

Nonrigid registration with tissue-dependent filtering of the deformation field

Marius Staring, Stefan Klein and Josien P W Pluim

Image Sciences Institute, University Medical Center Utrecht, PO Box 85500, 3508 GA, Room Q0S.459, Utrecht, The Netherlands

E-mail: marius@isi.uu.nl

Received 4 May 2007, in final form 18 October 2007

Published 8 November 2007

Online at stacks.iop.org/PMB/52/6879

Abstract

In present-day medical practice it is often necessary to nonrigidly align image data. Current registration algorithms do not generally take the characteristics of tissue into account. Consequently, rigid tissue, such as bone, can be deformed elastically, growth of tumours may be concealed, and contrast-enhanced structures may be reduced in volume. We propose a method to locally adapt the deformation field at structures that must be kept rigid, using a tissue-dependent filtering technique. This adaptive filtering of the deformation field results in locally linear transformations without scaling or shearing. The degree of filtering is related to tissue stiffness: more filtering is applied at stiff tissue locations, less at parts of the image containing nonrigid tissue. The tissue-dependent filter is incorporated in a commonly used registration algorithm, using mutual information as a similarity measure and cubic B-splines to model the deformation field. The new registration algorithm is compared with this popular method. Evaluation of the proposed tissue-dependent filtering is performed on 3D computed tomography (CT) data of the thorax and on 2D digital subtraction angiography (DSA) images. The results show that tissue-dependent filtering of the deformation field leads to improved registration results: tumour volumes and vessel widths are preserved rather than affected.

1. Introduction

Nonrigid image registration is an important technique in the field of medical image processing (Hill *et al* 2001, Lester and Arridge 1999, Maintz and Viergever 1998). It performs the task of finding the spatial correspondence between two images, i.e. it searches for a one-to-one mapping from voxels in one image to voxels in the other image. This mapping is called the transformation. Commonly, first a rigid or affine registration is performed to capture the global transformation. The rigid or affine transformation is then used as a starting point

for a nonrigid registration, which searches for local deformations. An example of a popular nonrigid registration algorithm is that of Rueckert *et al* (1999).

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

- (a) *Truly rigid tissue.* Bone tissue that is present in the image volume must not be allowed to deform. Treating the bones as deformable tissue can lead, for example, to undesired thickening and bending of ribs (Staring *et al* 2005).
- (b) *Temporal changes of tissue.* An example is the study of tumour progression in a patient. A successful nonrigid registration will achieve a perfect fit between tumours of two different time points, effectively concealing tumour growth or shrinkage. To visualize the tumour growth with a difference image, the tumours have to be kept rigid.
- (c) *Intensity changes of tissue.* Local intensity changes can be induced, for example, by injection of a contrast agent. Rohlfing *et al* (2003) and Tanner *et al* (2000) demonstrate that contrast-enhanced breast lesions in MR are compressed after intensity-based nonrigid registration. In Rohlfing and Maurer (2001) shrinkage is reported for contrast-enhanced vessels in CT digital subtraction angiography (CT-DSA).

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

Several methods are reported in the literature to address this issue. Some of these methods add a regularization or penalty term to the registration, thereby constraining the transformation. Rohlfing *et al* (2003) achieve incompressibility by penalizing a deviation of the determinant of the Jacobian of the transformation from 1. The method does not distinguish between different tissue types, nor does it enforce rigidity of stiff tissue. Loeckx *et al* (2004) employ a rigidity penalty term to certain structures by penalizing a deviation of the Jacobian from orthonormality. Other ways to enforce rigidity of structures have also been presented. Tanner *et al* (2000) propose to couple the control points of a B-spline deformation to enforce rigidity on certain structures. It is assumed that the rotational part of the deformation can be captured by the initial rigid registration. Another approach is taken by Little *et al* (1997), who constrain the nonlinear part of a deformation at rigid locations by multiplication with a weight function. Multiple rigid transformations are used by Arsigny *et al* (2005), each defined on different parts of the image, but influencing each other. Edwards *et al* (1998) model the deformation field using a three-component deformation model, consisting of rigid, elastic and fluid structures. Brock *et al* (2005) have employed a finite element method to introduce tissue-specific parameters. For the viscous fluid algorithm a locally adaptable regularization technique is proposed that spatially varies the viscosity parameter (Lester *et al* 1999).

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

Filtering of deformation fields has been used in other situations, for example in optical flow like registration algorithms. In these cases the filtering is typically used to regularize the deformation field to get more plausible results. Thirion (1998) uses a Gaussian kernel to smooth the deformation field. This kernel is not adapted to the underlying tissue type; it is fixed for the entire image volume. Cachier *et al* (1999) use the Demons algorithm, which they adapt with a filtering of the deformation field. They derive and employ *a priori* and *a posteriori* smoothing weights for the filtering; these weights vary over the image, but do not depend on the tissue type. Another method to smooth the deformation field is used by

Stefanescu *et al* (2004). They regularize the deformation field by smoothing it with a Gaussian kernel with a standard deviation depending on the spatial position. They propose to take a large standard deviation in areas where little deformation is expected (e.g., rigid areas), and vice versa. At boundaries of rigid and nonrigid tissue the deformation field within the rigid objects is influenced by the nonrigid deformations of its surroundings. Both the extent and strength of this influence are quite large, because a large standard deviation is used. Also, a scaling of the deformation field within rigid objects is not prevented with this method. Therefore, that approach is not suitable for the problem we address. To ensure rigidity of the deformation field, an asymmetric kernel is needed with a large standard deviation inside the rigid object and a small one outside.

The filter we propose has the property that undesired deformations at rigid structures are addressed, while the surroundings are free to deform. The deformation field is adapted by filtering the deformation field in a tissue-dependent way: the filter is different for each voxel, depending on the stiffness of the underlying tissue. In this way some parts of the image can be kept rigid, whereas other parts are allowed to deform freely. In section 2 the proposed tissue-dependent filter for regularizing the deformation field is described in detail. Incorporation of the filter in the registration algorithm is also described in that section.

For evaluation two applications are considered: the study of tumour progression in CT thorax imaging (an example of case (b)) and the study of the vasculature at several regions in the human body in DSA imaging (an example of case (c)). The validation of the algorithm is described in section 3, and the paper is concluded in section 4.

2. Method

Registration of a moving image $I_M(\mathbf{x}) : \Omega_M \subset \mathbb{R}^D \mapsto \mathbb{R}$ to a fixed image $I_F(\mathbf{x}) : \Omega_F \subset \mathbb{R}^D \mapsto \mathbb{R}$, both of dimension D , is the problem of finding the deformation field \mathbf{d} , or equivalently the transformation $\mathbf{T}(\mathbf{x}) = \mathbf{x} + \mathbf{d}(\mathbf{x})$, that spatially aligns $I_M(\mathbf{x} + \mathbf{d}(\mathbf{x}))$ and $I_F(\mathbf{x})$. Registration can be defined as a minimization problem:

$$\hat{\mathbf{d}} = \arg \min_{\mathbf{d}} \mathcal{C} [I_F(\mathbf{x}), I_M(\mathbf{x} + \mathbf{d}(\mathbf{x}))], \quad (1)$$

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

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

2.1. Filter design

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

$$\mathbf{m}(\mathbf{x}) \triangleq \sum_{\mathbf{x} \in \mathcal{N}_x} c(\mathbf{T}(\mathbf{x})) \mathbf{d}(\mathbf{x}) / \sum_{\mathbf{x} \in \mathcal{N}_x} c(\mathbf{T}(\mathbf{x})), \quad (2)$$

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

coefficient $c(\mathbf{x})$ is high, and a value close to the original deformation $\mathbf{d}(\mathbf{x})$ for low stiffness coefficients:

$$\mathcal{F}(\mathbf{d}(\mathbf{x})) \triangleq (1 - c(\mathbf{T}(\mathbf{x})))\mathbf{d}(\mathbf{x}) + c(\mathbf{T}(\mathbf{x}))\mathbf{m}(\mathbf{x}). \quad (3)$$

For neighbourhoods on the edge of rigid and nonrigid structures, the weighted mean $\mathbf{m}(\mathbf{x})$ is determined mainly by the voxels having a high stiffness coefficient. In other parts of the image equation (2) yields an averaging, also in the nonrigid parts. By defining the tissue-dependent filtering as in (3), it is controlled to what extent the weighted mean is used, such that it is only employed at rigid parts of the image. The result is a more homogeneous, linear deformation field at regions with a high $c(\mathbf{x})$, where tissue is more stiff. At regions with a low $c(\mathbf{x})$ the deformation field is unaffected by the filtering, and therefore those regions can deform freely, as desired.

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

We are interested in the deformed moving image. Therefore, we only have to take care that the transformation of that image, as found by the registration, obeys restrictions of rigidity. This means that only the relevant structures in the moving image have to be segmented.

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

$$\mathcal{F}^k(\mathbf{d}(\mathbf{x})) \triangleq \mathcal{F}(\mathcal{F}^{k-1}(\mathbf{d}(\mathbf{x}))), \quad (4a)$$

$$\mathcal{F}^0(\mathbf{d}(\mathbf{x})) \triangleq \mathbf{d}(\mathbf{x}). \quad (4b)$$

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

The neighbourhood \mathcal{N}_x in (2) is another parameter that can be controlled. It is chosen to be a square in 2D and a cube in 3D, of size $l \times l$ and $l \times l \times l$ voxels, respectively. For large l the adaptive filter will linearize larger parts of the image at once, thereby increasing the power of the filter. However, for large l , the deformation fields of two separate rigid structures lying close together might both be within the neighbourhood. This way, the deformation field of both rigid structures evolves to the average of them, which is undesired. Therefore, a small neighbourhood of $l = 3$ is used in this study. The power of the tissue-dependent filter is controlled with the number of successive applications k of the filter (see (4a)). In the following section a second parameter m is introduced, which also controls the power of the filter.

Figure 1(d) illustrates that the effect of the tissue-dependent filter is first observed near the boundary of rigid structures. After successive applications of the filter, the results are propagated to the inside of such regions. Strong edges in the stiffness coefficient image at the boundary of rigid and nonrigid tissue can result in sharp transitions in the deformation field after filtering. Invertibility of the deformation field is not guaranteed. This can be easily achieved by defining some rules on the displacement field (Stefanescu *et al* 2004). However, no problems were encountered in the experiments and therefore this step is omitted. The tissue-dependent filter is based on a weighted mean, therefore, the filtered deformation field will eventually, for large k , evolve to a translational field, not allowing rotations of rigid objects.

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

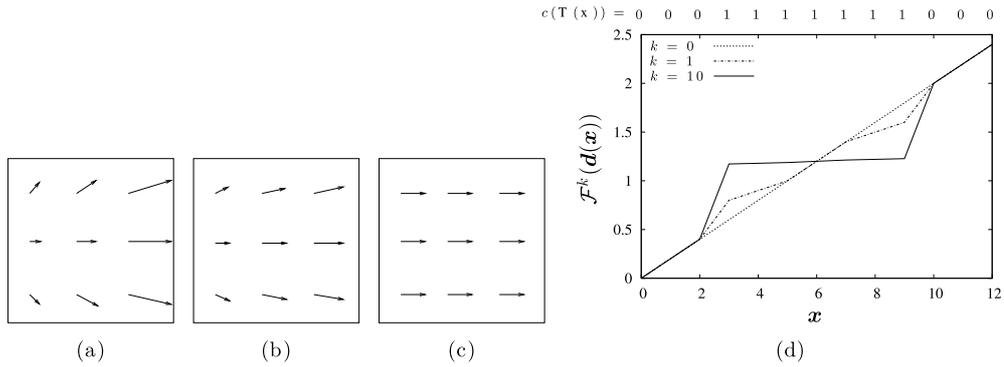


Figure 1. The original 2D deformation field (a) stretches the underlying tissue both vertically and horizontally. Applying the tissue-dependent filter, with $c(x) = 1$ everywhere, once to this field results in (b). This field still stretches tissue, but less. After applying the adaptive filter several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

2.2. Tissue-dependent nonrigid registration

This section describes how the tissue-dependent filter \mathcal{F} is incorporated in the registration. The registration algorithm is largely based on the papers of Rueckert *et al* (1999), Mattes *et al* (2003) and Thévenaz and Unser (2000).

The methods described by these authors use the well-known *mutual information* between the fixed image I_F and the moving image I_M as a cost function \mathcal{C} . In this study the implementation of Thévenaz and Unser (2000) is adopted, who use B-spline Parzen windows to describe the joint histogram. This way the cost function is a continuous, differentiable function of the transformation parameters p . The derivative of the cost function $\partial\mathcal{C}/\partial p$ can, therefore, be computed analytically. The deformation field d is parametrized by cubic B-splines (Unser 1999) with parameters p^{BS} . The B-splines modelling the deformation field have an inherent smoothness, which is desired in registration problems. They also have local support, which is beneficial both for modelling local transformations and for fast computation. The derivative $\partial d/\partial p^{BS}$, needed for computing $\partial\mathcal{C}/\partial p$, is available analytically.

An iterative optimization strategy is used to solve the minimization problem (1):

$$d_{t+1}(x) = d_t(x) + \Delta_t(x; p^{BS}), \tag{5}$$

where Δ_t is an incremental deformation field at iteration t , as found by the optimizer. Incorporating our tissue-dependent filtering of the deformation field in the optimization scheme is straightforward:

$$d_{t+1}(x) = \mathcal{F}^k(d_t(x) + \Delta_t(x; p^{BS})). \tag{6}$$

The resulting deformation $d_{t+1}(x)$ is a filtered version of a B-spline based deformation field for $t > 1$ and it cannot be parametrized by B-splines. It is described by a vector field. Only the additional deformation field $\Delta_t(x; p^{BS})$ is modelled by B-splines, which has fewer restrictions than modelling the entire transformation with B-splines. The optimization is performed on

$\Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$, i.e., on the B-spline coefficients only. This is computationally advantageous compared to optimizing over the entire deformation field.

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

- (i) perform $m > 0$ iterations of the optimizer, resulting in a new additional deformation field $\Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$;
- (ii) calculate the deformation field $\tilde{\mathbf{d}}_{t+1}(\mathbf{x}) = \mathbf{d}_t(\mathbf{x}) + \Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$;
- (iii) filter the deformation field, i.e., calculate $\mathbf{d}_{t+1}(\mathbf{x}) = \mathcal{F}^k(\tilde{\mathbf{d}}_{t+1}(\mathbf{x}))$ according to equations (2), (3) and (4a), (4b).

Repeat steps (i)–(iii) until convergence.

A stochastic gradient descent optimizer is used to determine the additional deformation field $\Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}})$ in the following way: $\Delta_t(\mathbf{x}; \mathbf{p}^{\text{BS}}) = -a_t \partial \mathcal{C} / \partial \mathbf{p}^{\text{BS}}$, where a_t is the size of the step taken in the direction $-\partial \mathcal{C} / \partial \mathbf{p}^{\text{BS}}$. The stochastic gradient descent optimizer uses an *approximation* of the derivative of the mutual information with respect to the B-spline parameters \mathbf{p}^{BS} . Approximation is done by using only a small random subset of voxels from the fixed image. This has been shown to accelerate registration significantly, without compromising registration accuracy (Klein *et al* 2007). At each iteration a step is taken towards the minimum of the cost function \mathcal{C} . The size of this step a_t is determined by a decreasing function of the iteration number t . This function is of the form $a_t = a / (t + A)^\alpha$, where a , A and α are user-defined constants. Following the suggestions in Spall (1998), in this work $A = 100.0$ and $\alpha = 0.602$ are used. The parameter a is related to the expected magnitude of the deformation, and is tuned for each application. As a stopping condition a user-defined number of iterations is used, upon which convergence is assumed. In order to avoid local minima, a multiresolution approach is taken. A Gaussian image pyramid is used with a subsampling factor of 2. Also a multigrid approach is taken: when the image resolution is doubled, the B-spline control point spacing is halved. Prior to the nonrigid registration an affine registration is performed in order to capture the global transformation between the fixed and the moving image.

3. Experiments and results

In order to evaluate the effectiveness of the proposed tissue-dependent nonrigid registration, it is compared with a general nonrigid registration approach based only on B-splines. This general approach does not apply the tissue-dependent filtering of the deformation field, but is similar to the proposed algorithm in all other respects (see the description in section 2.2). The methods will be referred to as ‘BS’ (B-spline only based) and ‘BSF’ (B-splines with tissue-dependent filtering of the deformation field).

The two methods are illustrated on a synthetic example (section 3.1), and compared on clinical data, namely 3D CT follow-up data of the thorax containing lung tumours (section 3.2) and 2D digital subtraction angiography (DSA) image data, see section 3.3. The applications are examples of cases (b) and (c) of the list given in the introduction. Both applications consider structures that are not physically rigid, but to prevent undesired volume changes, they should be regarded as undeformable. Special attention is given to the influence of the parameters k and m .

In all experiments with BSF, the size of the neighbourhood \mathcal{N}_x in equation (2) is 3×3 voxels in 2D and $3 \times 3 \times 3$ in 3D. For both BS and BSF, the number of bins for the calculation of the mutual information was set to 32. All experiments were performed with

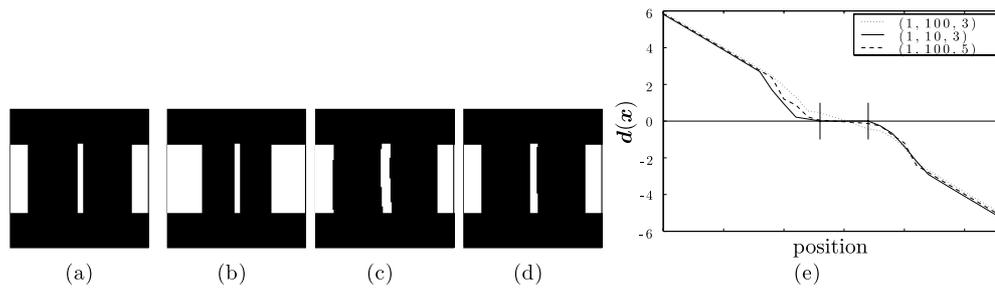


Figure 2. BS registration compared with the adaptive filtering of the deformation field. (a) The fixed image I_F , (b) the moving image I_M , (c) the result with BS and (d) the result with BSF ($k = 1, m = 100, l = 3$). When the neighbourhood size is increased to $l = 5$, or when more filtering is applied ($m = 10$), the resulting images are equal to the fixed image. In (e) a profile in the horizontal direction of the x -component of the deformation field is depicted. In the legend the parameters (k, l, m) refer to the corresponding BSF registration. The two vertical lines indicate the rigid part.

software developed by the authors (www.isi.uu.nl/Elastix). This registration package is largely based on the Insight Segmentation and Registration Toolkit (Ibáñez *et al* 2005). Filtering the deformation field took approximately 1 s. and 1.7 s. for $k = 1$ and $k = 10$, respectively, on 2D 512^2 data. For 3D 256^3 data this took 140 and 275 s. approximately. An AMD Opteron 2218 running at 2.6 GHz. was used, without multi-threading enabled.

3.1. 2D Synthetic example

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

3.2. 3D CT thorax data with lung tumours

Nonrigid registration is a valuable tool for following disease progress of patients over time. A possible way to detect changes is by analysing the deformation field found by the nonrigid registration, as is done by Rey *et al* (2002). Another way to evaluate this progress is to visually inspect the difference between a first scan, taken at time t_0 , and a registered

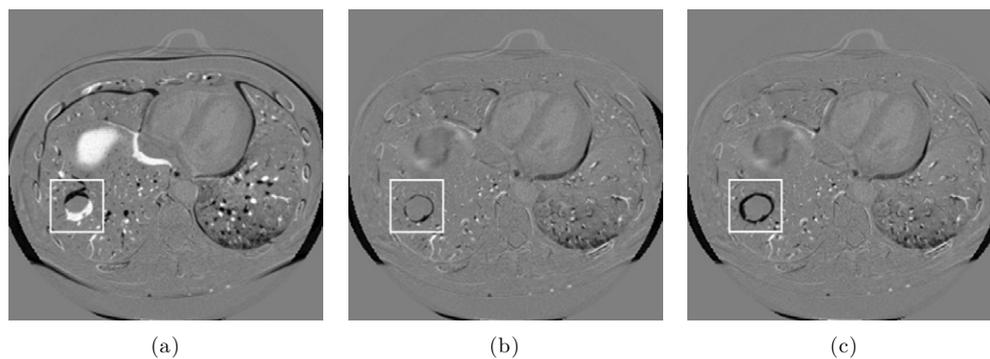


Figure 3. BS registration compared with BSF for the CT thorax follow-up data. Difference images of the registration result with the fixed t_0 image are shown. (a) for affine, (b) for BS and (c) for BSF registration. The white box shows a tumour location.

follow-up scan, taken at t_1 . For the case of patients suffering from lung tumours, there might be tumour growth between the first scan and the follow-up. Because of nonrigid motion occurring between scans, a rigid registration is not sufficient to achieve good alignment of the anatomy. However, standard nonrigid registration methods will minimize the difference between tumours at different time points, effectively concealing tumour growth (see figure 3). Therefore, the tumours should not be allowed to deform by the nonrigid registration algorithm. It is assumed that the tumour volume does not change between scans due to differences in the lung inspiration levels, because the protocol for acquiring the data was such that inspiration levels are approximately equal. In this section, BS registration of CT thorax follow-up scans of patients suffering from lung tumours is compared with BSF registration. Rigidity is evaluated with tumour volume measurements.

3.2.1. Data description. Registration was performed on CT follow-up data sets of the thorax of five patients having lung tumours. For each patient two images were taken at different time points. The data were acquired with a 16-slice spiral CT scanner (Mx8000 IDT 16, Philips Medical Systems, Best, The Netherlands). The images have an in-plane resolution of 512×512 pixels. The number of slices varies for the data sets, ranging from 400 to 550. The in-plane voxel size is around 0.7×0.7 mm. The slice thickness was always 1.0 mm and slices were reconstructed every 0.7 mm. Before registration, each data set was downsampled with a factor of two in each dimension by discarding odd rows, columns and slices, to reduce computer memory and computational load. The five data sets contain thirty-six tumours in total, with an average volume of 2.8 ± 3.4 ml for the first scan t_0 and 5.9 ± 7.2 ml for the follow-up t_1 . No new tumours had developed at t_1 .

3.2.2. Experiment setup. The CT image taken at time t_0 is set to be the fixed image, the CT image taken at time t_1 the moving image. In order to get a coarse alignment between fixed and moving images an affine registration was performed first. For both the affine and the nonrigid registration three resolutions were employed. For the nonrigid registration the B-spline grid spacing at the final level was 16 voxels. In the three resolutions 100, 100 and 300 iterations were used for the optimizer, respectively. In every iteration 5000 samples were selected to calculate (the derivative of) the mutual information, see section 2.2. The parameter a that defines the step size for the optimization algorithm was chosen 100 000.0, 70 000.0 and

Table 1. Average lung overlap. Between brackets the parameters (k, m) are given.

Before registration	Affine	BS	BSF (10, 1)	BSF (1, 1)	BSF (1, 10)
0.64 ± 0.22	0.92 ± 0.06	0.97 ± 0.02	0.97 ± 0.02	0.97 ± 0.02	0.97 ± 0.02

50 000.0 for the three resolutions. For BSF, a crude manual segmentation of the tumours was used to define $c(\mathbf{x})$, setting $c(\mathbf{x})$ to 1.0 for voxels within the tumour and to 0.0 elsewhere. The BSF method was tested with varying strength of the adaptive filtering, with $(k, m) = (1, 10)$, $(1, 1)$ and $(10, 1)$.

3.2.3. Results. The accuracy of the registration is measured by calculating the lung overlap of the registered image and the fixed image. For this purpose, automatic lung segmentations were made with an algorithm based on the method by Hu *et al* (2001), described in detail in Sluimer *et al* (2005). The segmentation does not include the lung vasculature and the tumours. The overlap measure is defined as

$$\text{overlap} \triangleq \frac{2|L_1 \cap L_2|}{|L_1| + |L_2|}, \quad (7)$$

where L_i is the set of all voxels within the lung and $|L_i|$ denotes the size of set L_i . The average and standard deviation values of lung overlap were calculated over all data, and are reported in table 1. The results show that a good lung overlap was achieved with both nonrigid registration algorithms. This is confirmed by visual inspection of the results (see figure 3 for an example).

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

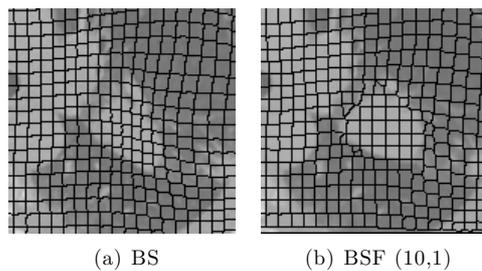


Figure 4. Deformation fields for the CT thorax follow-up data around a tumour, following from the (a) BS and the (b) BSF registration. Where BS compresses the tumour ($r_{BS} = 0.69$), BSF retains rigidity ($r_{BSF} = 0.98$).

Table 2. Geometric mean tumour volume ratios. Between brackets the parameters (k, m) are given.

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

of similarity. Especially when there are large tumour volume differences between t_0 and t_1 (the fourth group in table 2), volume preservation is harder to obtain. Examples of the deformation fields after BS and BSF registration are given in figure 4. Unlike BS, the deformation field for BSF is clearly rigid at the tumour location.

3.3. Digital subtraction angiography

Digital subtraction angiography (DSA) is an established modality for visualizing blood vessels in the human body. During image acquisition patient motion often occurs, due to breathing, heart beat, activity in the intestines or movement of the body. This motion results in artefacts in the subtraction images. In figure 5(a) a subtraction image is shown after rigid registration; the artefacts in the original subtraction image are even larger. In order to correct for patient motion after acquisition, a nonrigid registration is needed. Typically, a sequence of images is taken, where different parts of the vasculature are visible at different times. To see the whole imaged vasculature, all the images from the specific sequence have to be registered to some fixed image. The first image, acquired just before the arrival of the contrast bolus and known as the baseline image, is used as the fixed image. As reported in the literature (Rohlfing *et al* 2003, Tanner *et al* 2000), nonrigid registration of images containing contrast-enhanced structures can lead to significant compression of those structures. Switching the fixed and the moving image and simply using BS does not guarantee that compression is avoided, and requires the inverse of the transformation. Therefore, a nonrigid registration is required that can treat the vasculature different from other tissue to maintain vasculature size.

3.3.1. Data description. The 2D digital x-ray angiography image data were acquired with an Integris V3000 C-arm imaging system (Philips Medical Systems, Best, The Netherlands). In total, 26 image sequences of 12 different patients were obtained. The image sequences were

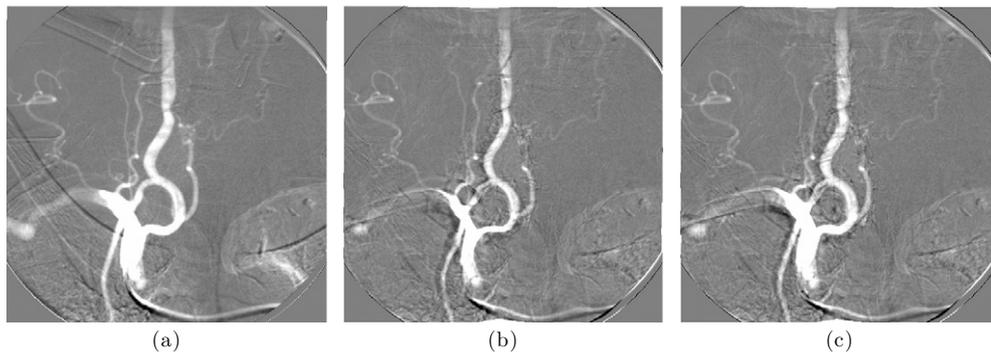


Figure 5. BS registration compared with BSF for the DSA data. Difference images of the registration result with the fixed image are shown. (a) for rigid, (b) for BS and (c) for BSF (2, 1) registration. Note the substantial change in vessel radius of the BS method, compared to the other algorithms. Whereas the BS method compresses the vasculature, the BSF method is much more capable of retaining vessel widths.

of size 512×512 pixels for 22 data sets and 1024 by 1024 pixels for 4 data sets; they contain about ten images each. Intensities in the DSA images range approximately from 100 to 950, with an arithmetic mean and standard deviation of 550 ± 180 . Images were taken of different locations in the body: abdomen (10), brain (5), hip and foot (4), heart (1), neck (5) and lungs (1). The first image in each sequence was taken before the arrival of the contrast bolus; in the following images the contrast agent is visible in parts of the vasculature.

3.3.2. Experiment setup. The baseline image was taken to be the fixed image. For our experiments one image from each sequence was registered to the fixed image. The image showing the most vasculature was manually selected as the moving image. To get a coarse alignment between fixed and moving images, a rigid registration was performed prior to nonrigid registration. Two resolutions were used, with a B-spline grid spacing of 32 and 16 pixels at the first and final resolution levels, respectively. For the two resolutions 600 and 300 iterations were used, respectively. At every resolution 5000 samples were used to calculate (the derivative of) the mutual information (see section 2.2). The parameter a that defines the step size for the optimization algorithm was chosen to be 6000.0 and 3500.0 for the two resolutions. For the BSF registration, a crude manual segmentation of the vessels was used to define $c(\mathbf{x})$, setting $c(\mathbf{x})$ to 1.0 for voxels within the vasculature and to 0.0 otherwise. The BSF method was tested with the parameter settings $(k, m) = (1, 5), (1, 1), (2, 1)$ and $(5, 1)$.

3.3.3. Results. The root mean square difference (RMSD) of $I_M(\mathbf{x} + \mathbf{d}(\mathbf{x}))$ with $I_F(\mathbf{x})$ at the background was calculated to see if the nonrigid registration indeed reduces motion artefacts. The background is defined as everything within the cone beam, but outside the manual vessel segmentation. Differences at vessel locations are not taken into account, since these differences are meant to be large. The arithmetic means and standard deviations of the RMSD were calculated for all 26 images. The results are reported in the top row of table 3. Rigid registration reduced the motion artefacts. Both BS and BSF nonrigid registration methods improved substantially on that. In figure 5 the difference image after rigid registration, and after the two nonrigid registration methods (BS and BSF (2, 1)) is shown for an example DSA image. Visual inspection confirms the reduction in motion artefacts, as measured by the RMSD.

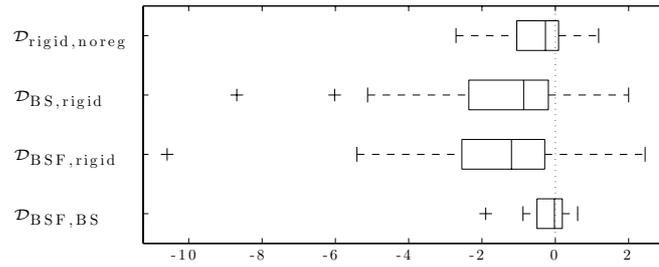


Figure 6. A box-and-whisker plot for the difference \mathcal{D} in background motion artefacts. The box represents the first and third quartiles. The horizontal line in this box is the median. The whiskers indicate the extent of the rest of the data, except for outliers. Outliers are points falling more than 1.5 times the interquartile range past the ends of the box, and are denoted by a dot.

Table 3. Results for DSA data. Arithmetic means and standard deviations for the RMSD are displayed in the top row. Geometric means and standard deviations for the vessel diameter ratios are shown in the bottom row. Between brackets the parameters (k, m) are given.

	No registration	Rigid	BS	BSF (1, 5)	BSF (1, 1)	BSF (2, 1)	BSF (5, 1)
RMSD	14.0 ± 6.0	13.5 ± 5.7	12.1 ± 4.2	11.8 ± 4.0	11.9 ± 4.0	11.9 ± 4.0	11.9 ± 4.0
Diameter	$1.00 \times / 1.00$	$1.00 \times / 1.00$	$0.85 \times / 1.16$	$0.91 \times / 1.11$	$0.96 \times / 1.06$	$0.97 \times / 1.04$	$0.98 \times / 1.03$

In order to compare the reduction in motion artefacts between the different registration methods, the difference of the RMSDs is defined as $\mathcal{D}_{i,j} = \text{RMSD}_i - \text{RMSD}_j$, where $i \neq j$ are the different registration methods. This difference was calculated for several combinations of registration methods and for all DSA data, resulting in the differences $\mathcal{D}_{\text{rigid,noreg}}$, $\mathcal{D}_{\text{BS,rigid}}$, $\mathcal{D}_{\text{BSF,rigid}}$ and $\mathcal{D}_{\text{BSF,BS}}$, where BSF is BSF (2, 1). In figure 6 a box-and-whisker plot for these differences in motion artefacts is shown. Rigid registration reduced the motion artefacts slightly, compared to no registration (top row). The second and third rows show that nonrigid registration indeed improved substantially on rigid registration (rigid versus BS and rigid versus BSF both have $p < 0.05$ in a two-tailed paired t -test). The bottom row indicates that the RMSD for the BS and BSF method are similar.

The rigidity of the vasculature is evaluated by manually measuring the vessel diameter vd at several locations. Six locations were selected for each of the 26 images, yielding a total of 156 diameter measurements. As in the previous section, the ratios r are used to evaluate vessel compression. The vessel diameter after registration is compared to the diameter before registration: $r_{\text{reg}} = vd_{\text{reg}}/vd_{\text{noreg}}$. The geometric mean of the vessel diameter ratios is reported in the bottom row of table 3. Standard nonrigid registration (BS) clearly resulted in the compression of the vessels, whereas the tissue-dependent filtering (BSF (2, 1)) retained the vessel widths. When the adaptive filtering strength is decreased (BSF (1, 1) and BSF (1, 5)) rigidity-preservation is decreased, and vice versa (BSF (5, 1)).

4. Discussion and conclusions

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

this work, and incorporated in a nonrigid registration framework: tissue-dependent nonrigid registration.

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

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

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

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

Acknowledgments

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

References

- Arsigny V, Pennec X and Ayache N 2005 Polyrigid and polyaffine transformations: a novel geometrical tool to deal with non-rigid deformations application to the registration of histological slices *Med. Image Anal.* **9** 507–23
- Brock K K, Sharpe M B, Dawson L A, Kim S M and Jaffray D A 2005 Accuracy of finite element model-based multi-organ deformable image registration *Med. Phys.* **32** 1647–59
- Cachier P, Pennec X and Ayache N 1999 Fast non-rigid matching by gradient descent: study and improvements of the demons algorithm *Research Report 3706*, INRIA
- Edwards P J, Hill D L G, Little J A and Hawkes D J 1998 A three-component deformation model for image-guided surgery *Med. Image Anal.* **2** 355–67
- Hill D L G, Batchelor P G, Holden M and Hawkes D J 2001 Medical image registration *Phys. Med. Biol.* **46** R1–45
- Hu S, Hoffman E A and Reinhardt J M 2001 Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images *IEEE Trans. Med. Imaging* **20** 490–8
- Ibáñez L, Schroeder W, Ng L and Cates J 2005 *The ITK Software Guide* 2nd edn (New York: Kitware, Inc.) (ISBN 1-930934-15-7)

- Klein S, Staring M and Pluim J P W 2007 Evaluation of optimisation methods for nonrigid medical image registration using mutual information and B-splines *IEEE Trans. Image Process.* **16** at press
- Lester H and Arridge S R 1999 A survey of hierarchical non-linear medical image registration *Pattern Recogn.* **32** 129–49
- Lester H, Arridge S R, Jansons K M, Lemieux L, Hajnal J V and Oatridge A 1999 Non-linear registration with the variable viscosity fluid algorithm *IPMI (Lect. Notes Comput. Sci.* vol 1613) (Berlin: Springer) pp 238–51
- Little J A, G. Hill D L and Hawkes D J 1997 Deformations incorporating rigid structures *Comput. Vis. Image Understand.* **66** 223–32
- Loeckx D, Maes F, Vandermeulen D and Suetens P 2004 Nonrigid image registration using free-form deformations with a local rigidity constraint *MICCAI (Lect. Notes Comput. Sci.* vol 3216) pp 639–46
- Maintz J B A and Viergever M A 1998 A survey of medical image registration *Med. Image Anal.* **2** 1–36
- Mattes D, Haynor D R, Vesselle H, Lewellen T K and Eubank W 2003 PET-CT image registration in the chest using free-form deformations *IEEE Trans. Med. Imaging* **22** 120–8
- Rey D, Subsol G, Delingette H and Ayache N 2002 Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis *Med. Image Anal.* **6** 163–79
- Rohlfing T and Maurer C R Jr 2001 Intensity-based nonrigid registration using adaptive multilevel free-form deformation with an incompressibility constraint *MICCAI (Lect. Notes Comput. Sci.* vol 2208) pp 111–9
- Rohlfing T, Maurer C R Jr, Bluemke D A and Jacobs M A 2003 Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint *IEEE Trans. Med. Imaging* **22** 730–41
- Rueckert D, Sonoda L I, Hayes C, Hill D L G, Leach M O and Hawkes D J 1999 Nonrigid registration using free-form deformations: application to breast MR images *IEEE Trans. Med. Imaging* **18** 712–21
- Sluimer I C, Prokop M and van Ginneken B 2005 Towards automated segmentation of the pathological lung in CT *IEEE Trans. Med. Imaging* **24** 1025–38
- Spall J C 1998 Implementation of the simultaneous perturbation method for stochastic optimization *IEEE Trans. Aerosp. Electron. Syst.* **34** 817–23
- Staring M, Klein S and Pluim J P W 2005 Nonrigid registration with adaptive, content-based filtering of the deformation field *SPIE Medical Imaging: Image Processing (Proceedings of SPIE* vol 5747) (Bellingham, WA: SPIE Optical Engineering Press) pp 212–21
- Stefanescu R, Pennec X and Ayache N 2004 Grid powered nonlinear image registration with locally adaptive regularization *Med. Image Anal.* **8** 325–42
- Tanner C, Schnabel J A, Chung D, Clarkson M J, Rueckert D, Hill D L G and Hawkes D J 2000 Volume and shape preservation of enhancing lesions when applying nonrigid registration to a time series of contrast enhancing MR breast images *MICCAI (Lect. Notes Comput. Sci.* vol 1935) pp 327–37
- Thévenaz P and Unser M 2000 Optimization of mutual information for multiresolution image registration *IEEE Trans. Image Process.* **9** 2083–99
- Thirion J P 1998 Image matching as a diffusion process: an analogy with Maxwell's demons *Med. Image Anal.* **2** 243–60
- Unser M 1999 Splines: a perfect fit for signal and image processing *IEEE Signal Process. Mag.* **16** 22–38