical shape models (SSMs) that model only the boundary surface displacements. Heimann and Meinzer [26]

give a thorough overview of the variety of SSMs that have been proposed. The key approach of SSMs (which has inspired the development of SDMs) is a reduction of the number of degrees of freedom of the transformation model in order to span only those modes of variation that are common in a training set. This resulting transformation model is used to generate shape instances in an optimization procedure to find the optimal mapping. Again, the restriction of the degrees of freedom can be a disadvantage in finding an accurate mapping [27].

The registration framework we propose consists of a B-spline deformation model that has many degrees of freedom, but with prior shape knowledge, in the form of an SSM, acting as a soft constraint during optimization. Compared with Albrecht et al. [1] the trained model is more compact and computationally attractive.

To assess registration quality, we will evaluate segmentations obtained from atlas-based segmentation of cervical MR images. The images are acquired for external-beam radiation therapy for cervical cancer. In this treatment the target volume should receive a high dose, while the neighboring organs should receive a dose as low as possible. For the treatment planning of cervical cancer, manual segmentations are made for the clinical target volume, bladder, rectum, sigmoid and bowel in our medical center. In our study, the manual segmentations of the clinical target volume and bladder are used for atlas-based segmentation. In atlas-based segmentation, registration is used to find the topological relation between the target image and an atlas image. This topological relation is used to propagate the segmentations of the atlas to the target.

Intra-patient registration of cervical images, using contrast-enhanced CT images [10] and MR images [58] has been reported, however inter-patient registration is not described in literature. The large and complex deformations due to differences in rectum and bladder fillings, the cervical position as well as sliding organ interfaces make these registrations a challenging task. Additionally, in our institute, the MR images for cervical cancer treatment are highly anisotropic with thick transversal slices.

Our goal is to improve inter-patient registration with application to cervical images. Registration quality is measured in terms of segmentation overlap and surface distance.

2.2 Method

The method is set up in the general context of organ segmentation in medical images. However, some choices depended on the imaging data available. Therefore the data is described first. Then, the registration framework chosen for this application and how prior shape knowledge can be incorporated is ex-

plained. Subsequently the statistical model that captures the shape knowledge is described. Finally, the expression for the actual penalty term is derived.

Data

The data used in this paper consists of 84 T2-weighted MR images of 17 patients that underwent radiation therapy for cervical cancer. Five images are available for each patient, acquired in five consecutive weeks, except for one patient who was scanned four times. In the first week a CT scan was also made for each patient.

The MR images were acquired with a 1.5 T scanner (Gyroscan NT Intera; Philips Medical Systems, Best, The Netherlands), in the transversal direction. The image dimensions are 512×512×30 voxels, with a voxel size of 0.625 mm ×0.625 mm×4.5 mm. The CT images (Aura; Philips Medical Systems, Best, The Netherlands) have a dimension of 512×512×118 voxels. The voxel size is 0.885 mm×0.885 mm×3 mm.

The manual delineations of the bladder and the clinical target volume for each MR image are used in this study. The clinical target volume is defined as the primary tumor, the uterus, the parametria, the cervix, and a part of the vagina, depending on tumor position. The delineations were made for clinical purposes by a radiation oncologist and reviewed by a radiologist. The rigid transform parameters that map MR coordinates to CT, which were used in the clinic, were provided. Figure 2.1 shows examples of the MR images.

Registration framework

In order to propagate the segmentations of a reference image I_{ref} to another subject image I_{target} a spatial transformation $\boldsymbol{T}$ that relates both images is obtained by registration. This transformation is defined as $\boldsymbol{T}(\mathbf{x}) : \Omega_{ref} \mapsto \Omega_{target}$, which maps one image domain to the other, where an image is defined as $I(\mathbf{x}) : \Omega \subset \mathbb{R}^3 \mapsto \mathbb{R}$.

For our application we initialize the transformation by alignment of bony structures in the field of view, since these provide a natural reference frame for the abdominal region in which the relevant organs are located. The global transform $\boldsymbol{T}_{\text{global}}$ captures the alignment in location, orientation and scale, as well as the coarse deformation of bony structures between the two patients. The transform $\boldsymbol{T}_{\text{local}}$ accounts for the complex organ deformations within the abdominal region. Accordingly, the transformation $\boldsymbol{T}$ is written as a (functional) composition of two parts: $\boldsymbol{T} = \boldsymbol{T}_{\text{global}} \circ \boldsymbol{T}_{\text{local}}$. The organ model we present in this paper will affect the local transform by acting as a regularizing penalty term in the registration process. The penalty term is based on a point distribution model [19] that acts on the vertices of the reference delineation and statistically describes the variations of deformations of the organ boundary in the sample population. The statistical model and the penalty term are

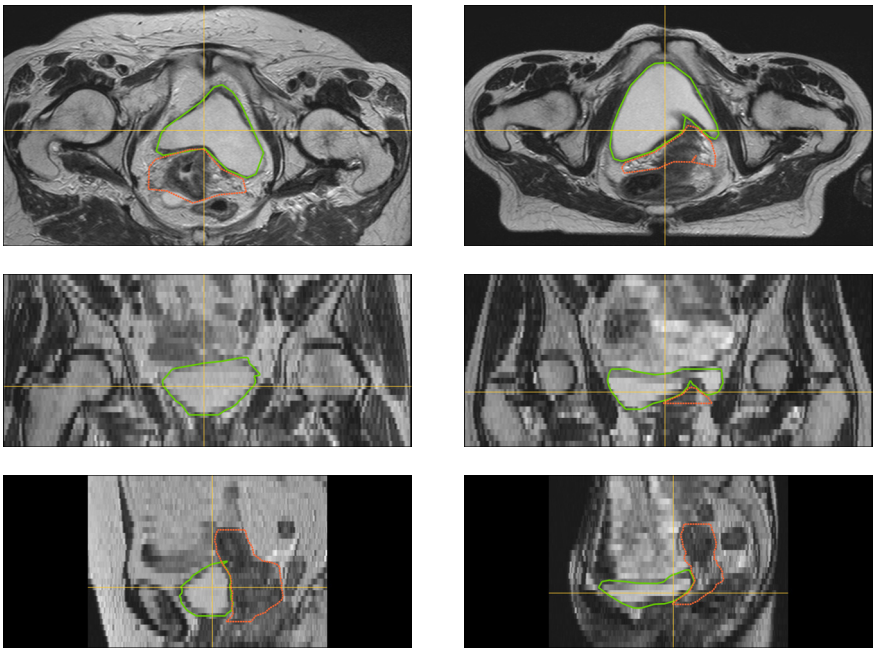

Figure 2.1: MR images of two subjects (numbered 5 and 7 respectively). The manual delineation of the bladder is shown as a solid, green contour and the delineation of the CTV is shown as a dashed orange contour. The lines denote the crossings of the viewing planes.

described in more detail in the next subsections 2.2 and 2.2. Since we use a B-spline transformation, which has no explicit inverse, we define the transformation of coordinates in the reference image as the transform $\boldsymbol{T}_{\text{local}}$. This implies that a target image is warped to the reference image during registration.

Since both CT and MR images are available for each patient, we choose to base the global transformation on bony structures visible in the CT images. For this transform inter-patient CT to CT registrations are performed, consisting of an affine and a coarse B-spline transform with a grid spacing of 80 mm. Because of the absolute nature of CT intensity values, a sum of squared differences is used as image similarity measure. The preprocessing of extracting bony structures is done by thresholding the image to select dense (i.e. intense) bone tissue and applying morphological operations to select the surrounding bone tissue that is less dense. The voxels in the CT image that are outside this selection are set to zero in Hounsfield units. This is approximately the value of the tissue that directly surrounds the bones.

Given this CT based inter-patient transformation, the global transformation

between MR images of two patients is the composition of respectively a rigid intra-patient (MR to CT), an affine+deformable inter-patient (CT to CT) and an intra-patient rigid (CT to MR) transform:

$$\boldsymbol{T}_{\text{global}} = \boldsymbol{T}^{-1\,\text{rigid}}_{\text{target}} \circ \boldsymbol{T}^{\text{affine+deformable}}_{\text{inter-patient}} \circ \boldsymbol{T}^{\text{rigid}}_{\text{ref}} \tag{2.1}$$

Because the statistical model is defined in the coordinate system of the reference image, the target I_{target} is transformed to the reference space using $\boldsymbol{T}_{\text{global}}$, resulting in I'_{target}. The residual transform between I'_{target} and I_{ref} is called $\boldsymbol{T}_{\text{local}}$. For the remainder of this paper $\boldsymbol{T}_{\text{local}}$ is denoted by $\boldsymbol{T}_{\boldsymbol{\mu}}$, with $\boldsymbol{\mu}$ the vector of B-spline parameters that defines the local transformation.

The registration procedure that is used to find this transformation can be expressed as a minimization of a cost function $\mathcal{C}$:

$$\hat{\boldsymbol{\mu}} = \arg\min_{\boldsymbol{\mu}} \mathcal{C}(\boldsymbol{T}_{\boldsymbol{\mu}}; I_{\text{ref}}, I'_{\text{target}}), \tag{2.2}$$

with the cost function in this registration scheme consisting of an image similarity measure $\mathcal{S}$ and penalty function $\mathcal{P}$:

$$\mathcal{C} = -\mathcal{S} + \alpha\mathcal{P}. \tag{2.3}$$

For the similarity measure we use normalized cross correlation, but other measures could be employed as well. The penalty term is defined by the statistical model, which is explained in the next subsections. A weighting factor α is used to balance the two terms.

Statistical model

The model is based on a point distribution model [19] that is made from a population of surface meshes. The model captures the variations in shape and position of the meshes in this population. Point correspondence of the meshes is obtained in a way similar to [23]. For this purpose, the delineation in the reference image is converted into a surface mesh. The population of meshes is obtained from the set of atlas delineations. By registering all binary images of the atlas delineations to the binary image of the reference, transformations are established that are used to generate a population of atlas meshes from the reference mesh. For this registration a kappa statistics metric is used, which maximizes overlap of binary images. The transformation model is a 3-level pyramid of B-splines with a grid spacing schedule of 40-20-10 mm.

A mesh consisting of N 3-D vertices $m = \{\mathbf{x}_i = [x_{1i}, x_{2i}, x_{3i}] : i = 1 \ldots N\}$ is represented as a concatenation of vertex coordinates:

$$\mathbf{p} = [x_{11}, x_{21}, x_{31}, \ldots, x_{1N}, x_{2N}, x_{3N}]^{\mathrm{T}}. \tag{2.4}$$

The space of all atlas shape vectors is characterized by a multivariate normal distribution defined by its mean $\bar{\mathbf{p}}_{\text{atlas}}$ and covariance Σ. The mean and the

covariance of the distribution are estimated by the arithmetic mean and the sample covariance.

When the number of vectors is smaller than the dimensionality of these vectors, this distribution is supported only in a subspace of possible shapes. To compute the penalty for an arbitrarily deformed reference shape, the inverse of the covariance matrix is needed, which does not exist in these cases. We circumvent this inversion problem by a covariance regularization technique called shrinkage estimation, as in [1]. The regularized covariance matrix Σ' is a linear combination of the estimated covariance matrix and a diagonal matrix:

$$\Sigma' = (1 - \beta)\Sigma + \beta\sigma_0^2 I, \tag{2.5}$$

with $0 < \beta \leq 1$ the shrinkage intensity and σ_0^2 an assumed variance. By this manipulation, the multivariate normal distribution will have a variance of at least $\beta\sigma_0^2$ in all directions. The matrix $\sigma_0^2 I$ is the assumed covariance matrix that defines the normal distribution if there was no knowledge available, except for the mean $\bar{\mathbf{p}}_{\text{atlas}}$. If the number of samples is large enough to have a reliable estimate, the shrinkage intensity should approach zero.

Penalty

With the estimated distribution, deformations of the reference mesh $\mathbf{p}_{\text{ref}}$ can be evaluated by means of the Mahalanobis distance, which is used as the penalty term during the registration. The penalty value can be interpreted as a measure of distance of the deformed reference mesh to the mean mesh, but with modes of deformation that are common in the population regarded as less distant.

Many free-form registration methods that rely on B-spline based deformations use gradient-based optimizers for speed and efficiency [32, 40, 54, 56]. Consequently, the derivatives of the penalty term with respect to the transform parameters are required. For the analytical derivative of the penalty function, the Mahalanobis distance to the mean of the distribution is expressed as a function of the transform parameters. First, the mean mesh vector is subtracted from the transformed mesh vector

$$\mathbf{p}(\boldsymbol{\mu}) = \boldsymbol{T}_{\boldsymbol{\mu}}(\mathbf{p}_{\text{ref}}) - \bar{\mathbf{p}}_{\text{atlas}}, \tag{2.6}$$

with $\boldsymbol{T}_{\boldsymbol{\mu}}(\mathbf{p}_{\text{ref}})$ the concatenated vector representation of the coordinates of each transformed element of the set m_{ref} that constitutes $\mathbf{p}_{\text{ref}}$. The difference mesh vector $\mathbf{p}(\boldsymbol{\mu})$ is now a function of the transform parameters. By including the difference vector, the Mahalanobis distance can be expressed as:

$$M(\boldsymbol{\mu}) = \sqrt{\mathbf{p}^{\mathrm{T}}(\boldsymbol{\mu})\Sigma'^{-1}\mathbf{p}(\boldsymbol{\mu})}. \tag{2.7}$$

The partial derivatives with respect to the transform parameters are:

$$\frac{\partial M}{\partial \boldsymbol{\mu}} = \frac{1}{M}\mathbf{p}^{\mathrm{T}}(\boldsymbol{\mu})\Sigma'^{-1}\frac{\partial \mathbf{p}(\boldsymbol{\mu})}{\partial \boldsymbol{\mu}} \tag{2.8}$$

From equation 2.6, it follows that $\frac{\partial \mathbf{p}(\boldsymbol{\mu})}{\partial \boldsymbol{\mu}} = \frac{\partial \boldsymbol{T}}{\partial \boldsymbol{\mu}}(\mathbf{p}_{\text{ref}})$, which are the derivatives of the displacements of the vertex coordinates with respect to the B-spline transform parameters. These derivatives are described in [40].

2.3 Experiments

The model as described in this paper was implemented as a module in `elastix` version 4.6 [35]. The method is evaluated on segmentation of two structures, i.e. the bladder and the clinical target volume, which were treated as independent registration experiments. Registration is restricted to a masked area of the reference image, with the mask created from the manual delineation of the structure with an extra margin of 4.5 mm (i.e. the height of one voxel) in all directions.

The quality of registration is measured both by the Dice similarity coefficient (DSC) and the Hausdorff distance of the propagated reference delineation and the manual segmentation in the target image. The Hausdorff distance *from* surface A to surface B is the maximum distance of the set of points on A to the nearest point on B. We use the Hausdorff distance *between* two surfaces, i.e the maximum of the distances of A to B and B to A. A propagated delineation of the target image is obtained by applying the local and global transform to the vertices of the reference mesh. The reference mesh, used for this purpose, is a densely triangulated version of the original reference delineation and typically consists of over 3000 vertices. This triangulated mesh can sufficiently represent non-linear deformations of the delineation.

The baseline of each experiment is a so-called conventional registration, which does not use the model. In order to keep the experiments comparable to the baseline, the parameters settings that were tuned to the conventional registrations are also used for model-enhanced registration (except for the model parameters). Both types of registration share the same global transformation based on bony anatomy.

Wilcoxon signed-rank tests for paired samples were performed on both the Dice scores and the Hausdorff distance of conventional and model-enhanced registrations. These tests show whether there is a significant difference in quality scores of both registration methods, with $p < 0.05$ regarded as significant and $p < 0.01$ as highly significant.

For all experiments, normalized cross-correlation was used as image similarity measure in local registration. The transformation was parameterized with a 3-level pyramid of cubic B-splines with a grid spacing schedule of 40-20-10 mm. The adaptive stochastic gradient descent optimizer [34] was used with 5000 randomly selected samples for 300 iterations each level.

A number of model parameters were determined empirically, solely based on experiments on the bladder. We found the resulting parameters suitable to perform an evaluation on the CTV as well.

To avoid biases in the experiments on the bladder, two sets of experiments were set up: one for parameter selection and one for evaluation. For these experiments two groups of data were made, one containing the first-week MR image of each patient and one that contained the final-week MR image of each patient. However, the atlas meshes used for training purposes were generated from all available images of all patients. In order to exclude intra-patient influence in the leave-one-out experiments, all five meshes of the target patient were excluded.

Parameter selection experiments

The parameters introduced are the weight factor α, which is the influence of the model in the cost function with respect to the image similarity and model regularization parameters σ_0^2 and β.

If no knowledge would be available about the shape variations of the bladder, we assume a normal distribution with uniform variance of $\sigma_0^2 = 1000$ around the mean shape for the position of vertices in the population. This means that each vertex has a spread with a standard deviation of about 3 centimeter. The shrinkage intensity is set to $\beta = 0.2$, which means that we rely on the training data for 80% and for 20% on the uniform variance assumption. In our experience with bladder registration, the method is not very sensitive to this parameter. By generating all references meshes with 81 vertices the shape characteristics of the bladders could be captured well enough, while the computational cost of using this number of vertices during registration is kept reasonably low.

As registration progresses, generally the model becomes less relevant. Therefore we decreased the model's influence on registration, in a step-wise fashion, by decreasing the weight α in the cost function. The registration process was divided into a pyramid of resolutions and a constant weight for each resolution was chosen.

The choice for the separate weighting values was experimentally determined as follows. Based on compactness [2] of the bladder shapes we ranked the data and chose one image of the upper and one of the lower half. The compactness measure separates sphere-like shapes from shapes that are for instance ellipsoidal or more complexly curved, by a ratio of surface area and volume. By selecting on compactness we have a good representation for the shape population. These two atlas images were registered to all other targets, where three-level weight sequences were evaluated with values ranging from 0.11 to 0.01 in steps of 0.02. There was no setting for the weights that showed an optimum for both atlases, but the weight sequence 0.09-0.07-0.05 was chosen as a compromise.

Additionally, it was investigated whether the model restricted registration too much. Possibly the optimal deformation is not captured by the model and the registration result might improve by allowing an iteration without the

penalty term. Therefore model-enhanced registrations were performed that were followed by a conventional registration at each resolution level. This registration, with weight sequence 0.09-0-0.07-0-0.05-0, was compared to a model-enhanced registration with 0.09-0.07-0.05 as weight sequence. When each registration level was extended in this way by another 300 iterations, a small improvement of 0.0048 mean DSC was obtained. Since this small improvement could also be addressed to a larger total number of iterations, experiments where both registration parts were set to 150 iterations per level were also performed. Because no significant improvement was found anymore, it was concluded that the chosen weight sequence 0.09-0.07-0.05 was not restricting the registration process and that is was a suitable choice for the evaluation experiments.

Evaluation experiments

For evaluation of the method, experiments were performed using the bladder and the CTV. The experiments with the bladder were done with MR images of the final week and experiments with the CTV were conducted using the first-week images. For each structure, the images of all 17 patients were registered to each other. For each patient selected as a target, 16 other patients served as reference images, which comprised 16×17 registrations per structure. However, as mentioned before, the shape models are trained by delineations of all weeks.

Although the aim of our method is to improve registrations between target and reference image, the behavior in a multi-atlas segmentation setting was investigated as well. A majority vote ($K \geq 8$) was calculated on the registered atlases for each target, which was evaluated by Dice similarity to the manual segmentation.

The conventional registration as was used in this paper does not have any explicit regularization. One could argue that adding any form of regularization would improve registration already. Since a bending energy penalty term is often used in the field of registration, a comparison to such a method was performed using the implementation available in [35]. The weighing parameter for the penalty term was set to 0.1, 1.0, 10 and 100 and the experiments were performed for the bladder on all image combinations.

As explained, the shrinkage intensity was set to 0.2 based on preliminary experiments on first-week images. The sensitivity to this parameter was also investigated for the last-week images. Model-enhanced registrations of the bladder were performed with the shrinkage intensity set to 0.1, 0.2, 0.5 and 1.0, respectively.

2.4 Results

The Dice overlap scores and Hausdorff surface distances for the bladder and the CTV are shown in figures 2.2+2.3 and 2.4+2.5, respectively. The overall mean

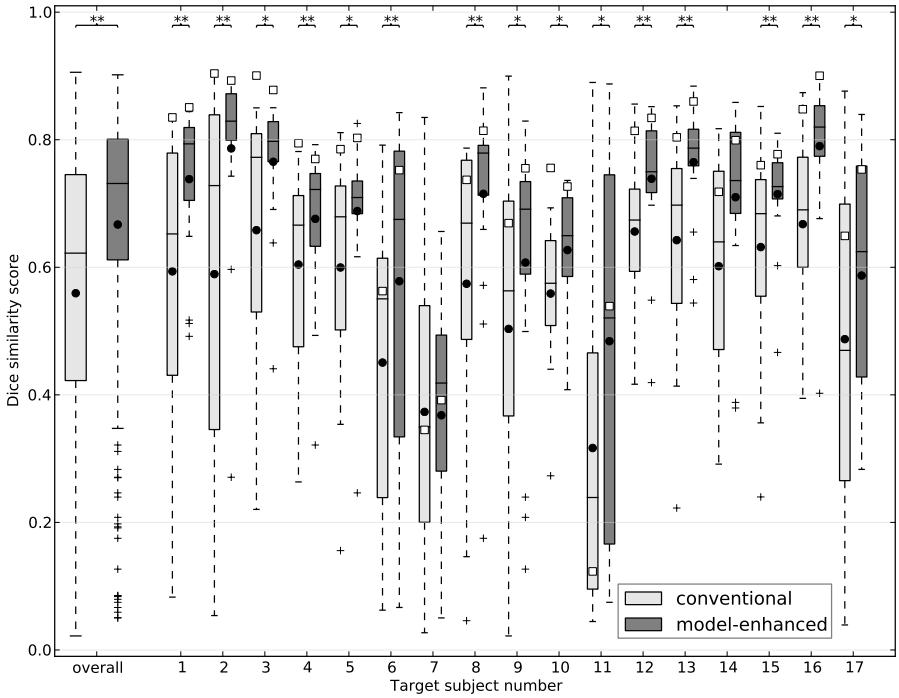

Figure 2.2: Box-and-whisker plots of bladder overlap scores for leave-one-patient-out registration experiments. The black discs (●) indicate mean values and the white squares (□) majority vote results. On the far left the results for all patients are combined. The remainder shows the results per patient. Each patient served once as a target, using the final-week image. The significance levels for $p < 0.05, 0.01$ are marked respectively *,**.

Structure	**Experiment**	**Dice**		**Hausdorff** [mm]	
		Median	**Mean**	**Median**	**Mean**
Bladder	Conventional	0.62	0.56	27	28
	Model-enhanced	0.73	0.67	17	20
CTV	Conventional	0.43	0.42	42	45
	Model-enhanced	0.57	0.55	35	36

Table 2.1: A summary of scores for all combinations of reference to target patient registrations.

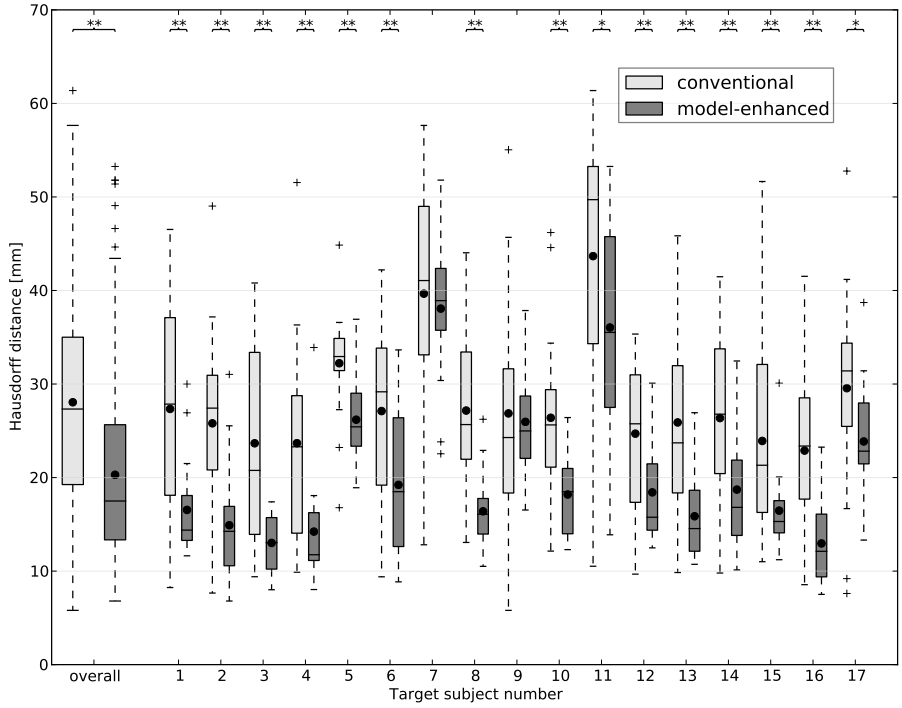

Figure 2.3: Box-and-whisker plots of bladder surface distance scores for leave-one-patient-out registration experiments. The setup and legend are explained in Fig. 2.2

and median scores of the experiments are summarized in Table 2.1. The registration results of the presented shape model were compared to conventional registration by Wilcoxon signed-rank tests for paired samples. The improvement of the average DSC for bladder and CTV are found to be significant, i.e. $p < 0.001$. The same level of significance is found for the improvement of the average Hausdorff distance for bladder and CTV.

When the results are regarded for the individual target patients as done in Fig. 2.2, the average improvement of DSC of is found to be significant for 15 out of 17 cases. For the average Hausdorff distances, as can be seen in Fig. 2.3, the improvement is also found to be significant for 15 cases. The same tests were performed on the CTV results of Figures 2.4 and 2.5, where the improvements are found to be significant for 14 and 15 out of 17 cases.

Additionally, the results are split over the individual reference patients rather than the target patients, to see whether this improvement was due to the

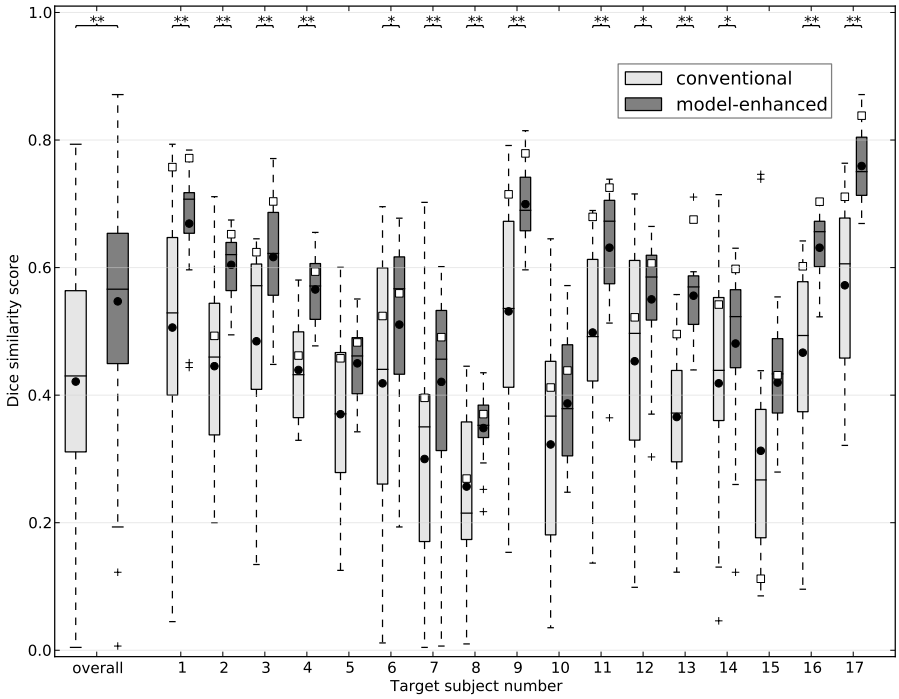

Figure 2.4: Box-and-whisker plots of CTV overlap scores for leave-one-patient-out registration experiments. The setup and legend are explained in Fig. 2.2

improvement of registration with only some of the references. For each of the four combinations of structure and score, there was a significant improvement for the majority of the 17 reference cases.

The results of the experiments with the bending penalty term and shrinkage intensity are shown in table 2.2 and table 2.3. The scores of the bending penalty experiments with weights 0.1, 1.0 and 10 are very similar to those of conventional registration, whereas with a weight of 100 the scores are a bit lower.

Registrations were done on a quad core desktop computer running at 2.83GHz. Computation times for conventional registration of one image pair is about 1 and a half minute, while with model it increases to about 3 minutes.

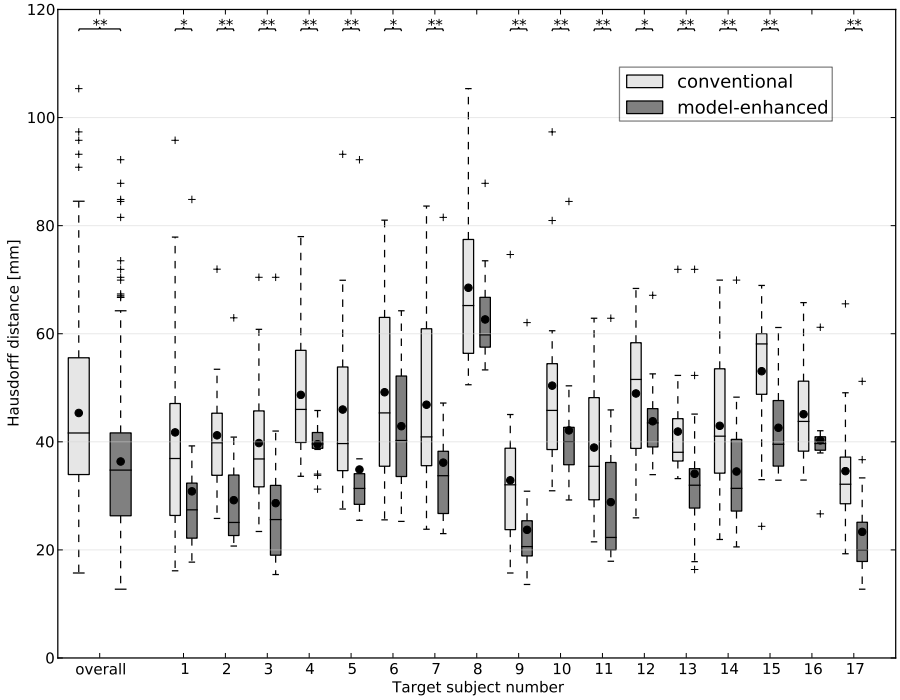

Figure 2.5: Box-and-whisker plots of CTV surface distance scores for leave-one-patient-out registration experiments. The setup and legend are explained in Fig. 2.2

Weight	**Median**	**Mean**
0.1	0.62	0.56
1.0	0.62	0.56
10	0.63	0.56
100	0.60	0.55

Table 2.2: Overall Dice overlap scores of conventional bladder registrations with various bending penalty weights.

Shrinkage intensity β	**Median**	**Mean**
0.1	0.73	0.66
0.2	0.73	0.67
0.5	0.74	0.68
1.0	0.71	0.68

Table 2.3: Overall Dice overlap scores of model-enhanced bladder registrations with various shrinkage intensities.

2.5 Discussion

The results of the experiments show that the inclusion of the statistical model significantly improves registration in terms of Dice similarity and Hausdorff distance of the segmentations. The Dice score for the bladder improves with 0.11 on average and for the CTV with 0.13 on average. Although Hausdorff distance can be regarded complementary to Dice similarity, the scores have a similar pattern of improvement. The bladder segmentation Hausdorff distances are improved by 8 mm and for the CTV segmentations by 9 mm on average.

The overall improvements are not caused by a just few large improvements, since figures 2.2 to 2.5 show that improvements are made for the majority of target patients. Additionally, for most of the targets, the box and whisker plots are shorter for model-enhanced registration than for conventional registration. The smaller spread in outcomes indicates that the impact of selecting a reference image for registration has decreased. Similarly, the evaluation per individual reference patient shows improvement for the majority of patients. However, it remains clear that some targets are more difficult to register than others and some references are more suitable as atlas than others.

Generally, the improved registrations also constitute a better segmentation that is based on majority voting. The assumption for the majority voting method is that the individual propagated segmentations have uncorrelated errors with respect to the target segmentation. By applying the voting rule on a sufficiently large number of propagated segmentations, these errors have a reduced influence on the final segmentation. The majority vote of the segmentations obtained by independent conventional registrations generally improves the Dice score over the average segmentation scores for that target. For model-enhanced registrations however, one could expect that majority voting does not give an improvement, since the individual errors with respect to the target segmentation might have become strongly correlated due to the influence of the model. Although in each model-enhanced registration the structure shapes of all other references are incorporated, the specific image characteristics of the image chosen as reference still has its influence on registration. Also the reference image defines the initial shape during registration.

The sensitivity of the registration outcome to the shrinkage intensity is low, which was also observed in preliminary experiments on first-week images. A small decrease in median Dice scores is observed for a shrinkage intensity of 1.0. As can been seen from equation 2.5, setting the shrinkage intensity σ_0^2 to 1 results in a model that still uses the mean shape, but where shape (co-)variations around this mean are penalized equally. Using this "reduced" model still gives good scores on average, but at a cost of an overall distribution that has more mass in the mid range of scores, resulting in a lower median. Modeling the shape covariance clearly has influence on the distribution of Dice scores.

Although the scores are substantially increased by inclusion of shape knowledge, the obtained median of 0.73 DSC for bladder segmentation indicates a quality of segmentation too low for clinical use. The quality is not yet at the level of *intra-patient* registration, of which a median Dice similarity coefficient of 0.81 for bladder alignment was reported [58]. However, shape and tissue variations for intra-patient are smaller than for inter-patient generally.

In registration there is a balance between the amount of flexibility that is allowed in a registration procedure and the risk of ending up with a completely failed registration. However, regularizing registration by a bending penalty term does not give a useful reduction of local minima. Bending penalty experiments, as summarized in table 2.2, generally show very similar distributions of Dice overlap compared with conventional registration.

By incorporation of the shape model the flexibility can be kept high, while resulting shapes remain consistent with the training population. Since the model incorporates a mean shape the method is not only a restricting term, but also a driving force in registration.

Figure 2.6 shows an example of conventional registration failure due to a local minimum. With conventional registration a part of the bladder edge of the reference is erroneously aligned with the edge in the target. Here normalized cross-correlation image similarity showed a local optimum for a mapping the bladder to the surrounding fat tissue of similar intensity, while aligning the edges.

A similar effect can be seen in Fig. 2.7, where the reference image is of patient 5, which is depicted in Fig. 2.1 left. This reference, which shows a bladder bulge at the right of the transversal image, defines the initialization shape for the registration process. Due to this initialization, conventional registration aligns the edges of the bulge to edges in neighboring tissue, creating an extreme bladder shape. For registration with the shape model on the other hand, deformations toward this unlikely shape are penalized.

The coronal view of Fig. 2.6 shows some spikes in the obtained delineation of both the conventional and the model-enhanced registration. As mentioned before, the manual delineations are made for clinical purposes. 2D polygons were drawn on transversal slices, which were connected automatically to form a 3D mesh. Because of this slice-based procedure, uncertainties in drawing

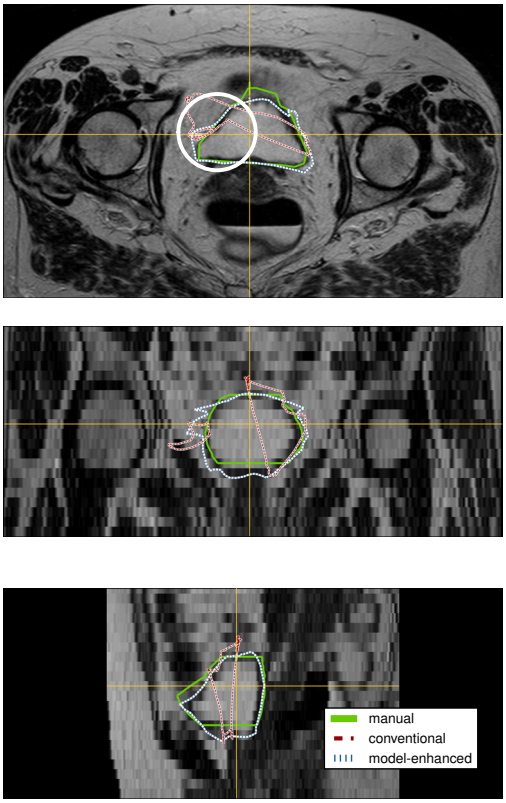

Figure 2.6: A transversal, coronal and sagittal cross section of an example registration. Registration with regularization by a shape model prevents an inside-out alignment of edges, as can be seen inside the white circle, because of the improbable shape it would result in.

delineations can produce zig-zag patterns when viewing different cross sections. Since the proposed delineations are the propagated versions of the 3D mesh, these patterns remain.

Although model-enhanced registration performs better that conventional registration for the CTV, the quality of segmentation of the CTV is considerably lower than that of the bladder, in both types of registration.

Poorer intensity-based registration of the CTV can be ascribed to multiple causes. First of all, there is a wide variation in size, shape and orientation of the uterus, which is the main part of the CTV. The uterus can have different tilted positions and has a high mobility due to bladder and rectum filling. Its sliding interface with multiple neighboring structures, i.e. the bladder, the rectum and the bowel, makes registration on surrounding tissue inconsistent.

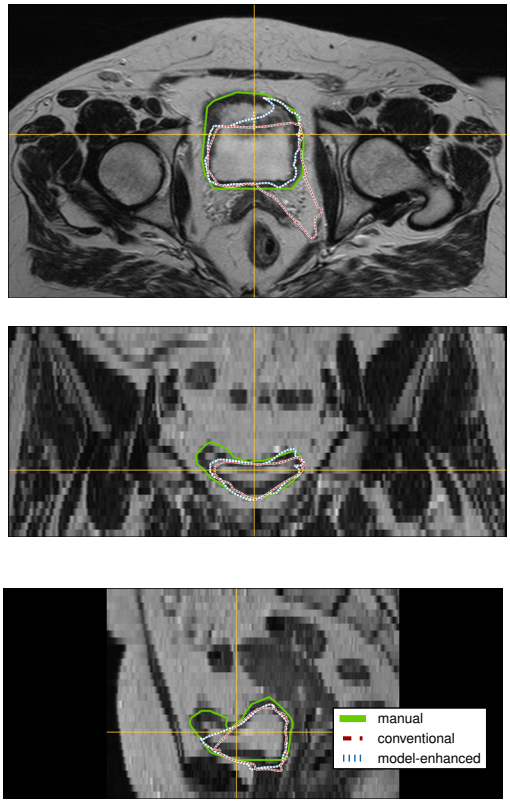

Figure 2.7: A transversal, coronal and sagittal cross section of an example registration. Conventional registration aligns the bladder to neighboring bright fat tissue.

The uterus does not have strong image features that are present in all subjects. Furthermore, the tumor size and appearance and its location within the CTV vary over patients.

Besides all the difficulties in registration, modeling the shape of the CTV is challenging too. Two causes could be identified. The point correspondence obtained by registration of binary images is prone to errors, because it does not take into account the different substructures that constitute the CTV. Secondly, due to tilting of the uterus, there are rotational parts in the transforms. But since our model treats the deformations as a linear combination of all the displacement vectors, it is less suitable to describe rotations.

2.6 Conclusion

Large and complex deformations form a great difficulty in free-form registration. In this paper we extend free-form registration by integrating prior knowledge obtained from boundary deformations of segmentations in set of atlas images. We demonstrate our approach by inter-patient registration of cervical images.

Inter-patient registration of cervical images has to deal with a number of difficulties. The topology and appearance of structures differ substantially between two persons. There is a wide variation in bladder sizes and shapes, as well in size and orientation of the uterus. Also the differences in the amount of (visceral) fat and rectum and bowel fillings between patients, make these registrations a difficult task.

To evaluate the proposed approach, we compare manual segmentations with the propagated segmentations by Dice similarity and Hausdorff distance. Registration experiments were performed on the bladder and the CTV (i.e. a combination of the uterus, the cervix, the parametria and the upper part of the vagina). We show that our proposed method significantly improves segmentations of both these structures in comparison with a conventional registration.

CHAPTER 3

Registration of Structurally Dissimilar Images in MRI-Based Brachytherapy

Abstract

A serious challenge in image registration is the accurate alignment of two images in which a certain structure is present in only one of the two. Such topological changes are problematic for conventional non-rigid registration algorithms. We propose to incorporate in a conventional free-form registration framework a geometrical penalty term that minimizes the volume of the missing structure in one image. We demonstrate our method on cervical MR images for brachytherapy. The intrapatient registration problem involves one image in which a therapy applicator is present and one in which it is not. By including the penalty term, a substantial improvement in the surface distance to the gold standard anatomical position and the residual volume of the applicator void are obtained. Registration of neighboring structures, i.e. the rectum and the bladder is generally improved as well, albeit to a lesser degree.

Published as: Floris. F. Berendsen, Alexis. N. T. J. Kotte, Astrid A. C. de Leeuw, Ina M. Jürgenliemk-Schulz, Max A. Viergever, and Josien P. W. Pluim. „Registration of structurally dissimilar images in MRI-based brachytherapy." In: *Physics in Medicine and Biology* 59.15 (Aug. 2014), pp. 4033–4045

3.1 Introduction

Anatomical correspondence is a common assumption in image registration, i.e. for every region in one image a corresponding region can be found in the other. However, for applications where a structure is present in one image and absent in another image, the assumption of a one-to-one mapping is violated. Missing structures occur for instance after brain resection or bone drillout, in laparoscopic surgery where the body is insufflated with gas, or after withdrawal of a brachy applicator.

The missing structure can affect the anatomy by either (or both) of these effects [44]:

(i) The missing structure leaves a void (filled with air or fluid);

(ii) The surrounding soft tissue fills in the missing structure.

Most approaches described in the literature address (i), with brain resection as the primary application. Since (manual) masks of the missing volume may not be present [48], [25], [49] or only probabilistic [16], [44], frameworks that perform segmentation and registration simultaneously have been proposed. Correctly retrieving the corresponding regions is crucial for reducing spurious deformation at the resection edges.

Zacharaki et al. [68] incorporate a tumor growth model to capture the displacement of surrounding brain tissue, which is a type (ii) effect. Finite element methods have been proposed for image registration during brain surgery, where instrument insertion [42] and brainshift [61] can also cause type (ii) effects. For type (ii) abdominal image registration, where deformations tend to be larger and more complex, Oktay et al. [45] proposed a method for laparoscopic surgery where the abdominal cavity is insufflated with gas to create room to maneuver instruments. Their two-step method uses a finite-element based registration to handle the insufflation of the abdominal cavity and a conventional registration to recover the complex deformations.

In this paper we focus on the type (ii) effect, with a method that deals with a missing structure and with complex deformations simultaneously. We propose a fast geometrical penalty term for a B-spline based registration framework, without the need of estimating mechanical parameters of the tissues involved. We apply our method on registration of cervical MR images. For the treatment of cervical cancer a combination of external beam radiotherapy and (internal) brachytherapy is often used. Before either treatment the patient is imaged (using MRI) for target delineation. The amount of dose that the combined therapies deliver to specific tissues can be investigated if the correspondence between the images of both therapies is found.

Work by Christensen et al. [17] on registration of follow-up CT-scans of brachytherapy interventions illustrates the challenging deformations present in cervical images. The presence or absence of a treatment applicator (see figure

3.1) can dramatically alter the geometry of the surrounding tissue. Additionally, registration is challenging because the relatively large volume occupied by the applicator in the brachytherapy image is not present in the image of external therapy. To cope with this we propose a penalty term that minimizes

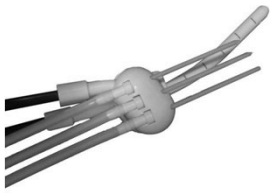

Figure 3.1: The Utrecht Interstitial CT/MR Applicator (Nucletron, Veenendaal). Together with an additional packing around the lead tubes this is the 'missing structure' in registration.

the volume of the applicator void in a general non-rigid registration framework. By including this prior knowledge, we expect better registration, because the final shape the void collapses into will be driven by the image data, subjected to minimization of the void volume. The proof of principle of this method was presented in [6].

3.2 Method

Our method exploits the fact that the missing structure can be considered to have a volume of zero after mapping it to the other image. By using a penalty term during registration that minimizes the volume of the missing structure, such a mapping can be established. The segmentation of the missing volume is assumed to be known in one image and its contour is represented by a surface mesh, denoted by U. The proposed penalty term is incorporated in the general purpose non-rigid registration framework from the `elastix` registration toolbox [35].

In the registration scheme a uniform cubic B-spline transformation model $T_\mu(\mathbf{x})$ with coefficients μ is used to map one image to the other. The so-called moving image $I_M(x) : \Omega_M$ is warped to the fixed image $I_F(x) : \Omega_F$ by interpolating the moving image at the transformed voxel coordinates of the fixed image $I'_M = I_M(T_\mu(\mathbf{x})) : \Omega_F$, with domains $\Omega_F, \Omega_M \subset \mathbb{R}^3$ in case of 3-dimensional images. The same transformation can be applied to the coordinates of the surface mesh $U' = T_\mu(U)$ which, contrarily, brings the mesh from the fixed domain to the moving domain. The cost function $\mathcal{C}$ in the registration scheme consists of an image similarity measure $\mathcal{S}$ and surface mesh penalty function $\mathcal{P}$:

$$\mathcal{C} = -\mathcal{S}(I_F, I'_M) + \alpha\mathcal{P}(U') \ , \tag{3.1}$$

where α balances the two terms. The image similarity is calculated on the images with the region of the missing structure (the applicator) masked out in the fixed image.

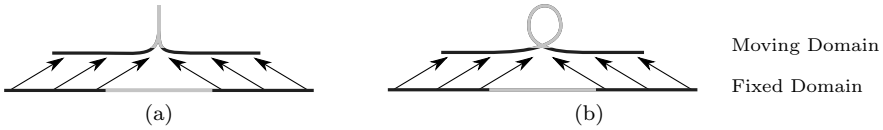

Figure 3.2: Illustration of a one-dimensional deformation field with a masked region denoted in gray. Sharp deformations in the masked region are required for full compression (a). We allow folding and negative volumes in the masked region, but impose an (absolute) minimum net volume (b).

The image that contains the volume that is minimized (the applicator) is chosen as the fixed image, since only this definition of domains allows modeling of a (surjective) mapping of a certain volume to zero by uniform B-splines. For the penalty term, we specifically choose not to force all Jacobian determinants within the applicator volume to be zero, since this would require a great number of costly evaluations. Additionally, the required sharp deformations need a very dense B-spline control grid, which is generally undesirable. Instead, we allow the region inside the applicator to fold, as is illustrated in figure 3.2. This distortion does not lead to difficulties with the image similarity metric, because the region is masked out for registration. Our missing structure penalty term (MSP) only acts on the boundary of the applicator void and is designed such that a minimal net volume is imposed. This term, where the applicator is represented by a triangulated surface mesh, allows computation of an analytical derivative (for gradient-based optimizers) and prevents self-intersection of the boundary in a computationally efficient way.

A standard technique to calculate the volume of a mesh is by subdividing the shape into tetrahedrons, which are the surface triangles connected to one central point $\mathbf{x}_c$. The oriented volume of a tetrahedron is calculated by

$$V(u) = \frac{1}{6} \det \left(\mathbf{x}_1 - \mathbf{x}_c, \mathbf{x}_2 - \mathbf{x}_c, \mathbf{x}_3 - \mathbf{x}_c \right) \quad , \tag{3.2}$$

where the ordering of the coordinates $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ of surface triangle u determines the sign of the oriented volume. By keeping the ordering of the coordinates consistent with the normal of the surface triangles, summing all oriented volumes provides the total volume (and is independent of $\mathbf{x}_c$). However, this definition fails when self intersections of the surface occur, since then parts of the shape turned inside-out will count as negative volume. To circumvent parts

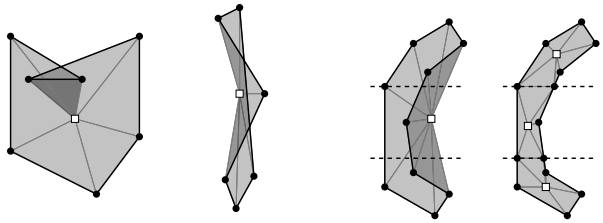

(a) Intersection (b) Subdivision

Figure 3.3: Illustrations of 2D self-intersection and of the subdivision of a non-radial convex shape. White squares are centroids; darker shades of gray indicate the surfaces (i.e. volumes in 3D) that are counted multiple times in the penalty term.

from turning inside-out during optimization, the penalty term is formulated as:

$$\mathcal{P} = \sum_{k=1}^{K} |V(u_k)| \ , \tag{3.3}$$

where K is the number of triangles in the surface mesh. The central point $\mathbf{x}_c$ is chosen as the centroid, i.e. the average coordinate of all vertices and is recalculated during optimization. The penalty term is not a true total volume, because multiple volume parts will be counted twice or more when intersection occurs, see figure 3.3. Consequently, self-intersections will be penalized, since they yield a higher cost. However, shapes that are not radially convex around $\mathbf{x}_c$ will also be penalized. By the nature of the penalty term, shapes are forced into radial convexity. This does not need to be a restriction in practice necessarily, but it prevents structures from collapsing into a curved center line, for instance. To accommodate for bending of elongated structures, a subdivision into multiple segments can be made in advance, sharing their vertices at the boundaries of the segments, as is illustrated in figure 3.3. Since the manipulation of vertex coordinates is controlled by one transformation model, the penalties can be calculated and added to the cost function per segment independently.

The gradient of the penalty term is defined by

$$\frac{\partial \mathcal{P}}{\partial \mu} = \sum_{k=1}^{K} \left[\text{sign}(V(T_\mu(u_k))) \frac{\partial V(T_\mu(u_k))}{\partial \mu} \right] \ , \tag{3.4}$$

where $T_\mu(u)$ denotes $\{T_\mu(\mathbf{x}_1), T_\mu(\mathbf{x}_2), T_\mu(\mathbf{x}_3)\}$. The partial derivative $\frac{\partial}{\partial \mu} V(T_\mu(u))$ is found by applying the chain rule on the determinant. The resulting term $\frac{\partial}{\partial \mu} T_\mu(\mathbf{x})$ is obtained analytically from the B-spline transformation model [35].

Table 3.1: Detailed properties of the data sets

Set name	No. of patients	Slice thickness fixed	Slice thickness moving	Interval
Study A	5	3 mm	3 mm	approx. 1 hour
Study B	5	4.5 mm	4.5 mm	multiple days
Evaluation	8	3 mm	4.5 mm	multiple days

In the calculation of the derivatives the dependence of the centroid on μ is ignored, since its influence is very small.

3.3 Experimental Setup

The proposed penalty term is implemented in the `elastix` registration toolbox [35] and the source code is publicly available through the `elastix` repository[1].

Datasets

Experiments are performed on 18 pairs of cervical cancer patient scans, grouped into three separate sets. Two sets, study set A and study set B, are used during development and for parameter tuning of the method. The third set is used for evaluation. Table 3.1 summarizes the details of the data sets.

For each patient there is an image pair consisting of a fixed image which is the image with the applicator in situ, i.e. before a brachytherapy fraction, and a moving image which is a scan without applicator. The moving images of study set A are acquired immediately after removal of the applicator and they are specifically made for deformation analysis. The images of study set B and the evaluation set are made during regular clinical procedures; the moving image was obtained in the week before brachytherapy during the planning of the external beam fraction.

All images are acquired with a 1.5 T MR scanner (Gyroscan NT Intera; Philips Medical Systems, Best, The Netherlands). The scans are made in transversal direction with an in-plane resolution of 0.625×0.625 mm and with varying slice thicknesses per set (see table 3.1). The dimensions of the slices are 512×512 voxels, but for study set A the images are cropped in anterior-posterior direction to remove some artifacts of signal loss.

The bladder and rectum are organs at risk since they are close to the target volume and they are subjected to deformations caused by the applicator. The gold standard for evaluation consists of manual delineations of bladder, rectum, and the surface of the void in the image without applicator, i.e. a combination

[1] `elastix.isi.uu.nl`

of the vaginal wall and the uterine inner wall. Additionally, manual delineations of the uterus are made for study set A. All gold standard delineations are drawn as contours on the transversal slices and are reviewed and approved by a radiation oncologist.

A manual delineation of the applicator is used in the proposed penalty term and as a mask. Only the main body of the applicator is delineated, leaving out the needles and the intra-uterine tube (tandem) since they are small. Fig. 3.4 shows an example image with delineations.

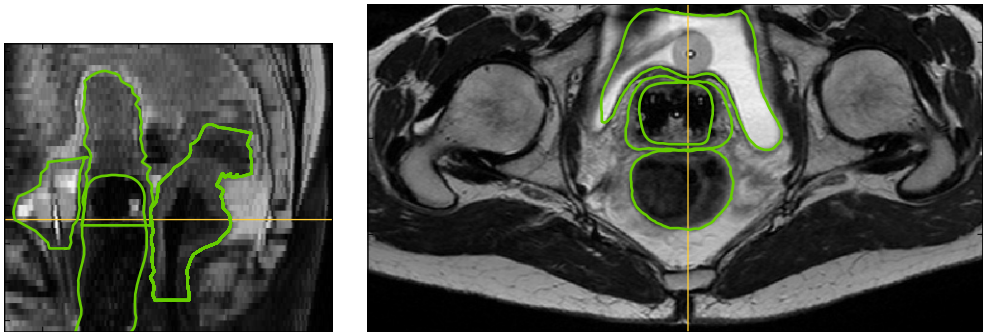

(a) With applicator

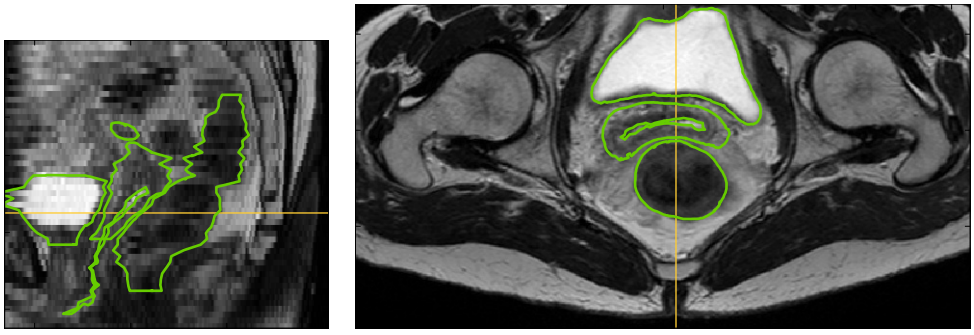

(b) Without applicator

Figure 3.4: Example image pair with manual delineations (gold standard). Delineated structures along the indicated line for top to bottom and from left to right: bladder, uterus and applicator(-void), and rectum.

Registration Experiments

Registration with the proposed penalty term is compared with registration without this term. For comparison, three types of registration are performed, ordered in increasing use of prior knowledge:

- conventional registration without applicator mask (CnM)
- conventional registration with applicator mask (CwM)
- the proposed Missing Structure Penalty (MSP)

All types of registration have the same parameter settings, which are chosen based on the two study data sets, and that are manually optimized for conventional registration. All images are T2-weighted, but because of their non-quantitative nature, we use normalized cross-correlation as a similarity measure. The registration experiments are performed in a coarse-to-fine manner with $R = 7$ resolution levels each with a factor $\sqrt{2}$ apart, approximately. Although the images of the three data sets have different slice thicknesses, one image pyramid schedule is employed that is a compromise between the amount of resolution reduction and the anisotropic nature of the data. The image scale-space resolutions are set to $\{16, 11, 8, 6, 4, 3, 2\}$ in the in-plane direction and $\{4, 3, 2, 1, 1, 1, 1\}$ in the out-of-plane direction. The B-spline grid spacing is $\{40.0, 28.0, 20.0, 14.0, 10.0, 7.0, 5.0\text{mm}\}$, isotropically. For each resolution level the optimizer is set to run for 150 iterations.

The masks that are used for experiments CwM and MSP are obtained from the applicator delineations. For experiment MSP the applicator shape is subdivided into four parts of equal height, leading to four meshes with a total number of 645 vertices on average. Physically, the applicator extends outside the field of view of the image, but obviously it could be delineated within the image domain only. To avoid an artificial shape edge at the border of the image domain, the applicator surface mesh was extended with one extra point outside the image domain. Figure 3.5 illustrates an example of an applicator mesh. The

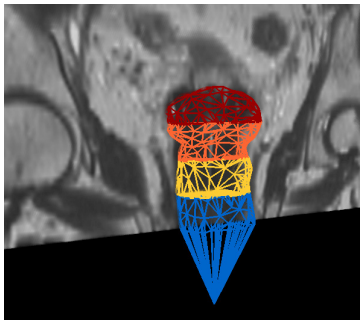

Figure 3.5: Surface mesh of an applicator split into four parts and with an extension of the lowest part outside the image domain.

weighting factor of the penalty term is $\alpha = 10^{-8}\text{mm}^{-3}$, this brings the penalty of unit mm^3 into the range of normalized cross correlation ($[-1, 1]$). This value of α was determined by tuning on study set A, initially. Subsequently when

study set B was included it was tested again by varying it with factors of two, but we found no improvement for this set.

Evaluation measures

To evaluate registration accuracy, the mapping established by each experiment is used to transform all manual delineations of the fixed image to the moving image. A transformed delineation U' is then compared with the gold standard U_{gold}. The surface distances from U_{gold} to U' are constructed by measuring the distances of the closest plane or point on U' for all points on U_{gold} [18]. Similarly, the surface distances are obtained for U' to U_{gold}. These distances are combined and they are characterized by three percentiles:

- Δ_{100}: the maximum distance (Hausdorff distance)
- Δ_{90}: the 90^{th} percentile distance
- Δ_{50}: the median distance

The rectum is the last portion of the large intestine and has no sharp definition of its start and end, intrinsically. To avoid errors due to the varying extent in head-feet direction of the delineations of the rectum, evaluation of the surface distance is done within the range of slices in which the surface meshes of all experiments were present. For the evaluation of the registered applicator, we obviously could not define an exact gold standard, because the structure is not present in the image. We have chosen to use a manual delineation of the vaginal and uterine wall for evaluation, which is the anatomical structure in which the deformed applicator should be positioned. In this case, the surface distances are calculated for U' to U_{gold} only. Similarly to the rectum, the slices of the propagated applicator below the delineated vaginal wall are excluded from evaluation.

A second measure of evaluation is the Dice similarity score, which indicates the amount of volume overlap. This is calculated for all structures, except for the applicator. Since the volume of the applicator void is near zero, calculating an overlap is infeasible, therefore the residual volume of the propagated applicator is used as a measure instead. The residual volume is obtained by voxel counting the volume of the applicator after transformation. This measure is not identical to the penalty term, but it is a true volume where negative volume due to self-intersection contributes positively. Similarly to the penalty term, this volume is expected to be very small for good registrations.

3.4 Experimental Results

The surface distance measures are summarized for each dataset by the triple $\overline{\Delta}_{100}$-$\overline{\Delta}_{90}$- $\overline{\Delta}_{50}$, i.e. the averages of the distribution characterizations. These

are reported in table 3.2 for each experiment and structure. In table 3.4 the average initial volume of the applicator and the residual volumes of the experiments are listed. The initial volume of the delineated applicator varied due to differences in field of view and the amount of packing around the lead tubes. The set averages for the Dice similarity scores are shown in table 3.3.

Although the set sizes are relatively small, we performed statistical tests for completeness. For all structures and measures, two-sided paired t-tests were performed for CnM compared with MSP and CwM compared with MSP. Significant differences ($p < 0.05$) are indicated by a star (*) in the respective columns of the conventional registration.

A visual result of the propagated delineations can be found in figure 3.6. Additionally, to provide insight in the established mapping of our proposed method, an example deformation field is shown in figure 3.7. Here, the deformations of the non-rigid part are illustrated by a deformed grid. As a reference, the contours of the structures delineated in the moving image are shown. The contours are obtained from the cross sections of the structures with the transformed grid, but with discarding out-of-plane displacements.

All registrations were done on a quad-core desktop computer running at 2.83GHz. Computation time for CnM was 6 minutes and 59 seconds and for CwM 7 minutes and 6 seconds on average. With inclusion of the penalty term the registration time increased to 8 minutes and 0 seconds on average.

3.5 Discussion and Conclusion

We have proposed a method for free-form registration of two images where a structure is present in one of them, but absent in the other. We employ a penalty term for registration that minimizes the volume of the void, as for instance left by the removal of an applicator in brachytherapy of the cervix. Our penalty term primarily affects the registration of the tissue surrounding the missing structure.

Compared with conventional registration, either using a mask or not, the use of the penalty term improves the surface distance of the propagated applicator to the gold standard anatomical position. The average Hausdorff distance, the average 90th percentile distance, and the average median distance are greatly improved by our proposed MSP method over the conventional methods CnM and CwM for all sets.

The ability of the penalty term to impose a minimum applicator void volume can best be seen by the very low residual volumes that were obtained relative to both conventional registrations. Without imposing the penalty term, the registration process could not find a mapping that has a minimum applicator void volume. The residual volume must be regarded complementary to the measure of surface distance, since a small residual volume only does not infer that the correct position of the void was obtained. Since the surface distance

Table 3.2: Average surface distance [mm] at three percentiles of the propagated delineations to the gold standard delineations.

	Experiment	**CnM** $\overline{\Delta}_{100}$-$\overline{\Delta}_{90}$-$\overline{\Delta}_{50}$	**CwM** $\overline{\Delta}_{100}$-$\overline{\Delta}_{90}$-$\overline{\Delta}_{50}$	**MSP** $\overline{\Delta}_{100}$-$\overline{\Delta}_{90}$-$\overline{\Delta}_{50}$
Applicator	Study set A	34*- 22*- 11*	27*- 17*- 7	11 - 7 - 3
	Study set B	25 - 19 - 11	28*- 21*- 11*	19 - 11 - 7
	Evaluation set	23*- 17*- 9*	25*- 17*- 9*	18 - 11 - 4
Rectum	Study set A	25 - 10 - 2	23 - 8 - 1	13 - 4 - 1
	Study set B	23 - 13 - 4	23 - 13*- 3	21 - 9 - 2
	Evaluation set	32 - 13 - 2	29 - 11 - 2	30 - 11 - 2
Bladder	Study set A	23*- 6 - 1	18 - 4 - 1	18 - 5 - 1
	Study set B	25 - 12 - 2	24 - 12 - 2	24 - 10 - 2
	Evaluation set	25 - 13 - 4	26 - 13 - 4	24 - 12 - 3
Uterus	Study set A	26 - 9 - 2	24 - 8 - 2	24 - 8 - 2

Table 3.3: Average Dice similarity of the propagated delineations to the gold standard delineations

	Experiment	**CnM**	**CwM**	**MSP**
Rectum	Study set A	0.62	0.74	0.88
	Study set B	0.52	0.60*	0.70
	Evaluation set	0.77	0.75	0.76
Bladder	Study set A	0.85	0.88	0.88
	Study set B	0.69	0.69	0.70
	Evaluation set	0.68	0.65	0.68
Uterus	Study set A	0.77	0.79	0.77

Table 3.4: Average applicator residual volume [mm^3]

Experiment	**Initial**	**CnM**	**CwM**	**MSP**
Study set A	66	28*	42*	2
Study set B	43	12*	23	0
Evaluation set	75	12*	27*	1

was improved, the smaller residual volumes are regarded as an improvement of registration.

The differences in registration are also reflected in the close neighborhood of the missing structure. The bladder, the rectum, and the uterus are delineated structures near the applicator. It must be noted that the propagated delineations for both the bladder and the rectum at the opposite side of the applicator are virtually identical for the three methods, as can be noticed in figure 3.6.However, the measures of surface distance and Dice similarity that apply to the entire structures were improved by using the MSP approach in two out of three sets.

The general differences between the outcomes of the three sets can be related to the differences in the data. In study set A, in contrast to the other sets, the image without applicator was acquired directly after treatment, at the end of the session. For most patients in set A the rectum expanded at removal of the applicator due to gas in the bowel accumulated during treatment. A gas-filled rectum is imaged as a dark region and is easily misaligned with the dark applicator, especially when the applicator is not masked for registration (see the relative large surface distances of CnM in study set A in table 3.2) In the other sets the rectum fillings consisted of less gas and were more diverse, since the images were acquired at different sessions with multiple days in between. For the same reason, more diversity in the bowel topology was observed. This has

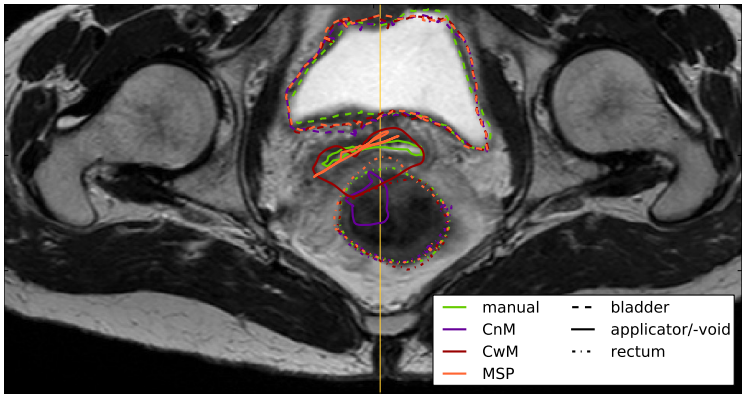

(a)

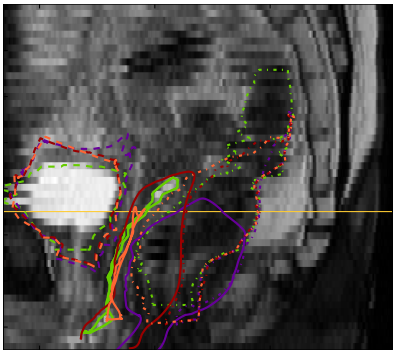

(b)

Figure 3.6: A transversal (a) and sagittal (b) view of a moving image with an overlay of the propagated delineations of the bladder, the applicator and the rectum. Green delineations denote the gold standard of the bladder, the rectum and the applicator void that is defined by the vaginal and uterine wall in which the applicator should be positioned. The propagated delineations of the experiments CnM, CwM and MSP are shown as an overlay in purple, red and orange.

an influence on the bladder shape and the registration thereof, which is reflected in poorer scores for bladder alignment in study set B and the evaluation set compared with study set A.

The amount of folding in (and smoothness of) the deformation field is often used as a quality measure of registration. Although the total volume of folded voxels outside the applicator region is not larger when adding the proposed penalty term, folding does occur. In figure 3.7 folding can be observed between the spine and the upper part of the rectum. Anatomically, the rectum pro-

gresses into the sigmoid, which is a highly movable structure that is loosely connected to the backside of the abdominal cavity and therefore discontinuous deformations can be expected. The folding at the lower part of the bladder can be attributed to the presence of a bladder catheter with a balloon in the fixed image but not in the moving image. These types of folding are typically caused by complex aspects of the data and are also present in the conventional registrations. However, the folding of the applicator region by our proposed method caused also some folding inside the uterus near the applicator. Topology changes due to sliding movement of the uterus and the bowel remain a challenge in registration.

In our applications no substantial air or fluid pockets were observed that replaced the applicator volume after removal. However, the method could allow for small pockets to exist if these are clear from the image, since our volume minimizing term is only a soft constraint. These situations, though, would require more attention to the setting of the balancing term α.

The method is presented in the context of radiotherapy, where a proper mapping between both images allows for the propagation of the associated dose images in order to estimate a cumulative (equivalent) dose for each organ tissue. First of all this dose is an estimate, because in (current) practice there is a time window between MRI scanning and dose delivery in which motion of organs can occur. Secondly, it is an estimate because registration cannot guarantee a perfect mapping. We hypothesize that a more accurate registration will lead to a better estimate of the cumulative radiation dose. We do not have the means to investigate this quantitatively in terms of dose, but we assume that the local geometry and strength of the gradients of the dose distribution play an important role. At the locations with a high dose gradient registration accuracy has the most impact on the accuracy of the cumulative dose. High dose gradients typically occur near the applicator (8-15%/mm [17]), which emphasizes the obtained improvement of the proposed method.

We have presented a penalty term to minimize the volume of a missing structure in medical image registration. In brachytherapy image registration, the penalty term substantially improves the alignment of the missing applicator structure to the gold standard anatomical position. In two out of three groups of data the alignment of neighboring structures is improved as well.

Acknowledgements

The authors are indebted to C. Nomden from the Department of Radiotherapy, University Medical Center Utrecht for providing manual delineations.

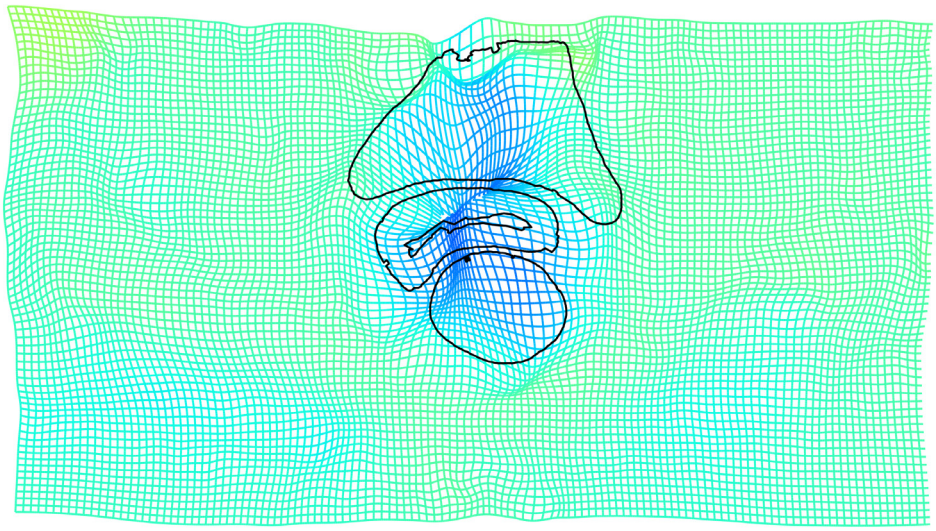

(a) Transversal plane with an original grid spacing of 2.5×2.5 mm.

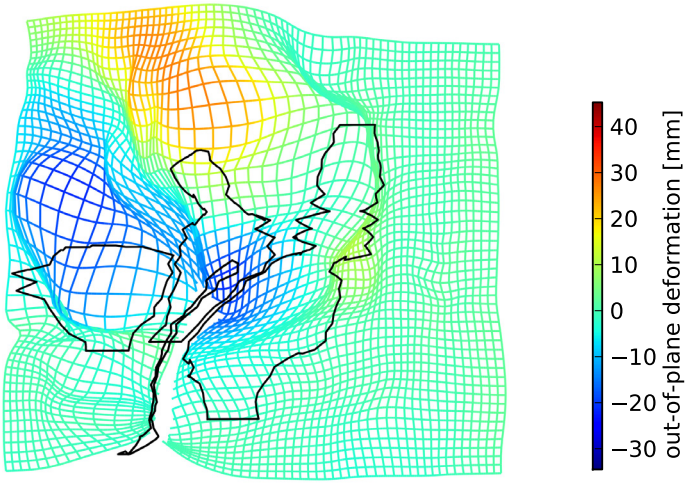

(b) Sagittal plane with an original grid spacing of 2.5×3.0 mm.

Figure 3.7: Example deformation field for MSP. A transversal (a) and a sagittal plane (b) in the fixed image were selected and the deformation field was applied to a rectangular grid. The contours of the structures as delineated in the moving image are shown for zero out-of-plane displacement.

CHAPTER 4

Registration of organs with sliding interfaces and changing topologies

Abstract

Smoothness and continuity assumptions on the deformation field in deformable image registration do not hold for applications that concern objects that slide along each other. Recent extensions to deformable image registration that accommodate for sliding motion of organs are limited to sliding motion along approximately planar surfaces or are incapable of modeling sliding that changes the topological configuration in case of multiple organs.

We propose an extension to free-form image registration that is not limited in this way. Our method uses a transformation model that consists of uniform B-spline transformations for each organ region separately, which is based on segmentation of one image. Since this model can create overlapping regions or gaps between regions, we introduce a penalty term that minimizes this undesired effect. The penalty term acts on the surfaces of the organ regions and is optimized simultaneously with the image similarity.

To evaluate our method registrations were performed on publicly available inhale-exhale CT scans for which performances of other methods are known. Target registration errors are computed on dense landmark sets that are available with these datasets. On these data our method outperforms the other methods in terms of target registration error and, where applicable, also in terms of overlap and gap volumes.

Based on: Floris F. Berendsen, Alexis N. T. J. Kotte, Max A. Viergever, and Josien P. W. Pluim. „Registration of organs with sliding interfaces and changing topologies." [submitted for review]

4.1 Introduction

Commonly, in deformable image registration, the spatial transformation from one image to the other is subjected to smoothness and continuity properties. These properties are imposed by a regularization term or are intrinsic in the transformation model and are assumed to find plausible solutions for the ill-posed optimization problem. However, these assumptions cause inaccuracies at the boundaries of regions that show sliding motion. Examples are the abdominal organs that slide along each other or the lungs that slide along the thoracic cage. An additional challenge is when the configuration of organ topology changes. This can be the case for example in the female pelvic area, where the uterus, bladder and bowel have sliding interfaces. The uterus can be in a position tilted forward, in between the bladder and the bowel or it can be in a upright position, with the bowel resting on the bladder.

Related work

The need for handling sliding motion in registration was recognized by many authors [13, 24, 33, 46, 49, 52, 64, 66] who have proposed deformation field regularizations for (non-parametric) image registration. Among the recent approaches [22, 51, 55, 63] that explicitly deal with sliding motion improvements in registration of lung images have been shown. These approaches can be categorized as B-spline-based [22, 63] and non-parametric (Demons-based) [51, 55]. Wu et al. [63] require accurate sliding region segmentations of both the fixed and the moving image. Delmon et al. [22] and Risser et al. [51] only need a segmentation in the fixed image and Schmidt-Richberg et al. [55] do not need a segmentation at all. The underlying idea of these methods is that all deformation vectors can be decomposed into a normal and two parallel directions relative to the closest surface. By allowing a discontinuity across the surface boundary in the amplitudes of deformation in the parallel directions, the regions can slide along this surface. The deformation vector in the normal direction is still forced to be smooth, which allows the surface boundary to move while keeping the regions touching.

However, in both Delmon et al. [22] and Schmidt-Richberg et al. [55], the bases of decomposition of the deformation vectors are determined in the fixed domain and is therefore fixed during registration. By fixing the bases of decomposition, sliding motion along a surface with relatively high curvature cannot be modeled adequately. In the case of a region that slides along a static region with a highly curved surface, the decomposition of the deformation of a certain point at the surface of the moving region will result in a large normal component that is incorrectly linked to the normal component of the static region. Enforcing smoothness and continuity between the normal component of the moving region and the normal component of the static region results in errors in the deformation field. As stated in their discussions, sliding along

curved boundaries cannot be handled adequately, but for registration of chest images their approximation was deemed sufficient. Although Risser et al. [50] originally modeled a moving region boundary, in later work [51] they assume the image region boundaries are fixed after a global non-sliding registration step.

Given their underlying assumptions none of the above methods can handle the complex sliding motion required for topology changes of organs. Due to the limitations of either the transformation model or the regularization terms, one region that shares its surface with a second region cannot slide into a position where it shares the surface with a third region.

Aim

We propose an extension to a B-spline registration algorithm that accounts for the sliding of interfaces of multiple regions where curved boundaries and topology changes of the regions are explicitly allowed. These properties of the method are demonstrated on synthetic images. The performance is evaluated on lung data and an example of registration of abdominal images is given.

A preliminary version of this method was presented previously [7].

4.2 Method

In our registration method we adopt a transformation model in which each organ region is transformed by an independent set of parameters. Consequently, different regions in the fixed image can be mapped to the same location in the moving image (causing overlap) or some locations in the moving image are not mapped by any region of the fixed image (causing gaps). Because it is unlikely in most medical applications that tissue disappears or is created, we introduce a penalty term that minimizes overlap and gaps during registration. The penalty term acts on the boundaries of the independent regions and only applies to the normal direction of the boundary (in the moving domain) to allow sliding motion.

More formally we define: the fixed image $I_F(x) : \Omega_F$, the moving image $I_M(x) : \Omega_M$ and the image domains $\Omega_F, \Omega_M \subset \mathbb{R}^3$. In registration the transformation model $\boldsymbol{T}_\mu(x)$ with parameter vector μ is optimized to find the match $I_F \approx \boldsymbol{T}_\mu(I_M)$. We assume that the fixed image is segmented into K sliding regions.

Transformation Model

The fixed image domain Ω_F is subdivided into K disjunct regions Ω_k, $k \in \{1 \dots K\}$, representing the anatomical structures. Each subdomain Ω_k is associated with a uniform cubic B-spline transform controlled by its own coefficients

which form a subset of μ.

$$\boldsymbol{T}_\mu(x) = \boldsymbol{T}_\mu^{\mathrm{k}}(x) \qquad \text{if} \qquad x \in \Omega_k \tag{4.1}$$

Penalty

Gaps and overlap of the object regions are penalized based on the misalignment of the region boundaries. The boundary of each region ($\partial\Omega_k$) is represented by a point set $\{\mathbf{p}_{k,i}\}$ with i the index of the point in the set. To avoid a sliding penalty due to the outer boundary of the image, the regions are covered by points except at the boundary of the image ($\mathbf{p}_{k,i} \in \partial\Omega_k \setminus \partial\Omega_F$). At these points the normals of the surface are denoted by the set $\{\hat{\mathbf{n}}_{k,i}\}$.

The boundary points are mapped to the moving domain by their designated transforms: $\mathbf{p}'_{k,i} = \boldsymbol{T}_\mu^{\mathrm{k}}(\mathbf{p}_{k,i})$. The orientation of the surface in the moving domain is calculated by the covariant transform of the surface normal vector followed by a renormalization:

$$\hat{\mathbf{n}}'_{k,i} = \frac{\hat{\mathbf{n}}_{k,i} J_x^k(\mathbf{p}_{k,i})^{-1\mathrm{T}}}{C},$$

using a positive normalization factor C such that $||\hat{\mathbf{n}}'|| = 1$. Since a B-spline transformation model is used, the spatial Jacobian $J_x^k = \frac{\partial \boldsymbol{T}_\mu^{\mathrm{k}}}{\partial x}$ is analytically available [35, 59], which provides an efficient implementation for propagating the normals.

The penalty term consists of all squared surface distances of locally opposing surfaces. At point $\mathbf{p}'_{k,i}$ of surface k, the perpendicular distance $D_{i,k}^{\perp}$ to the opposing surfaces is calculated by a weighted mean of the projection distances of all the points of the opposing surfaces:

$$D_{k,i}^{\perp} = \frac{\sum_{m,j|m\neq k} w_{k,i,m,j} \left\langle \hat{\mathbf{n}}'_{k,i}, \mathbf{p}'_{m,j} - \mathbf{p}'_{k,i} \right\rangle}{\sum_{m,j|m\neq k} w_{k,i,m,j}} \tag{4.2}$$

with $w_{k,i,m,j} = \exp\left(-\frac{||\mathbf{p}'_{m,j} - \mathbf{p}'_{k,i}||^2}{2\sigma_w^2}\right)$ a weight that gives the closest point the most influence and $\langle \cdot, \cdot \rangle$ the vector inner product to calculate the perpendicular distance of $\mathbf{p}'_{m,j}$ to the plane at $\mathbf{p}'_{k,i}$.

Figure 4.1 provides a geometric illustration of the penalty term which is formulated as:

$$\mathcal{P} = \frac{1}{2} \sum_{k,i} {D_{k,i}^{\perp}}^2. \tag{4.3}$$

The factor $\frac{1}{2}$ compensates for effectively measuring the distance between two opposing surfaces twice, i.e. once perpendicular to one surface and once perpendicular to the other.

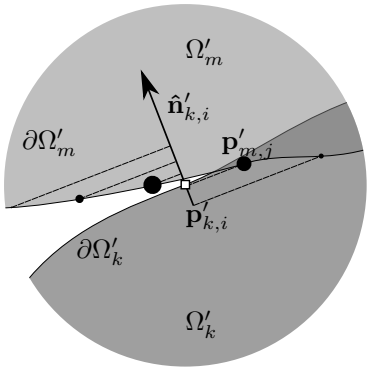

Figure 4.1: Illustration of sliding surface penalty term. The penalty contribution of one point (□) at the transformed surface of Ω'_k is considered. The points (•) of the opposing surface Ω'_m are weighted (illustrated by their radius) and projected onto the normal $\mathbf{n}'_{k,i}$ at $\mathbf{p}'_{k,i}$. The weighted average of the projections determines the penalty contribution.

The penalty term is scaled and added to the image (dis)similarity measure which then is minimized by the optimizer of the registration procedure. The optimizer is gradient-based, therefore the partial derivative of the penalty term was derived analytically. For the partial derivative the following approximation was used:

$$\frac{\partial \mathcal{P}}{\partial \mu} \approx 2 \sum_{k,i} D^{\perp}_{k,i} \frac{\sum_{m,j|m\neq k} w_{k,i,m,j} \left\langle \hat{\mathbf{n}}'_{k,i}, -\frac{\partial \boldsymbol{T}^{\mathrm{k}}_{\mu}}{\partial \mu}(\mathbf{p}_{k,i}) \right\rangle}{\sum_{m,j|m\neq k} w_{k,i,m,j}} \tag{4.4}$$

In the approximation it was assumed that the derivative of the weighted mean with respect to μ equals zero, since, although w is dependent on μ, most of its influence is reduced by the normalization which involves this weight as well. The dependence of $\hat{\mathbf{n}}'$ on μ is also assumed to be of negligible influence on the partial derivative of the penalty term compared with the main contribution via the partial derivative to the point difference $\mathbf{p}'_{m,j} - \mathbf{p}'_{k,i}$ in equation 4.2. By our definition of the transformation model in equation 4.1, the derivatives of $\mathbf{p}'_{m,j}$ and $\mathbf{p}'_{k,i}$ are never non-zero for the same parameter of μ. But, because in equation 4.3 each distance is measured from both sides once, the term $\frac{\partial \boldsymbol{T}^{\mathrm{m}}_{\mu}}{\partial \mu}(\mathbf{p}_{m,j})$, following from the chain rule, is dropped and replaced by a factor of 2. This approximation is valid if $\left\langle \hat{\mathbf{n}}'_{k,i}, -\frac{\partial \boldsymbol{T}^{\mathrm{k}}_{\mu}}{\partial \mu}(\mathbf{p}_{k,i}) \right\rangle \approx \left\langle \hat{\mathbf{n}}'_{m,j}, \frac{\partial \boldsymbol{T}^{\mathrm{m}}_{\mu}}{\partial \mu}(\mathbf{p}_{m,j}) \right\rangle$, which is true if $\hat{\mathbf{n}}'_{k,i} \approx -\hat{\mathbf{n}}'_{m,j}$. This assumption is reasonable at least at the optimum since then opposing surfaces are aligned and the normals have opposite signs.

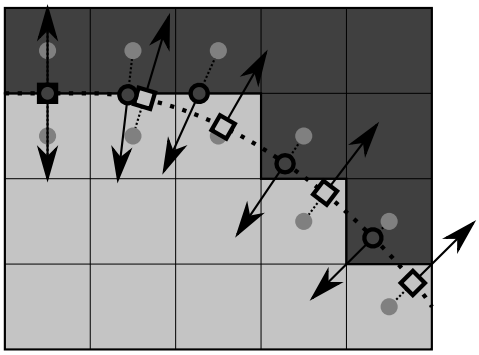

Figure 4.2: Dark gray circles and light gray squares denote the point sets of two sliding regions, each with a normal vector (black arrow). The points originate from the voxel centers (medium gray dots) and are placed at the smooth implicit surface boundary (dashed curve), but not outside the voxel

Implementation details

Surface point sets

The surface of each sliding region is represented by a point set with surface normals. The point sets with normals are constructed once in the fixed domain. At each surface a point set is created by sequentially augmenting the set with a randomly selected point on the surface, such that it exceeds a threshold distance to each other point in the set. This minimum distance parameter controls the balance between the amount of surface detail that is captured and the computational cost of having a large number of points during registration. The point sets and normals for the sliding regions are constructed from a label image. For simplicity, we start with point sets that consider the centers of the voxels at the edge of the selected label region only. However, by using voxel centers, the distance between two aligned region surfaces will be at least one times the voxel size instead of zero. Therefore, the point coordinates, initially at a voxel center, are updated to a sub-voxel location by one iteration of Newton's method such that the points lie on a smooth implicit region boundary. The implicit region boundary is obtained from Gaussian blurring the label image and the derivatives used in Newton's method are the Gaussian derivatives of the label image. This is illustrated in figure 4.2. The normal vector at the selected point equals the normalized Gaussian derivative of the label image.

Neighborhood search

The calculation of the sliding surface penalty term is based on distances between points on two sliding surfaces. In each iteration of the registration pro-

cess the distances between all the points of different surfaces are calculated. To save computation time, only the point pairs below a cut-off distance σ_c are considered in the calculation of the metric. This discards all point pairs that have a negligible contribution due to the Gaussian neighborhood kernel, i.e. the majority of point pairs. As formulated in equation 4.3 each pair of points is used twice, but in its implementation the weight and the difference vector is calculated only once for each pair.

4.3 Experiments

The transformation model and the penalty term were implemented as modules in the `elastix` [35] registration toolbox. Our method was evaluated on CT lung data that were also used by Schmidt-Richberg et al. [55], Wu el al. [63] and Delmon et al. [22]. In addition, to test sliding along curved boundaries and changing topology configurations, we applied the method to two synthetic images. As an example case, to demonstrate sliding motion with topology changes in patient data, a pair of cervical MR images was used.

Data

The CT lung data were obtained from the publicly available DIR-Lab[1] database as described in R. Castillo et al. [15] and E. Castillo et al. [14]. The database consists of 4D CT lung images of ten patients. We selected the end-exhale and end-inhale time point images as fixed and moving images. The DIR-Lab database provides expert-annotated landmarks for these images that correspond between the two time points. The spatial resolution of the data ranges from 0.97 to 1.16 mm in the transversal plane and is 2.5 mm in the orthogonal direction.

The sliding regions for the registration were defined by motion masks. A motion mask, as described by Wu el al. [63], subdivides the thorax in two sliding regions. One region contains the thoracic cage and the backbone, while the other region contains the lungs, the mediastinum and the abdomen, which move together during breathing. The motion masks for the DIR-lab data were created by Delmon et al. [22], who used the method of Vandemeulebroucke et al. [60] and they were made publicly available.

Two synthetic image pairs with motion masks were created. The first pair consists of a 3D cylindrical region that was rotated to simulate sliding along a curved boundary and is depicted in figures 4.3a and 4.3b. In the second pair three spherical shapes slide into a different topology, such that the dark regions do not share an interface anymore and that the white region creates an extra interface with the outer region. This pair is shown in figures 4.3d and 4.3e. The images of both pairs have a dimension of 256^3 voxels.

[1] `www.dir-lab.com`

The scans of the cervical cancer patient (figures 4.4a and 4.4b) were made before weekly radiotherapy treatments; one before the first, the other before the fourth treatment. The images have a high resolution in the sagittal plane directions (0.625 mm) and a low resolution in the other direction (4.5 mm). In these images sliding motion is observed between the bladder and the abdominal wall, between the bladder and the uterus and between the uterus and the rectum. The topological configuration of the organ regions changes: in the first image the bladder is surrounded by the uterus and the abdominal wall, whereas in the second image the bladder is expanded into the space between the uterus and the abdominal wall. For the registration four sliding regions were defined by a label image, which is shown in figure 4.4c. The sliding regions are the bladder, the bowel, and the cervix plus uterus. The remaining region defines the body surrounding the abdominal cavity. Since the rectum is attached to the backside of the abdominal cavity, we included this structure in the remaining region.

The label image assigns each voxel in the fixed image to one of the regions. Any visceral fat between the organs is assigned to the nearest organ, such that fully connected regions are created with relatively smooth surfaces.

Registration settings

For the thorax a 6-level resolution registration was performed with the Gaussian blur sequence $\sigma_r = \{16, 8, 4, 2, 1, 1\}$ voxels in the transversal plane directions and $\sigma_r = \{8, 4, 2, 1, 0, 0\}$ voxels in the other direction, where $\sigma = 0$ denotes no blurring. The B-spline grid spacings for the resolution levels are $\{160, 160, 80, 40, 20, 10\}$ mm in all directions. The motion masks were converted into point sets with normal vectors with the minimal threshold distance set to 10 mm. The total number of points per surface depended on its size, but consisted of 612 points on average. The neighborhood kernel width σ_w was set to 10 mm with a cut-off width σ_x of 40 mm. The optimizer was run for 2000 iterations per level.

For the synthetic experiments three levels were used with $\sigma_r = \{8, 4, 2\}$ voxels and a B-spline grid spacing of $\{64, 32, 16\}$ voxels in all directions. Both the lung data and the synthetic image registrations use normalized correlation as image similarity and a weighting factor of 10^{-6} for the sliding interface penalty.

The settings of the registration of the cervical images are similar to those for the thorax. However, the out-of-plane (i.e. left-right axis) Gaussian blur sequence is set to $\sigma_r = \{4, 2, 1, 0, 0, 0\}$ and the optimizer was run for 1000 iterations per level.

Evaluation

The registrations of CT lung data were evaluated on two criteria. The first criterion is based on the 300 expert annotated landmarks that are located

throughout the lungs in each image. Target registration error was calculated as the mean Euclidean distance between the transformed landmarks from the fixed image and their corresponding landmarks in the moving image.

The second criterion is the amount of gap and overlap between the sliding regions. Gap and overlap volumes were calculated as in Delmon et al. [22]. The region label images were converted to a 3D surfaces mesh and transformed by the transformation of each region separately. By converting the transformed meshes back into binary masks the overlap and gap volumes were determined.

A visual assessment was done for the deformation fields of the registrations of the synthetic images as well as for the cervical case.

4.4 Results

Target registration errors (TRE, Table 4.1) and gap and overlap volumes (Table 4.2) are reported together with the scores of other methods as reported by their authors (note that the scores of the method of Wu et al. [63] were taken from Delmon et al. [22]). Without registration the average TRE is 8.46 ± 5.48 mm and with our proposed method 1.36 ± 0.99 mm. With our proposed method the average gap volume is 76.5 ± 29.1 cm^3 and the average overlap volume is 37.4 ± 10.3 cm^3.

Note that the overall TRE is reported as the mean TRE of all cases with the mean of the standard deviations after the $\pm$-sign. Additionally, the TRE distribution within each case is highly skewed, making standard deviations less informative, but to conform with literature we report the mean and standard deviation.

The results of the synthetic image registrations are shown as deformation fields of the separate sliding regions in Fig. 4.3.

Registration results of the scans of the patient with cervical cancer are shown in figure 4.4.

4.5 Discussion

We developed a registration method that deals with sliding motion of organs. The sliding motion we targeted encompasses sliding of regions along curved surfaces and sliding where the topology configuration can change such that two adjacent regions become separated by a third region.

Experiments with a publicly available set of CT thorax data showed that our proposed method yields an average target registration error of 1.36 ± 0.99 mm, which compares favorably to the B-spline-based methods of Wu et al. (2.13 ± 1.82 mm) and Delmon et al. (1.47 ± 0.96 mm), and the non-parametric method of Schmidt-Richberg et al (1.66 ± 1.14 mm). These results can be found in Table 4.1. For the B-spline-based methods (including our method) gaps and overlaps of the different region can occur. The average gap and overlap volume

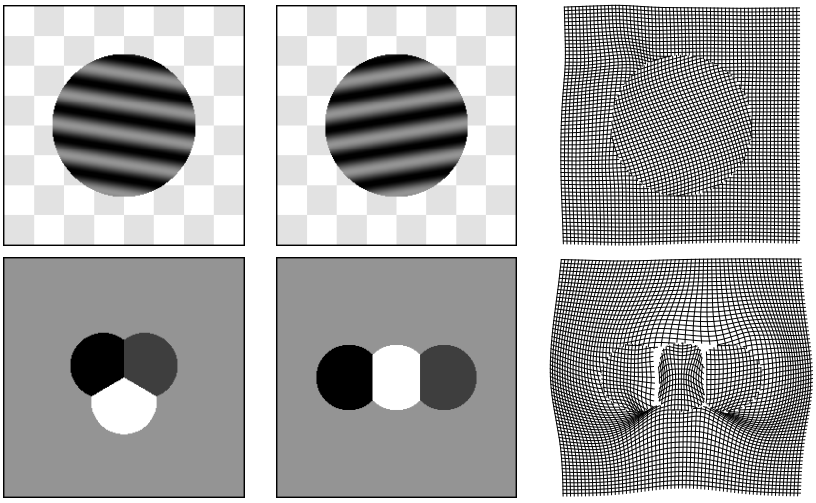

Figure 4.3: Synthetic images to test sliding along curved boundaries (top row) and sliding with changing topology (bottom row). The figures in the last column show the deformation field from the fixed to the moving image as obtained by the proposed registration procedure.

is decreased to 76.5 ± 29.1 and 37.4 ± 10.3 cm^3 by our method compared with 88.2 ± 23.7 and 52.5 ± 23.8 cm^3 of Wu et al. and 80.0 ± 23.2 and 55.1 ± 21.9 cm^3 of Delmon et al.

Additionally, comparisons were made with a conventional (non-sliding) B-spline registration, and a registration with independent regions (as in our proposed method) but without the penalty for gaps and overlaps. Both methods were performed with the same parameters as our proposed method. The conventional registration obtained a mean TRE of 2.43 mm with an average standard deviation of 2.67 mm. The difference in scores compared with the proposed method are only caused by whether sliding is modeled or not, as is illustrated in figure 4.5. From this figure it can be seen that if conventional registration is used, large landmark registration errors are made near the surface of the sliding region which are reduced in our proposed method. The landmark errors far from the surface are nearly identical for both methods, which confirms that our method only has influence at the surfaces of the regions. The deformation fields (of case 3) as depicted in figure 4.6 provides a visual comparison between the two methods. For the registration with independent regions TREs similar to the proposed method were obtained, i.e. 1.35 ± 0.99 mm on average, but with a larger gap volume (138.0 ± 77.7 cm^3) and a larger overlap volume (49.017 ± 27.020 cm^3). This indicates that our penalty term is

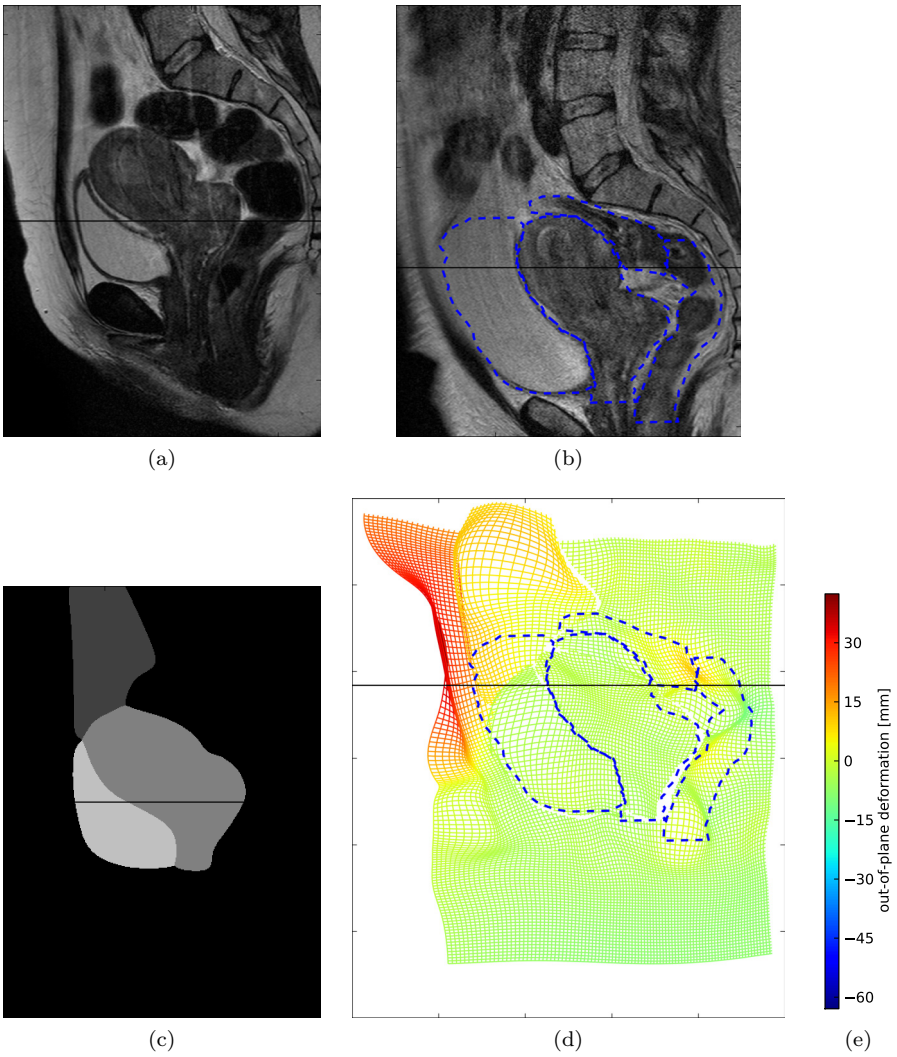

Figure 4.4: Example of cervical registration. (a) Fixed image; (b) Moving image with manual contours of bladder, uterus/cervix, rectum/sigmoid; (c) Label image defining four sliding regions, from light to dark: bladder, bowel, uterus/cervix, surrounding body + rectum/sigmoid; (d) Deformation field visualized as a grid transformed from the fixed to the moving image. The manual contours of the moving image are overlaid for reference.

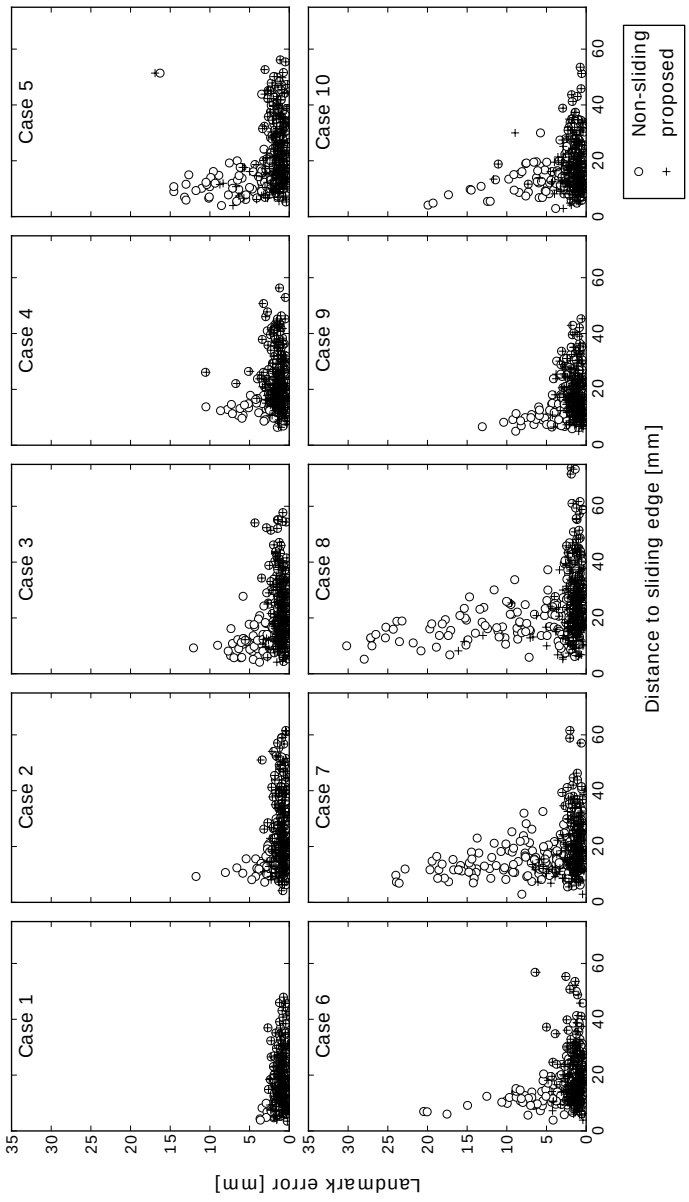

Figure 4.5: Scatter plots of landmark error against distance to the boundary of the motion mask

useful.

By using the publicly available DIR-lab data base for evaluation the different methods for registration with sliding interfaces could be compared directly. Although our B-spline-based method shares parameters such as grid spacing and image blur resolution with the method of Delmon et al., we chose different parameter settings. Delmon et al. did not use a B-spline grid spacing smaller than 32 mm, otherwise an additional bending penalty to maintain the regularity of the deformation field would have been required. Despite a B-spline grid spacing resolution schedule with a denser final spacing (10 mm), we did not find any undesired deformations such as folding or other irregularities inside the lungs. However, when our proposed method ran until a B-spline grid spacing of 32 mm, we obtained similar results as were obtained by Delmon et al. In this case the obtained scores are a mean TRE of 1.62 mm with a mean standard deviation of 1.13 mm, an overlap of 38.9 ± 14.6 cm^3 and a gap of 111.7 ± 49.6 cm^3. Since our method is not restricted to a minimal grid spacing of 32 mm, our method can improve registration by using a denser grid.

The sliding motion in the thorax largely consists of a planar movement of the lungs along the thoracic cage. To test our method for sliding along curved boundaries a registration experiment with synthetic images of figures 4.3a and 4.3b was performed. The resulting deformation field (figure 4.3c) shows smooth deformations inside both motion regions. At the interface of both regions virtually no gap or overlap is observed. Although the synthetic images have been created with a rigid transformation of the cylinder and a static background, the deformation field found contains other deformation components as well. These minor deviations from the optimum can be related to the random sampling of the optimizer, the random position of surface points, and the fact that the sinusoidal pattern of the cylinder has contrast in one direction only.

The second experiment with synthetic images simulates deformation of organs with large topology configuration changes, see figures 4.3d and 4.3e. Our proposed method was able to register the images with relatively smooth deformations within the four regions, see figure 4.3f. In the outer region a larger expansion in the upper right part than in the upper left part can be noticed. The asymmetry of this solution can be attributed to any small distortion due to random sampling and that, given the original shape of the outer boundary, this asymmetric solution provides better alignment of the outer boundary with the boundaries of the inner shapes than a symmetric one. The gaps and overlaps between the outer region and the three inner shapes are small, however, between the inner shapes some gaps have emerged.

The results of the registration of the cervical images are shown in Figure 4.4. Although the final mapping found by registration is not very accurate, it clearly shows the potential of the proposed method in a multi-organ situation. The different regions as defined by the label image can be recognized in the deformation field by the discontinuous grid lines. From the deformation field it can be seen that the abdominal wall moved outward and the bladder and

bowel slide along it. The effect of topology change can be seen by the touching surfaces of the bladder and bowel which were separated by the uterus initially. The deformation field shows that our proposed method for sliding surfaces can handle these complex deformation such that gaps and overlaps between the sliding regions are reasonably small. The inaccuracies in the registration result may be attributed to many causes, such as a poor spatial resolution in the out-of-plane direction (4.5 mm versus 0.625 mm in-plane) and some imaging artifacts. As can be noticed in the moving image, due to motion artifacts the edges of the bladder wall at the upper part are poorly imaged, which hampers a good registration. Another difficulty in the registration is noticed at the abdominal wall at the edge of the image. Due to the signal loss in this area in the moving image the registration algorithm folds back the area to a darker nearby muscle region. At the upper part of the uterus, the registration fails to find the large compression of the sigmoid and maps part of the sigmoid to the uterus which has a similar intensity. Clearly, these presented results are not sufficient for this complex example. Despite this, we think that it provides good insight in the geometric capabilities of our proposed method namely that it can handle multiple sliding regions and the topology changes between these regions.

A general limitation of our method and the other methods that partition the volume into regions that slide along each other, is that it defines all boundaries of the organs to be sliding surface which might not always be correct physically. In many applications sliding motion will only occur at part of an organ boundary. For instance in the motion mask (as defined by Vandemeulebroucke et al. [60]) the lungs are considered connected to the heart, but the heart is entirely disconnected from the thoracic cage. By introducing these artificial sliding surfaces the intrinsic regularization of the B-spline transformation model is not applied across this boundary anymore, which makes these areas more susceptible to errors in registration. These effects can be observed in figure 4.6, where the upper part of the motion mask passes through the trachea, and in figure 4.4 at the lower part of the labeled regions of the bladder and the cervix. While this vulnerability does not seem to play a big role in the current applications, this may need to be addressed in other applications.

4.6 Conclusion

A method for image registration that accommodates sliding motion between multiple regions was presented. Quantitative results were obtained from the experiments with the freely available DIR-lab lung datasets [14, 15] that have 300 manually annotated landmarks on corresponding locations throughout the lungs. On this data our method outperforms other recent methods that were developed for planar sliding motion between lungs and rib cage, in terms of mean target registration error. The total volumes of gap and overlap, which

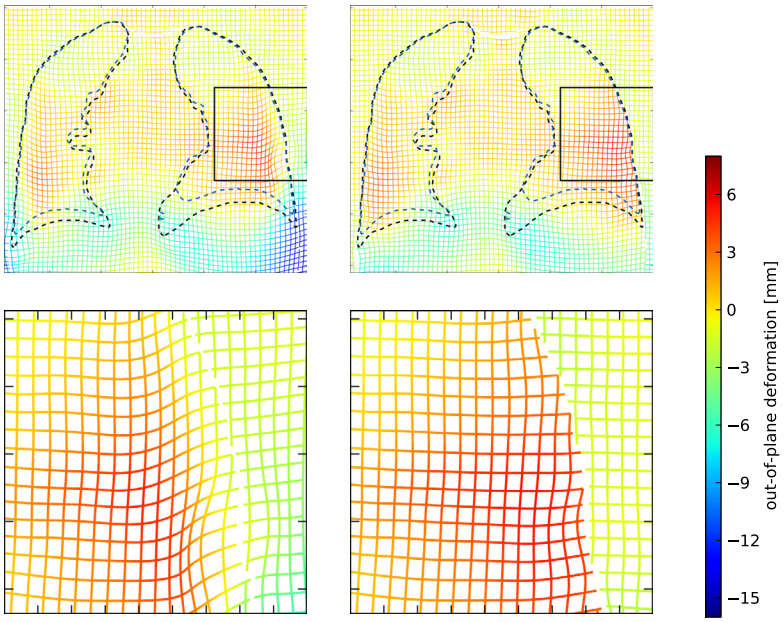

Figure 4.6: On the left the deformation found by a conventional (non-sliding) B-spline registration, on the right by the proposed method. The positions of the lungs as imaged in the inhale and in the exhale phase are overlaid as contours for reference.

are present in the B-spline based methods, are also smaller than those of the other methods. Furthermore, qualitative results from experiments on synthetic images provide evidence that our proposed method can handle sliding motion along curved boundaries and is capable of transforming regions without restricting a fixed topology configuration, i.e. in which the other methods are limited by design. This was further supported by the results of registration of cervical images in which multiple regions have sliding surfaces and where topology changes occurred.

Table 4.1: Target registration errors [mm] on publicly available DIR-lab lung datasets [14, 15]

Case	**w/o Reg**	**Schmidt-Richberg et al. [55]**	**Wu et al. [63]***	**Delmon et al. [22]**	**Proposed**
1	3.89 ± 2.78	1.22 ± 0.64	1.1 ± 0.5	1.2 ± 0.6	1.00 ± 0.52
2	4.34 ± 3.90	1.14 ± 0.65	1.0 ± 0.5	1.1 ± 0.6	1.02 ± 0.57
3	6.94 ± 4.05	1.36 ± 0.81	1.3 ± 0.7	1.6 ± 0.9	1.14 ± 0.89
4	9.83 ± 4.86	2.68 ± 2.79	1.5 ± 1.0	1.6 ± 1.1	1.46 ± 0.96
5	7.48 ± 5.51	1.57 ± 1.23	1.9 ± 1.5	2.0 ± 1.6	1.61 ± 1.48
6	10.9 ± 6.97	2.21 ± 1.66	1.6 ± 0.9	1.7 ± 1.0	1.42 ± 0.89
7	11.0 ± 7.43	3.81 ± 3.06	1.7 ± 1.1	1.9 ± 1.2	1.49 ± 1.06
8	15.0 ± 9.01	3.42 ± 4.25	1.6 ± 1.4	2.2 ± 2.3	1.62 ± 1.71
9	7.92 ± 3.98	1.83 ± 1.19	1.4 ± 0.8	1.6 ± 0.9	1.30 ± 0.76
10	7.30 ± 6.35	2.06 ± 1.92	1.6 ± 1.2	1.7 ± 1.2	1.50 ± 1.31
mean	8.46 ± 5.48	2.13 ± 1.82	1.47 ± 0.96	1.66 ± 1.14	1.36 ± 0.99

Table 4.2: Gap and overlap volumes in cm^3.

Case	Wu et al. [63]*		Delmon et al. [22]		Proposed	
	gap	overlap	gap	overlap	gap	overlap
1	38	26	39	15	23	18
2	78	46	67	60	74	34
3	99	28	83	33	57	30
4	75	34	66	44	66	28
5	110	38	78	52	61	32
6	100	86	119	77	130	50
7	105	79	108	77	119	45
8	96	91	92	93	85	53
9	61	34	54	44	70	51
10	120	63	94	56	80	43
mean	88.2 ± 23.7	52.5 ± 23.8	80.0 ± 23.2	55.1 ± 21.9	76.5 ± 29.1	37.4 ± 10.3

*As presented by Delmon et al. [22]

CHAPTER 5

Efficient normalized cross-correlation for groupwise registration

Abstract

A template-free groupwise registration method can be formulated as simultaneous registration of all image pairs in the group considered. Typically, such a groupwise metric for registration has computational complexity that scales quadratically with the number of images in the group. To save computation time, a groupwise metric was proposed with linear computational complexity that is mathematically equivalent to the sum of squared differences within each image pair. We propose a metric with linear computational complexity based on normalized cross-correlation between all pairs of images. Our metric is suited for the alignment of time series and populations that have sufficient overlap in the region of interest, since our metric is defined for the common overlap between all images only. Additionally, a variation of the metric is proposed in which the overlap of the images may vary among all pairs, but at the cost of an image normalization that might be approximate. The analytic derivative of the metric is provided for gradient-based optimization.

Based on: Floris F. Berendsen and Josien P. W. Pluim. „Efficient normalized cross-correlation for groupwise registration." [in preparation]

5.1 Introduction

Groupwise registration is the process of aligning multiple images simultaneously into a common coordinate frame. It is a template-free registration, such that it is unbiased by the choice of a specific image as reference. Its applications are, among others, the alignment of time series, mean shape template construction, and population (deformation) analysis and fusion.

Due to increased amount of image data and dimensionality of the groupwise registration problem the computational efficiency of the method, in particular the metric, is a great concern. Various similarity metrics for groupwise registration have been proposed, either general metrics or ones more specific to the type of data in an application. A groupwise information-based metric is proposed by Learned-Miller et al. [36], which minimizes entropy of the distribution of intensities. Bhatia et al. [9] use a similarity that models the average intensity within homogeneous structures. Other works, such as by Cordero-Grande et al. [20] and Mahapatra et al. [38], define a metric that focuses on the temporal characteristics in the image sequence to perform a groupwise registration of a time series. Huizinga et al. [29] propose a groupwise metric for the alignment of a series of different diffusion direction acquisitions by maximizing the explained variance of the pca-decomposed image series. A groupwise metric based on the pairwise sum of squared differences is proposed by Joshi et al. [31] and used by Metz et al. [41]. This metric calculates a voxel-wise average and minimizes the voxel-wise variance at each iteration of the registration process. By using this formulation the computational complexity of the metric is linear with respect to the number of images in the group, but is mathematically equal to the calculation of the sum of squared differences between all pairs of images [62]. Wachinger et al. [62] formulated a framework of accumulated pair-wise estimates (APE) for groupwise registration. In their work the objective function is formulated as the sum of all pairwise similarities or as the sum of squared pairwise similarities. Experiments were performed with various pairwise similarities, such as sum of squared differences, normalized cross-correlation, and mutual information. Except for the groupwise sum of squared differences, the other metrics in APE were defined to have a quadratic computational complexity. Ying et al. [67] considered the images in the group nodes in a graph with the similarity calculations between images represented by edges. To reduce computation time for a large set of images, a minimal number of edges is used for registration.

We provide a new formulation of a sum of pairwise normalized cross-correlation coefficients that has linear computational complexity with respect to the number of images, but is mathematically similar to the full pairwise implementation. Normalized cross correlation between images has been one of the classic metrics used in image registration. Up to now, this type of metric was regarded computationally expensive in a groupwise registration due to its quadratic computational complexity [20].

5.2 Transformation model

In groupwise registration the spatial correspondences between all images $I_n(x) : \mathbb{R}^d \to \mathbb{R}$ in the group $n \in \{1 \dots N\}$ are modeled by their transformations $T_n(\boldsymbol{\mu}_n; x) : \mathbb{R}^d \to \mathbb{R}^d$ from a fixed (reference) domain Ω_F to the moving image domains Ω_n. The fixed domain serves as the mean space to which all images are warped $I_n(T_n(\boldsymbol{\mu}_n; x))$. In a so-called template free registration the fixed domain is not based on an individual image; this avoids a bias to the chosen reference frame. The transform parameters $\boldsymbol{\mu}_n$ of all images are concatenated into one vector $\boldsymbol{\mu} = \left(\boldsymbol{\mu}_1^{\mathrm{T}} \dots \boldsymbol{\mu}_N^{\mathrm{T}}\right)^{\mathrm{T}}$ and are optimized over in a single (groupwise) registration process. The optimal parameters $\hat{\boldsymbol{\mu}}$ that align all images to each other into the reference domain are estimated by maximization of the image similarity metric $\mathcal{S}$: $\hat{\boldsymbol{\mu}} = \arg\max_{\boldsymbol{\mu}} \mathcal{S}(\boldsymbol{\mu})$.

For notational simplicity of some equations, the arguments $\boldsymbol{\mu}$ and x might be omitted from the transformation if it is clear from the context.

5.3 Image similarity metric

The groupwise image similarity metric we reformulate is the average of normalized cross-correlation coefficients between all image pairs in the group. The metric can be implemented directly [62], but that requires $\frac{1}{2}\left(N^2 - N\right)$ pairwise metric evaluations for a group of N images. We propose a reformulation of this metric to reduce its complexity from $O(N^2)$ to $O(N)$.

Normalized cross correlation between two images using the groupwise transformation model is given by:

$$\mathcal{S}_{\mathrm{NCC}}(I_1, I_2) = \frac{1}{|\Omega_F^{\cap}|} \frac{\sum_{x \in \Omega_F^{\cap}} (I_1(T_1(x)) - \overline{I_1})(I_2(T_2(x)) - \overline{I_2})}{\sigma_1 \sigma_2} \tag{5.1}$$

Since a (pairwise) similarity between images can be evaluated for the overlapping part only, the reference domain is $\Omega_F^{\cap} = \Omega_1' \cap \Omega_2'$, with $\Omega_n' = \left\{x \in \mathbb{R}^d \mid T_n(x) \in \Omega_n\right\}$ the domain of the warped image. The cardinality of a domain is denoted by $|\Omega|$. Note that the formulation of the metric and transformation model reverts to the common pairwise registration framework if T_2 is replaced by an identity transform. The image mean and variance are calculated each iteration based on the current $\Omega_F^{\cap}$ and are defined:

$$\overline{I_n} = \frac{1}{|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} I_n(T_n(x)) \tag{5.2}$$

$$\sigma_n^2 = \frac{1}{|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} (I_n(T_n(x)) - \overline{I_n})^2 \tag{5.3}$$

By using equations 5.2 and 5.3 to normalize the images

$$\widehat{I_n}(x) = \frac{I_n(x) - \overline{I_n}}{\sigma_n}, \tag{5.4}$$

equation 5.1 can be written as a voxel-wise product of normalized images

$$\mathcal{S}_{\mathrm{NCC}}(I_1, I_2) = \frac{1}{|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} \widehat{I_1}(T_1(x)) \widehat{I_2}(T_2(x)). \tag{5.5}$$

Using this formulation, a groupwise normalized cross-correlation of, for example, three images can be expressed as:

$$\mathcal{S}_{\mathrm{NCC}}^{\cap}(I_1, I_2, I_3) = \frac{1}{3|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} \left(\widehat{I_1}(T_1)\widehat{I_2}(T_2) + \widehat{I_1}(T_1)\widehat{I_3}(T_3) + \widehat{I_2}(T_2)\widehat{I_3}(T_3) \right). \tag{5.6}$$

In this case, the reference domain is defined as $\Omega_F^{\cap} = \Omega_1' \cap \Omega_2' \cap \Omega_3'$, which imposes some assumptions and therefore the symbol '$\cap$' is added to the metric. Namely, this definition of the domain restricts the evaluation of the similarity between two images of the group to the domain where all warped images are present, even if the individual pair has a larger overlapping domain. This could be a limiting aspect of our approach, however, for many applications where the same region of interest is present in all images the overlap between the images will be large enough to perform a good registration. In section 5.5 we address the situation where varying amounts of pairwise overlap between images are considered.

The computational complexity of equation 5.6 is $O(N^2)$, but it can be reduced to $O(N)$ by using the substitution rule

$$a \cdot b + a \cdot c + b \cdot c = \frac{(a+b+c)^2 - a^2 - b^2 - c^2}{2}, \tag{5.7}$$

resulting in

$$\mathcal{S}_{\mathrm{NCC}}^{\cap}(I_1, I_2, I_3) = \frac{1}{3|\Omega_F^{\cap}|} \frac{1}{2} \sum_{x \in \Omega_F^{\cap}} \Bigg(\left(\widehat{I_1}(T_1) + \widehat{I_2}(T_2) + \widehat{I_3}(T_3) \right)^2 - \widehat{I_1}(T_1)^2 - \widehat{I_2}(T_2)^2 - \widehat{I_3}(T_3)^2 \Bigg) \tag{5.8}$$

Generalizing from three to N images results in

$$\mathcal{S}_{\mathrm{NCC}}^{\cap}(I_1, \ldots, I_N) = \frac{1}{2N|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} \left(\left(\sum_{n=1}^{N} \widehat{I_n}(T_n(x)) \right)^2 - \sum_{m=1}^{N} \widehat{I_m}(T_m(x))^2 \right), \tag{5.9}$$

with

$$\Omega_F^{\cap} = \bigcap_{n=1}^{N} \Omega_n'. \tag{5.10}$$

Equation 5.9 can be simplified into

$$\mathcal{S}_{\mathrm{NCC}}^{\cap}(I_1, \ldots, I_N) = \frac{1}{2N|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} \left(\sum_{n=1}^{N} \widehat{I_n}(T_n(x)) \right)^2 - \frac{1}{2}, \tag{5.11}$$

by using equation 5.4 which defines

$$\frac{1}{|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} \widehat{I_n}(T_n(x))^2 = 1. \tag{5.12}$$

5.4 Metric Derivative

Many optimizers in medical image registration rely on gradient information of the optimization space. The partial derivative of the proposed metric with respect to the transform parameters $\boldsymbol{\mu}$ is

$$\frac{\partial \mathcal{S}_{\mathrm{NCC}}^{\cap}(I_1, \ldots, I_N)}{\partial \boldsymbol{\mu}} = \frac{1}{N|\Omega_F^{\cap}|} \sum_{x \in \Omega_F^{\cap}} \sum_{n=1}^{N} \frac{\partial \widehat{I_n}(T_n(x))}{\partial \boldsymbol{\mu}} \left(C(x) - \widehat{I_n}(T_n(x)) \right) \tag{5.13}$$

with, by using the chain rule subsequently on equations 5.2 and 5.3,

$$\begin{aligned} \frac{\partial \widehat{I_m}(T_m(x))}{\partial \boldsymbol{\mu}} = {} & \frac{1}{\sigma_m} \left(\frac{\partial I_m}{\partial \boldsymbol{\mu}} - \frac{\partial \overline{I_m}}{\partial \boldsymbol{\mu}} \right) \\ & + \left(I_m - \overline{I_m} \right) \frac{-1}{\sigma_m^3 |\Omega|} \sum_{x \in \Omega_F^{\cap}} \left(I_m - \overline{I_m} \right) \left(\frac{\partial I_m}{\partial \boldsymbol{\mu}} - \frac{\partial \overline{I_m}}{\partial \boldsymbol{\mu}} \right), \end{aligned} \tag{5.14}$$

and with

$$C(x) = \sum_{n=1}^{N} \widehat{I_n}(T_n(x)). \tag{5.15}$$

For a time-efficient computation of the derivative of the metric, the computation of equation 5.14 can be performed once for all images prior to the computation of equation 5.13 and the computed $C(x)$ can be reused in equation 5.11.

5.5 Partial overlaps

In section 5.3 the groupwise normalized cross correlation metric was defined only for those parts of the images that have overlap with all other images in the group, i.e. the domain $\Omega_F^{\cap}$. This definition could be a restriction for situations where not all images in the group cover the entire region of interest, or when one image in the group is severely misregistered, since then the effective domain for calculating the similarity can become very small. To handle these situations, we propose an alternative formulation (denotes by $\mathcal{S}_{\mathrm{NCC}}^{\cup}$) that has the same complexity of $O(N)$, but has different assumptions.

Since all pairs can have different amounts of overlap, the groupwise metric is defined as a weighted average of the pairwise normalized cross correlation by ratio of the pairwise overlap. Let $\Omega'_{n,m} = \Omega'_n \cap \Omega'_m$ be the domain of the overlap of the warped images I_n and I_m, then the groupwise metric is defined as:

$$\mathcal{S}_{\mathrm{NCC}}(I_1, \ldots, I_N) = \frac{1}{K} \sum_{n=1}^{N} \sum_{m=n+1}^{N} \left|\Omega'_{n,m}\right| \mathcal{S}_{\mathrm{NCC}}(I_n, I_m), \tag{5.16}$$

with $K = \sum_{n=1}^{N} \sum_{m=n+1}^{N} \left|\Omega'_{n,m}\right|$, which equals the total amount of sampled voxels in the metric.

Among the pairwise evaluations of the normalized cross-correlation the normalization of one image might vary depending on the amount of overlap it has with the other image in the pair. Instead of restricting all normalizations to a single reference domain based on intersections $\Omega_F^{\cap}$, we define the reference domain as the union of all pairwise intersections, i.e.

$$\Omega_F^{\cup} = \bigcup_{n=1}^{N} \bigcup_{m=n+1}^{N} \Omega'_{n,m}, \tag{5.17}$$

and define the normalization of each image (as in equations 5.2 and 5.3) within the domain $\Omega_F^{\cup} \cap \Omega'_n$. In other words, an image is normalized on the largest domain providing overlap with at least one other image. By the assumption of a single groupwise normalization for each image, equation 5.16 can be written as

$$\mathcal{S}_{\mathrm{NCC}}^{\cup}(I_1, \ldots, I_N) = \frac{1}{K} \sum_{n=1}^{N} \sum_{m=n+1}^{N} \sum_{x \in \Omega'_{n,m}} \widehat{I_n}(T_n(x)) \widehat{I_m}(T_m(x)). \tag{5.18}$$

Manipulating this equation, by extending the domain of summation $\Omega'_{n,m}$ to $\Omega_F^{\cup}$ while simultaneously defining the normalized images outside their original domain by

$$\widehat{I_n}(T_n(x)) = 0 \quad \text{if} \quad T_n(x) \notin \Omega_n, \tag{5.19}$$

results in an equivalent formulation

$$\mathcal{S}_{\mathrm{NCC}}^{\cup}(I_1,\ldots,I_N) = \frac{1}{K}\sum_{n=1}^{N}\sum_{m=n+1}^{N}\sum_{x\in\Omega_F^{\cup}} \widehat{I_n}(T_n(x))\widehat{I_m}(T_m(x)), \tag{5.20}$$

since adding zeros does not change a sum. From this equation the same substitution of equation 5.7 can be made and with a similar step as in equation 5.12 the metric becomes:

$$\mathcal{S}_{\mathrm{NCC}}^{\cup}(I_1,\ldots,I_N) = \frac{1}{K}\frac{1}{2}\sum_{x\in\Omega_F^{\cup}}\left(\sum_{n=1}^{N}(\widehat{I_n}(T_n(x)))\right)^2 - \frac{1}{2}, \tag{5.21}$$

The implicit assumption in this metric is that the normalization of each of the images can be done on the set of samples of a larger reference domain $\Omega_F^{\cup}$, instead of performing multiple normalizations of an image based on the overlap with each other image. This assumes that the images are stationary such that the mean and variance do not change with a varying domain. Small discrepancies in the estimation of the mean and variance will generally lead to a less steep maximum of the metric compared to a metric that performs pairwise image normalization. For larger discrepancies the intended optimum might not be the global optimum anymore.

5.6 Discussion and Conclusion

In this chapter we propose a new formulation of the normalized cross correlation for a groupwise registration framework. Compared to the implementations of this metric currently known from literature, our formulation reduces the computational complexity from $O(N^2)$ to $O(N)$, with N the number of images in the group.

The only concession that was made for this reduction is that the similarity between all image pairs is calculated within the largest space in which *all* images are present. This excludes the metric for the use in image mosaicing [43], where no common overlap between all images is present. But, for many other applications, such as temporal alignment or population alignment, the region of interest is likely to be present in all images and this provides a sufficiently large common space. Additionally, an alternative implementation of the metric was proposed that relaxes the constraint of groupwise image overlap, but assumes the images to be stationary.

The effective speedup of the proposed implementation on the total registration time will depend on the number of images used. Since the calculation of the metric is only a part of the registration algorithm, the speedup of the total registration time will scale less (Amdahl's law). The computation time

of other parts in the registration algorithm, such as loading the images, sampling the intensities, managing the transform, calculating the update in the optimizer, should scale linearly with the number of images, but the efficiency of their implementation makes it hard to compare different groupwise registration algorithms. Memory can be a limitation as well, since all images must be loaded simultaneously and 3D or high resolution 2D images can consume a lot of space.

CHAPTER 6

General Summary and Discussion

Image registration is an important task in medical image processing. Among its applications are inter-patient registration to perform segmentation of organs, registration of follow-up scans to propagate the in tissue accumulated radiation dose of a radiotherapy, and registration to perform deformation analysis over time or within a population. Generally, depending on the type of images and the nature of the spatial variation between the images, application specific registration settings need to be chosen, such as the image similarity metric, the type of optimizer and the number of degrees of freedom for the transformation model. For various sorts of applications the options available in a general registration algorithm are limited to obtain good registration results since they do not exploit application specific geometry knowledge. Specific applications can benefit from prior shape knowledge of a population, geometric properties of structures in the image, or knowledge about discontinuities in the deformation field. In this thesis various extensions to a general registration algorithm are proposed to tailor the algorithm to the issues in the applications involved.

Chapter 2 Deformable registration is prone to errors when it involves large and complex deformations, since the procedure can easily end up in a local minimum. To reduce the number of local minima, and thus the risk of misalignment, regularization terms based on prior knowledge can be incorporated in registration. We propose a regularization term that is based on statistical knowledge of the deformations that are to be expected. A statistical model, trained on the shapes of a set of segmentations, is integrated as a penalty term in a free-form registration framework. For the evaluation of our approach, we perform inter-patient registration of MR images, which were acquired for planning of radiation therapy of cervical cancer. The manual delineations of structures such as the bladder and the clinical target volume are available. For both structures, leave-one-patient-out registration experiments were performed. The propagated atlas segmentations were compared to the manual target segmentations by Dice similarity and Hausdorff distance. Compared with registration without the use of statistical knowledge, the segmentations were significantly improved, by 0.1 in Dice similarity and by 8 mm Hausdorff distance on average for both structures.

Chapter 3 A serious challenge in image registration is the accurate alignment of two images in which a certain structure is present in only one of the two. Such topological changes are problematic for conventional non-rigid registration algorithms. We propose to incorporate in a conventional free-form registration framework a geometrical penalty term that minimizes the volume of the missing structure in one image. We demonstrate our method on cervical MR images for brachytherapy. The intrapatient registration problem involves one image in which a therapy applicator is present and one in which it is not. By including the penalty term, a substantial improvement in the surface distance to the gold standard anatomical position and the residual volume of the applicator void are obtained. Registration of neighboring structures, i.e. the

rectum and the bladder is generally improved as well, albeit to a lesser degree.

Chapter 4 Smoothness and continuity assumptions on the deformation field in deformable image registration do not hold for applications that concern objects that slide along each other. Recent extensions to deformable image registration that accommodate for sliding motion of organs are limited to sliding motion along approximately planar surfaces or are incapable of modelling sliding that changes the topological configuration in case of multiple organs. We propose an extension to free-form image registration that is not limited in this way. Our method uses a transformation model that consists of uniform B-spline transformations for each organ region separately, which is based on segmentation of one image. Since this model can create overlapping regions or gaps between regions, we introduce a penalty term that minimizes this undesired effect. The penalty term acts on the surfaces of the organ regions and is optimized simultaneously with the image similarity. To evaluate our method registrations were performed on publicly available inhale-exhale CT scans for which performances of other methods are known. Target registration errors are computed on dense landmark sets that are available with these datasets. On these data our method outperforms the other methods in terms of target registration error and, where applicable, also in terms of overlap and gap volumes.

Chapter 5 A template-free groupwise registration method can be formulated as simultaneous registration of all image pairs in the group considered. Typically, such a groupwise metric for registration has computational complexity that scales quadratically with the number of images in the group. To save computation time, a groupwise metric was proposed with linear computational complexity that is mathematically equivalent to the sum of squared differences within each image pair. We propose a metric with linear computational complexity based on normalized cross-correlation between all pairs of images. Our metric is suited for the alignment of time series and populations that have sufficient overlap in the region of interest, since our metric is defined for the common overlap between all images only. Additionally, a variation of the metric is proposed in which the overlap of the images may vary among all pairs, but at the cost of an image normalization that might be approximate. The analytic derivative of the metric is provided for gradient-based optimization.

Image registration has many applications and it forms often the basis for further processing or analysis. The chapters of this thesis present various registration methods each extended to a specific application. While this thesis focuses on the registration methodology with a broader applicability in mind, little effort was put in other (additional) image processing techniques to optimize the performance to the underlying clinical goal specifically. Each chapter in this thesis independently proposes an extension to the registration algorithm that addresses a specific issue in registration. However, in various applications of registration multiple issues can play a role simultaneously to a larger or

smaller degree. By targeting a specific aspect of the application only, other aspects that are ignored may limit registration quality. Since the proposed methods are formulated as a penalty term that adds to the cost function (Chapters 2 to 4) and/or as a variation to conventional methods (Chapters 4 and 5), the next step could be to combine multiple methods in a single registration process.

In Chapter 2 the registration of an organ to the target image is guided by a shape model. To avoid hinder of the conventional transformation model, which is smooth and continuous across organ boundaries, for the organ to slide along its surrounding organs, the experiments in this chapter are confined to the organ regions by a mask. However, by excluding the surrounding tissue in registration, no structural context is provided for the alignment of the organ of interest, which may limit registration quality. A registration setup that combines the statistical shape model with the transformation model and sliding surface penalty term, proposed in Chapter 4, might address this limitation. In such a registration process the organ of interest can deform into a shape that is plausible given its training and given that there is minimal overlap with surrounding tissue which is simultaneously aligned as well. Additionally, shape models for multiple structures (possibly trained independently to avoid overtraining) might even be used in a single registration process including the sliding interface method.

In registration of cervical images involving a missing applicator structure (Chapter 3) the sliding of some surfaces was considered a source of error as well. Combining the method for a missing structure with that for sliding organs in a single registration might be a solution. Although the methods may be conceptually similar to each other, i.e. both try to minimize the space between surfaces, they have a distinct difference in implementation. The missing structure penalty acts on the points within a single structure, i.e. the applicator, whereas the sliding surface penalty acts between points of different structures exclusively. In the MSP the relation between surface points is fixed via the centroid point to attract opposing surfaces, whereas the sliding surfaces penalty uses only nearby points and the (transformed) normal to define the direction of attraction. Both methods can be used in registration simultaneously by adding their penalty terms to the cost function, however, generalizing both concepts into a single geometric formulation would be an interesting direction of research.

Similar to Chapter 2, groupwise registration of Chapter 5 could be used to segment a target image. Whereas in Chapter 2 segmentation is performed by registration of a single atlas that is guided by a shape model summarizing the variation of shapes of all other atlasses, a multi-atlas segmentation can be performed by a groupwise registration including all atlas images and the target image. The idea is that by simultaneously registering all images of a group into a common space, the spatial variation of anatomy that is present at the start of the process might gradually converge to a single solution and guides the target image to that same solution. Fusion of the propagated segmentations might be

performed to discard those atlasses that did not converge properly. Additionally, convergence to a single solution might be enforced by including a measure of segmentation overlap of the atlasses in the global cost function. How this will affect the alignment of the atlasses to the target image, which does not have a segmentation in advance, is a topic for future research. Compared with registration using a shape model, the disadvantage of groupwise registration for this purpose is its relatively large computational cost and memory requirements. Registration using a shape model is typically more computationally intensive than a pairwise registration but in the same order of magitude. Although the training of the shape model is computationally intensive, it is done in advance and it involves only pairwise registrations.

Due to the current formulation of groupwise registration it is not possible to include the proposed methods of the other chapters directly. These proposed methods define geometrical objects in coordinates of the fixed domain of registration. In groupwise registration however, geometry in the fixed domain is unknown initially, since this is where the spatial average of the images is constructed during the registration process.

Combining multiple of the proposed penalty terms with an image similarity term might require the effort of tuning of all weighting constants precisely. The increasing number of parameters that need to be set for each method that is incorporated might be a practical concern as well. Since the proposed terms alter the optimization space in different ways the interactions of the various methods and their effect on the optimization procedure may need to be investigated.

All the proposed methods are implemented in the publicly available toolbox for registration `elastix` [35], which makes it possible for other researchers to use and extend the methods for their own applications. The author hopes that, by broadening the range of applications of registration, it becomes an even more versatile instrument to use in clinical research and eventually clinical practice.

Bibliography

[1] T. Albrecht, M. Luthi, and T. Vetter. „A statistical deformation prior for non-rigid image and shape registration." In: *IEEE Conference on Computer Vision and Pattern Recognition* (2008), pp. 1–8 (cited on pages 9, 10, 14).

[2] A. Alyassin, J. Lancaster, J. Downs III, and P. Fox. „Evaluation of new algorithms for the interactive measurement of surface area and volume." In: *Medical Physics* 21.6 (1994), pp. 741–752 (cited on page 16).

[3] F. F. Berendsen and J. P. W. Pluim. „Efficient normalized cross-correlation for groupwise registration." [in preparation] (cited on page 62).

[4] F. F. Berendsen, A. N. T. J. Kotte, M. A. Viergever, and J. P. W. Pluim. „Registration of organs with sliding interfaces and changing topologies." [submitted for review] (cited on page 44).

[5] F. F. Berendsen, U. A. van der Heide, T. R. Langerak, A. N. T. J. Kotte, and J. P. W. Pluim. „Free-form image registration regularized by a statistical shape model: application to organ segmentation in cervical MR." In: *Computer Vision and Image Understanding* 117.9 (2013), pp. 1119 –1127 (cited on page 8).

[6] F. F. Berendsen, A. N. T. J. Kotte, A. A. C. Leeuw, M. A. Viergever, and J. P. W. Pluim. „Free-Form Registration Involving Disappearing Structures: Application to Brachytherapy MRI." In: *Abdominal Imaging. Computation and Clinical Applications*. Vol. 8198. Lecture Notes in Computer Science. Berlin, Heidelberg, 2013, pp. 136–144 (cited on page 30).

[7] F. F. Berendsen, A. N. T. J. Kotte, M. A. Viergever, and J. P. W. Pluim. „Registration of organs with sliding interfaces and changing topologies." In: *Proceedings of SPIE, Medical Imaging* 9034 (2014), 90340E (cited on page 46).

[8] F. F. Berendsen, A. N. T. J. Kotte, A. A. C. de Leeuw, I. M. Jürgenliemk-Schulz, M. A. Viergever, and J. P. W. Pluim. „Registration of structurally dissimilar images in MRI-based brachytherapy." In: *Physics in Medicine and Biology* 59.15 (2014), pp. 4033–4045 (cited on page 28).

[9] K. K. Bhatia, J. Hajnal, A. Hammers, and D. Rueckert. „Similarity metrics for groupwise non-rigid registration." In: *Medical Image Computing and Computer-Assisted Intervention*. Vol. 4792. Lecture Notes in Computer Science. Berlin, Heidelberg, 2007, pp. 544–552 (cited on page 63).

[10] L. Bondar, M. S. Hoogeman, E. M. Vasquez Osorio, and B. J. Heijmen. „A symmetric nonrigid registration method to handle large organ deformations in cervical cancer patients." In: *Medical Physics* 37.7 (2010), pp. 3760–3772 (cited on page 10).

[11] K. Brock, M. Sharpe, L. Dawson, S. Kim, and D. A. Jaffray. „Accuracy of finite element model-based multi-organ deformable image registration." In: *Medical Physics* 32.6 (2005), pp. 1647–1659 (cited on page 9).

[12] L. Brown. „A survey of image registration techniques." In: *ACM computing surveys* 24.4 (1992), pp. 325–376 (cited on page 2).

[13] N. D. Cahill, J. A. Noble, and D. J. Hawkes. „A Demons Algorithm for Image Registration with Locally Adaptive Regularization." In: *Medical Image Computing and Computer-Assisted Intervention*. Vol. 5761. Lecture Notes in Computer Science. Berlin, Heidelberg, 8, 2009, pp. 574–581 (cited on page 45).

[14] E. Castillo, R. Castillo, J. Martinez, M. Shenoy, and T. Guerrero. „Four-dimensional deformable image registration using trajectory modeling." In: *Physics in Medicine and Biology* 55.1 (2010), pp. 305–327 (cited on pages 50, 57, 59).

[15] R. Castillo, E. Castillo, R. Guerra, V. E. Johnson, T. McPhail, A. K. Garg, and T. Guerrero. „A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets." In: *Physics in Medicine and Biology* 54.7 (2009), pp. 1849–1870 (cited on pages 50, 57, 59).

[16] N. Chitphakdithai and J. S. Duncan. „Non-rigid Registration with Missing Correspondences in Preoperative and Postresection Brain Images." In: *Medical Image Computing and Computer-assisted Intervention*. Vol. 6361. Lecture Notes in Computer Science. Berlin, Heidelberg, 2010, pp. 367–374 (cited on page 29).

[17] G. E. Christensen, B. Carlson, K. S. Chao, P. Yin, P. W. Grigsby, K. Nguyen, J. F. Dempsey, F. A. Lerma, K. T. Bae, M. W. Vannier, and J. F. Williamson. „Image-based dose planning of intracavitary brachytherapy: registration of serial-imaging studies using deformable anatomic templates." In: *International Journal of Radiation Oncology*Biology*Physics* 51.1 (2001), pp. 227–243 (cited on pages 29, 41).

[18] F. Commandeur, J. Velut, and O. Acosta. „A VTK Algorithm for the Computation of the Hausdorff Distance." In: *Midas Journal* (2011) (cited on page 36).

[19] T. Cootes, C. Taylor, D. Cooper, and J. Graham. „Training Models of Shape from Sets of Examples." English. In: *British Machine Vision Conference*. Ed. by D. Hogg and R. Boyle. 1992, pp. 9–18 (cited on pages 11, 13).

[20] L. Cordero-Grande, S. Merino-Caviedes, S. Aja-Fernandez, and C. Alberola-Lopez. „Groupwise elastic registration by a new sparsity-promoting metric: application to the alignment of cardiac magnetic resonance perfusion images." In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.11 (2013), pp. 2638–2650 (cited on page 63).

[21] W. R. Crum, T. Hartkens, and D. L. G. Hill. „Non-rigid image registration: theory and practice." In: *British journal of radiology* 77.2 (2004), pp. 140–153 (cited on pages 2, 9).

[22] V. Delmon, S. Rit, R. Pinho, and D. Sarrut. „Registration of sliding objects using direction dependent B-splines decomposition." In: *Physics in Medicine and Biology* 58.5 (2013), pp. 1303–1314 (cited on pages 45, 50, 52, 59, 60).

[23] A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen. „Automatic 3D ASM Construction via Atlas-Based Landmarking and Volumetric Elastic Registration." In: *Proceedings of the 17th International Conference on Information Processing in Medical Imaging*. London, UK, 2001, pp. 78–91 (cited on page 13).

[24] M. Freiman, S. D. Voss, and S. K. Warfield. „Demons registration with local affine adaptive regularization: application to registration of abdominal structures." In: *IEEE International Symposium on Biomedical Imaging*. Piscataway, NJ, USA, 2011, pp. 1219–1222 (cited on page 45).

[25] Z. Gu and B. Qin. „Nonrigid Registration of Brain Tumor Resection MR Images Based on Joint Saliency Map and Keypoint Clustering." In: *Sensors* 9.12 (2009), pp. 10270–10290 (cited on page 29).

[26] T. Heimann and H.-P. Meinzer. „Statistical shape models for 3D medical image segmentation: A review." In: *Medical Image Analysis* 13.4 (2009), pp. 543 –563 (cited on page 9).

[27] T. Heimann, S. Münzing, H.-P. Meinzer, and I. Wolf. „A shape-guided deformable model with evolutionary algorithm initialization for 3D soft tissue segmentation." In: *Information Processing in Medical Imaging*. Vol. 4584. Lecture Notes in Computer Science. 2007, pp. 1–12 (cited on page 10).

[28] D. L. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes. „Medical image registration." In: *Physics in Medicine and Biology* 46.3 (2001), pp. 1–45 (cited on page 2).

[29] W. Huizinga, C. T. Metz, D. H. J. Poot, M. de Groot, W. J. Niessen, A. Leemans, and S. Klein. „Groupwise Registration for Correcting Subject Motion and Eddy Current Distortions in Diffusion MRI Using a PCA Based Dissimilarity Metric." In: *Computational Diffusion MRI and Brain Connectivity*. Mathematics and Visualization. 2014, pp. 163–174 (cited on page 63).

[30] B. F. Hutton, M. Braun, L. Thurfjell, and D. Y. Lau. „Image registration: an essential tool for nuclear medicine." In: *European Journal of Nuclear Medicine and Molecular Imaging* 29.4 (2002), pp. 559–577 (cited on page 2).

[31] S. Joshi, B. Davis, M. Jomier, and G. Gerig. „Unbiased diffeomorphic atlas construction for computational anatomy." In: *Neuroimage* 23.1 (2004), pp. 151–160 (cited on page 63).

[32] S. Kabus, T. Netsch, B. Fischer, and J. Modersitzki. „B-Spline Registration of 3D Images with Levenberg-Marquardt Optimization." In: vol. 5370. 2004, pp. 304–313 (cited on page 14).

[33] S. Kiriyanthan, K. Fundana, and P. C. Cattin. „Discontinuity Preserving Registration of Abdominal MR Images with Apparent Sliding Organ Motion." In: *Proceedings of the 3rd International Conference on Abdominal Imaging*. Vol. 7029. Lecture Notes in Computer Science. Berlin, Heidelberg, 2011, pp. 231–239 (cited on page 45).

[34] S. Klein, J. P. W. Pluim, M. Staring, and M. A. Viergever. „Adaptive stochastic gradient descent optimisation for image registration." In: *International Journal of Computer Vision* 81.3 (2009), pp. 227–239 (cited on page 15).

[35] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. W. Pluim. „elastix: A toolbox for intensity-based medical image registration." In: *IEEE Transactions on Medical Imaging* 29.1 (2010), pp. 196–205 (cited on pages 5, 15, 17, 30, 32, 33, 47, 50, 75).

[36] E. G. Learned-Miller. „Data driven image models through continuous joint alignment." In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28.2 (2006), pp. 236–250 (cited on page 63).

[37] D. Loeckx, F. Maes, D. Vandermeulen, and P. Suetens. „Non-rigid image registration using a statistical spline deformation model." In: *Information Processing in Medical Imaging*. Vol. 2732. Lecture Notes in Computer Science. Berlin, Heidelberg, 2003, pp. 463–474 (cited on page 9).

[38] D. Mahapatra. „Joint segmentation and groupwise registration of cardiac perfusion images using temporal information." In: *Journal of Digital Imaging* 26.2 (2013), pp. 173–182 (cited on page 63).

[39] J. B. Maintz and M. A. Viergever. „A survey of medical image registration." In: *Medical Image Analysis* 2.1 (1998), pp. 1–36 (cited on page 2).

[40] D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank. „PET-CT Image Registration in the Chest Using Free-form Deformations." In: *IEEE Transactions on Medical Imaging* 22.1 (2003), pp. 120–128 (cited on pages 14, 15).

[41] C. T. Metz, S. Klein, M. Schaap, T. van Walsum, and W. J. Niessen. „Nonrigid registration of dynamic medical imaging data using nD + t B-splines and a groupwise optimization approach." In: *Medical Image Analysis* 15.2 (2011), pp. 238–249 (cited on page 63).

[42] M. I. Miga, D. W. Roberts, F. E. Kennedy, L. A. Platenik, A. Hartov, K. E. Lunn, and K. D. Paulsen. „Modeling of retraction and resection for intraoperative updating of images." In: *Neurosurgery* 49.1 (2001), pp. 75–84 (cited on page 29).

[43] H. W. Mulder, M. van Stralen, C. Ren, F. F. Berendsen, J. G. Bosch, and J. P. W. Pluim. „Simultaneous pairwise registration for image mosaicing of TEE data." In: *IEEE International Symposium on biomedical imaging*. Piscataway, NJ, USA, 2013, pp. 242–245 (cited on page 68).

[44] S. Nithiananthan, S. Schafer, D. J. Mirota, J. W. Stayman, W. Zbijewski, D. D. Reh, G. L. Gallia, and J. H. Siewerdsen. „Extra-dimensional Demons: a method for incorporating missing tissue in deformable image registration." In: *Medical Physics* 39.9 (2012), pp. 5718–5731 (cited on page 29).

[45] O. Oktay, L. Zhang, T. Mansi, P. Mountney, P. Mewes, S. Nicolau, L. Soler, and C. Chefd'hotel. „Biomechanically Driven Registration of Pre-to Intra-Operative 3D Images for Laparoscopic Surgery." In: *Medical Image Computing and Computer-Assisted Intervention*. Vol. 8150. Lecture Notes in Computer Science. Berlin, Heidelberg, 2013, pp. 1–9 (cited on page 29).

[46] D. F. Pace, M. Niethammer, and S. R. Aylward. „Sliding Geometries in Deformable Image Registration." In: *Proceedings of the 3rd International Conference on Abdominal Imaging*. Vol. 7029. Lecture Notes in Computer Science. Berlin, Heidelberg, 2011, pp. 141–148 (cited on page 45).

[47] X. Pennec, R. Stefanescu, V. Arsigny, P. Fillard, and N. Ayache. „Riemannian elasticity: a statistical regularization framework for non-linear registration." In: *Medical Image Computing and Computer-Assisted Intervention*. Vol. 3750. Lecture Notes in Computer Science. Berlin, Heidelberg, 2005. Chap. 116, pp. 943–950 (cited on page 9).

[48] S. Periaswamy and H. Farid. „Medical image registration with partial data." In: *Medical Image Analysis* 10.3 (2006), pp. 452–464 (cited on page 29).

[49] P. Risholm, E. Samsett, I. F. Talos, and W. Wells. „A non-rigid registration framework that accommodates resection and retraction." In: *Information Processing in Medical Imaging*. Vol. 5636. Lecture Notes in Computer Science. Berlin, Heidelberg, 2009, pp. 447–458 (cited on pages 29, 45).

[50] L. Risser, H. Y. Baluwala, and J. A. Schnabel. „Diffeomorphic registration with sliding conditions: Application to the registration of lungs CT images." In: *International Conference on Medical Image Computing and Computer-Assisted Intervention - Worshop on Pulmonary Image Analysis*. 2011, pp. 79–90 (cited on page 46).

[51] L. Risser, F. X. Vialard, H. Y. Baluwala, and J. A. Schnabel. „Piecewise-diffeomorphic image registration: application to the motion estimation between 3D CT lung images with sliding conditions." In: *Medical Image Analysis* 17.2 (2013), pp. 182–193 (cited on pages 45, 46).

[52] D. Ruan, S. Esedoglu, and J. A. Fessler. „Discriminative Sliding Preserving Regularization in Medical Image Registration." In: *IEEE International Symposium on Biomedical Imaging*. Piscataway, NJ, USA, 2009, pp. 430–433 (cited on page 45).

[53] D. Rueckert and P. Aljabar. „Nonrigid Registration of Medical Images: Theory, Methods, and Applications." In: *IEEE Signal Processing Magazine* 27.4 (2010), pp. 113–119 (cited on page 2).

[54] D. Rueckert, A. F. Frangi, and J. A. Schnabel. „Automatic construction of 3D statistical deformation models using non-rigid registration." In: *IEEE Transactions on Medical Imaging* 22.8 (2003), pp. 77–84 (cited on pages 9, 14).

[55] A. Schmidt-Richberg, R. Werner, H. Handels, and J. Ehrhardt. „Estimation of slipping organ motion by registration with direction-dependent regularization." In: *Medical Image Analysis* 16.1 (2012), pp. 150–159 (cited on pages 45, 50, 59).

[56] R. W. K. So and A. C. S. Chung. „Multi-modal non-rigid image registration based on similarity and dissimilarity with the prior joint intensity distributions." In: *From nano to Macro*. Proceedings of IEEE international conference on Biomedical imaging. Piscataway, NJ, USA, 2010, pp. 368–371 (cited on page 14).

[57] A. Sotiras, C. Davatzikos, and N. Paragios. „Deformable medical image registration: a survey." In: *IEEE Transactions on Medical Imaging* 32.7 (2013), pp. 1153–1190 (cited on page 2).

[58] 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." In: *IEEE Transactions on Medical Imaging* 28.9 (2009), pp. 1412–1421 (cited on pages 10, 23).

[59] M. Unser. „Splines: A Perfect Fit for Signal and Image Processing." In: *IEEE Signal Processing Magazine* 16.6 (1999). IEEE Signal Processing Society's 2000 magazine award, pp. 22–38 (cited on page 47).

[60] J. Vandemeulebroucke, O. Bernard, S. Rit, J. Kybic, P. Clarysse, and D. Sarrut. „Automated segmentation of a motion mask to preserve sliding motion in deformable registration of thoracic CT." In: *Medical Physics* 39.2 (2012), pp. 1006–1015 (cited on pages 50, 57).

[61] L. M. Vigneron, L. Noels, S. K. Warfield, J. G. Verly, and P. A. Robe. „Serial FEM/XFEM-Based Update of Preoperative Brain Images Using Intraoperative MRI." In: *International Journal of Biomedical Imaging* 2012 (2012), p. 872783 (cited on page 29).

[62] C. Wachinger and N. Navab. „Simultaneous registration of multiple images: similarity metrics and efficient optimization." In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.5 (2013), pp. 1221–1233 (cited on pages 63, 64).

[63] Z. Wu, E. Rietzel, V. Boldea, D. Sarrut, and G. C. Sharp. „Evaluation of deformable registration of patient lung 4DCT with subanatomical region segmentations." In: *Medical Physics* 35.2 (2008), pp. 775–781 (cited on pages 45, 50, 52, 59, 60).

[64] Y. Xie, M. Chao, and G. Xiong. „Deformable Image Registration of Liver With Consideration of Lung Sliding Motion." In: *Medical Physics* 38.10 (2011), pp. 5351–5361 (cited on page 45).

[65] Z. Xue, D. Shen, and C. Davatzikos. „Statistical representation of high-dimensional deformation fields with application to statistically constrained 3D warping." In: *Medical Image Analysis* 10.5 (2006), pp. 740–751 (cited on page 9).

[66] Y. Yin, E. A. Hoffman, and C. L. Lin. „Lung lobar slippage assessed with the aid of image registration." In: *Medical Image Computing and Computer-Assisted Intervention*. Vol. 6362. Lecture Notes in Computer Science. Berlin, Heidelberg, 2010, pp. 578–585 (cited on page 45).

[67] S. Ying, G. Wu, Q. Wang, and D. Shen. „Hierarchical unbiased graph shrinkage (HUGS): a novel groupwise registration for large data set." In: *Neuroimage* 84 (2014), pp. 626–638 (cited on page 63).

[68] E. I. Zacharaki, C. S. Hogea, D. Shen, G. Biros, and C. Davatzikos. „Non-diffeomorphic registration of brain tumor images by simulating tissue loss and tumor growth." In: *Neuroimage* 46.3 (2009), pp. 762–774 (cited on page 29).

Nederlandse samenvatting

Beeldregistratie is een belangrijke taak in de medische beeldverwerking. Voorbeelden van toepassingen zijn: (1) registratie tussen patienten om handmatige orgaansegmentaties van de ene patient over te brengen op de andere patient, (2) registratie van de scans uit een reeks radiotherapie behandelingen om een gecombineerde weefselspecifieke stralingsdosis te kunnen bepalen, en (3) registratie over tijd of binnen populaties om vervormingsanalyses uit te kunnen voeren. Over het algemeen moeten er instellingen gekozen worden voor het registratiealgoritme, afhankelijk van het type beeld en de soort ruimelijke vervorming tussen de beelden. De beeldgelijkenismaat, het type optimalisatie en het aantal graden vrijheid van het transformatiemodel moeten bijvoorbeeld gekozen worden. Voor verscheidene toepassingen zijn de instellingskeuzes van een algemeen registratiealgoritme te gelimiteerd om een goed registratieresultaat te verkrijgen, omdat ze geen gebruik maken van toepassingsspecifieke geometrische kennis. Bepaalde toepassingen kunnen profiteren van voorkennis over vormen in een populatie, over de geometrische eigenschappen van structuren in de afbeelding of over discontinuïteiten in het te verwachten deformatieveld. In dit proefschrift worden verschillende uitbreidingen op een algemeen registratiealgoritme voorgesteld om het geschikt te maken voor de problemen die in verscheidene toepassingen voorkomen.

Hoofdstuk 2 Deformeerbare registratie is foutgevoelig als er grote en complexe vervormingen in het spel zijn. De registratieprocedure zal namelijk waarschijnlijk in een lokaal minimum belanden. Regularisatie op basis van voorkennis kan opgenomen worden in de registratie om daarmee het aantal lokale minima te verminderen en dus het risico op misregistratie te verkleinen. Wij stellen een regularisatieterm voor die gebaseerd is op statistische voorkennis over de te verwachten deformaties. Een statistisch model wordt daarvoor getraind op de vormen uit verzameling segmentaties en wordt geïntegreerd als een strafmaat in de registratieprocedure. Ter evaluatie van onze methode registreren we de MR-beelden, die gemaakt zijn voor de planning van de bestraling van baarmoederhalskanker, van verschillende patiënten naar elkaar. Handmatige intekeningen van structuren zoals de blaas en het klinische doelvolume zijn hierbij beschikbaar. Voor beide structuren werden leave-one-patient-out registratie-experimenten uitgevoerd. De segmentaties verkregen met onze me-

thode werden vergeleken met de handmatige segmentaties op basis van Dicescore en Hausdorffafstand. In vergelijking met de registratie zonder het gebruik van statistische kennis waren de segmentaties significant verbeterd. Gemiddeld was dit voor beide structuren met 0,1 punt Dicescore en met 8 mm op Hausdorffafstand.

Hoofdstuk 3 Een grote uitdaging in beeldregistratie is het nauwkeurig uitlijnen van de twee beelden als een bepaalde structuur alleen aanwezig is in één van beide beelden. Dergelijke topologische veranderingen zijn problematisch voor conventionele deformeerbare registratiealgoritmes. Wij stellen voor om een geometrische strafterm op te nemen in het conventionele 'free-form' registratiekader die het volume van de ontbrekende structuur in één van de beelden minimaliseert. We demonstreren onze methode aan de hand van baarmoederhalskanker MR-beelden voor brachytherapie. In dit intra-patiënt registratieprobleem is er één beeld waarin een therapie-applicator aanwezig is en een ander waarin deze niet aanwezig is. Het toevoegen van de strafterm resulteerde in een aanzienlijke verbetering van de oppervlakteafstandsmaat tot de anatomische positie van de gouden standaard en van het restvolume van de leegte die de applicator achterlaat. Registratie van aangrenzende structuren, gemeten aan de hand van het rectum en de blaas, wordt over het algemeen verbeterd, maar in mindere mate.

Hoofdstuk 4 Door gladheids- en continuïteitsaannames omtrent het deformatieveld is deformeerbare beeldregistratie niet geschikt voor toepassingen waar objecten langs elkaar schuiven. Recente uitbreidingen die deformeerbare beeldregistratie geschikt maken voor de glijdende beweging van organen zijn beperkt tot glijdende bewegingen langs bij benadering rechte oppervlaktes of zijn niet in staat tot het modelleren van verschuivingen waarbij de topologische configuratie van meerdere organen verandert. Wij stellen een uitbreiding van 'free-form' beeldregistratie voor die deze beperkingen niet heeft. Onze methode bestaat uit een vervormingsmodel dat is opgebouwd uit voor iedere orgaanregio afzonderlijke uniforme B-spline transformaties die zijn gebaseerd op de segmentatie van één van de afbeeldingen. Omdat met dit model overlappende regio's of leegtes tussen de gebieden kunnen ontstaan, introduceren we een strafterm die dit ongewenste effect minimaliseert. De strafterm werkt op de oppervlakten van de orgaanregios en wordt gelijktijdig geoptimaliseerd met de beeldgelijkenismaat. Om onze methode te evalueren werden registraties uitgevoerd op publiekelijk beschikbare in- en uitademhalings-CT-scans waarop de prestaties van andere methoden bekend is. De registratiefouten worden gemeten aan de hand van de verzameling landmarkpunten die bij de datasets horen. Op deze data presteert onze methode beter dan de andere methoden in termen van registratiefout en, indien van toepassing, ook in termen van de volumes van overlap en leegte.

Hoofdstuk 5 Een template-vrije groepsgewijze registratiemethode kan worden geformuleerd als een gelijktijdige registratie van alle beeldpaarcombinaties mogelijk in een groep. Een dergelijke groepsgewijze beeldgelijkenismaat

voor registratie heeft typisch een computationele complexiteit die kwadratisch schaalt met het aantal beelden in de groep. Er bestaat al een groepsgewijze beeldgelijkingsmaat met lineaire computationele complexiteit die mathematisch gelijk is aan de som van de kwadratische verschillen binnen elk beeldpaar. Wij breiden dit uit door een beeldmaat voor te stellen met lineaire computationele complexiteit die gebaseerd is op de genormaliseerde kruis-correlatie tussen alle beeldparen. Onze beeldgelijkenismaat is geschikt voor de registratie van de tijdreeksen en beeldgroepen waar er voldoende overlap is in het gebied van interesse, aangezien onze beeldgelijkenismaat alleen gedefiniëerd is voor het gemeenschappelijke deel van *alle* beelden. Daarnaast stellen we ook een beeldgelijkenismaat voor waarbij de overlap van de beelden wel mag variëren tussen de paren, maar hierbij is de beeldnormalisatie alleen bij benadering correct. Voor bij het gebruik van gradiënt gebaseerde optimalisaties geven we ook de analytische afgeleides van de beeldmaat.

Beeldregistratie kent vele toepassingen en vormt vaak de basis voor verdere verwerking of analyse. De hoofdstukken van dit proefschrift presenteren diverse registratiemethoden, elk uitgebreid voor een specifieke toepassing. Omdat dit proefschrift zich richt op registratiealgoritmes met een bredere toepasbaarheid voor ogen, hebben beeldverwerkingstechnieken voor een onderliggend specifiek klinisch doel minder aandacht gekregen. Elk hoofdstuk in dit proefschrift biedt een losstaande uitbreiding op het registratiealgoritme om een specifiek probleem in de registratie aan te pakken. Echter, in verscheidene toepassingen kunnen meerdere van die problemen in meer of mindere mate tegelijkertijd een rol spelen. Door zich alleen op een specifiek aspect van de toepassing te richten, kan het zijn dat registratie kwaliteit beperkt wordt door de andere aspecten die genegeerd worden. Aangezien de voorgestelde methoden als een strafterm geformuleerd zijn die bijdraagt aan de kostenfunctie (hoofdstukken 2 tot en met 4) en/of als een variant op de conventionele methode (hoofdstukken 4 en refchapter:groupwise), zou vervolgonderzoek zich kunnen richten op het combineren van meerdere methoden tot een enkel registratieproces.

Alle voorgestelde methodes zijn geïmplementeerd in de open source registratietoolbox `elastix`, waardoor het mogelijk wordt voor andere onderzoekers om het te gebruiken voor of aan te passen aan hun eigen toepassingen. De auteur hoopt dan ook dat door het uitbreiden van de toepassingsmogelijkheden van registratiealgoritmes, beeldregistratie een nog veelzijdiger instrument is voor het gebruik in klinisch onderzoek en uiteindelijk in de klinische praktijk.

Dankwoord

Dit proefschrift was niet tot stand gekomen zonder hulp van velen. Ik wil iedereen die, op wat voor een wijze dan ook, aan dit proefschrift heeft bijgedragen bedanken. Met name wil ik de volgende personen noemen. Allereerst een woord van dank aan Josien. Ik waardeer enorm je altijd snelle en constructieve terugkoppelingen op mijn werk. Ik heb veel van je geleerd in het proces van idee naar wetenschappelijke publicatie. Dat het Image Science Institute zo'n goed geoliede machine is waarin kwaliteit en fijne werkomgeving samen gaan, is boven iedereen te danken aan Max. Bedankt ook voor al je inbreng en voor het meedenken tijdens de diverse moeilijkheden van het promoveren. Robin, bedankt voor al je productiviteit als mede-auteur én mede-bandlid. Ik erg ben blij dat we elkaar nog steeds niet alleen voor werk ontmoeten, maar ook daarbuiten. Voor de fijne samenwerking met de afdeling Radiotherapie ben ik een aantal mensen dankbaar: Alexis, Uulke, Astrid en Ina. Jullie nieuwe inzichten, inzet om mij van hoge kwaliteit intekeningen te voorzien en verhelderingen omtrent klinische aspecten, hebben mij enorm geholpen. Paul, als kantoorgenoot van begin tot eind heb ik veel aan je gehad. Vooral door je welkome, met name niet werk gerelateerde afleidingen tussendoor. Desondanks heb ik veel van jouw onderwerp geleerd, ik hoop jij ook van het mijne. Colleague-next-door Mitko: thanks for all the technical and scientific discussions we had. I always felt welcome in your office to chat about anything, not excluding our shared interest in music and electronics. Alle collega's door de jaren heen wil ik bedanken voor de gezellige koffiepauses en (spontane) borrels, voor discussies (met taart of chocolade erbij) en de lol tijdens het filmen van 'stukjes'. Ook wil ik alle vrienden & studiegenoten (indertijd) uit Enschede bedanken voor alle gezellige activiteiten en weekendjes weg op regelmatige en soms minder regelmatige basis. Daarnaast natuurlijk ook m'n ouders, Evert en Louise, en m'n broers Jasper, Rikkert en Arnout, en Oma: bedankt voor het nog altijd warme welkom 'thuis'. Als afsluiter het mooiste resultaat van m'n promotietijd in Utrecht; Nynke, je bent geweldig.

List of publications

Papers in international journals

- **Floris F. Berendsen**, Josien P. W. Pluim: 'Effcient normalized cross-correlation for groupwise registration', [in preparation]
- **Floris F. Berendsen**, Alexis N. T. J. Kotte, Max A. Viergever, Josien P. W. Pluim: 'Registering organs with sliding interfaces and changing topologies', [submitted for review]
- Thomas R. Langerak, Uulke A. van der Heide, Alexis N. T. J. Kotte, **Floris F. Berendsen**, Josien P. W. Pluim:'Improving label fusion in multi-atlas based segmentation by locally combining atlas selection and performance estimation', Computer Vision and Image Understanding (CVIU), In Press
- **Floris F. Berendsen**, Alexis N. T. J. Kotte, Ina M. Jürgenliemk-Schulz, Max A. Viergever, Josien P. W. Pluim: 'Registration of structurally dissimilar images in MRI-based brachytherapy', Physics in Medicine and Biology (PMB), 59(15):4033–4045, 2014
- Thomas R. Langerak, Uulke A. van der Heide, Alexis N. T. J. Kotte, **Floris F. Berendsen**, Marco van Vulpen, Josien P. W. Pluim: 'Expert-driven label fusion in atlas-based segmentation using weighted atlases', International Journal of Computer Assisted Radiology and Surgery (CARS), 8(6):929–936, 2013
- **Floris F. Berendsen**, Uulke A. van der Heide, Thomas R. Langerak, Alexis N. T. J. Kotte, Josien P. W. Pluim: 'Free-form image registration regularized by a statistical shape model: application to organ segmentation in cervical MR', Computer Vision and Image Understanding (CVIU), 117(9):1119–1127, 2013
- Thomas R. Langerak, **Floris F. Berendsen**, Uulke A. van der Heide, Alexis N. T. J. Kotte, Josien P. W. Pluim: 'Multi-atlas-based segmentation with pre-registration atlas selection', Medical Physics, 40(9):091701, 2013

Papers in conference proceedings

- **Floris F. Berendsen**, Alexis N. T. J. Kotte, Max A. Viergever, Josien P. W. Pluim: 'Registration of organs with sliding interfaces and changing topologies'[1], SPIE Medical Imaging, 90340E, San Diego, United States, 2014

- **Floris F. Berendsen**, Alexis N. T. J. Kotte, Astrid A. C. de Leeuw, Max A. Viergever, Josien P. W. Pluim: 'Free-Form Registration Involving Disappearing Structures: Application to Brachytherapy MRI', MICCAI Workshop on Abdominal Imaging: Computational and Clinical Applications, Lecture Notes in Computer Science 8198:136–144, Nagoya, Japan, 2013

- Harriët W. Mulder, Marijn van Stralen, Claire Ren, **Floris F. Berendsen**, Johan G. Bosch, Josien P. W. Pluim: 'Simultaneous pairwise registration for image mosaicing of TEE data', IEEE International Symposium on Biomedical Imaging (ISBI), 242–245, San Francisco, United States, 2013

- **Floris F. Berendsen**, Uulke A. van der Heide, Thomas R. Langerak, Alexis N. T. J. Kotte, Josien P. W. Pluim: 'Segmentation of Cervical Images by Inter-subject Registration with a Statistical Organ Model', MICCAI Workshop on Abdominal Imaging: Computational and Clinical Applications, Lecture Notes in Computer Science, 7029:240–247, Toronto, Canada, 2011

- Thomas R. Langerak, Uulke A. Van der Heide, Alexis N. T. J. Kotte, **Floris F. Berendsen**, Josien P. W. Pluim: 'Multi-atlas-based segmentation with pre-registration atlas selection', MICCAI Workshop on Multi-Atlas Labeling and Statistical Fusion, 1–9, Toronto, Canada, 2011

- Thomas R. Langerak, Uulke A. van der Heide, Alexis N. T. J. Kotte, **Floris F. Berendsen**, Josien P. W. Pluim: 'Local atlas selection and performance estimation in multi-atlas based segmentation', IEEE International Symposium on Biomedical Imaging (ISBI), 669–672, Chicago, United States, 2011

- Thomas R. Langerak, Uulke A. van der Heide, Alexis N. T. J. Kotte, **Floris F. Berendsen**, Josien P. W. Pluim: 'Label fusion in multi-atlas based segmentation with user-defined local weights', IEEE International Symposium on Biomedical Imaging (ISBI), 1480–1483, Chicago, United States, 2011

- Ferdinand van der Heijden, **Floris F. Berendsen**, Luuk J. Spreeuwers, E. Schippers: 'Particle Smoothing for Solving Ambiguity Problems in

[1] Nominated for SPIE 2014 Robert F. Wagner All-Conference Best Student Paper Award

One-shot Structured Light Systems', International Conference on Computer Vision Theory and Applications (VISAPP), 531–537, Vilamoura, Portugal, 2011

- Thomas R. Langerak, Uulke A. Van der Heide, Alexis N. T. J. Kotte, **Floris F. Berendsen**, Josien P. W. Pluim: 'Evaluating and improving label fusion in atlas-based segmentation using the surface distance', SPIE Medical Imaging, 796226, San Diego, United States, 2010

Curriculum vitae

Floris F. Berendsen was born on July 28th, 1983 in Hengelo (O), the Netherlands. From 2001 to 2009 he studied Electrical Engineering at the University of Twente in Enschede, the Netherlands. His MSc project was on 3-dimensional object reconstruction from structured light by particle filtering methods. Since 2009 he has been working as a PhD candidate at the University Medical Center in Utrecht, The Netherlands on the topic of image registration. The results from this research are presented in this thesis. Since 2014, he is working at the Division of Image Processing (LKEB) of Leiden University Medical Center, the Netherlands.