Medical Image Analysis 15 (2011) 238–249

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Nonrigid registration of dynamic medical imaging data using nD + t B-splines and a groupwise optimization approach

C.T. Metz [a,*], S. Klein [a], M. Schaap [a], T. van Walsum [a], W.J. Niessen [a,b]

[a] Departments of Medical Informatics and Radiology, Erasmus MC – University Medical Center Rotterdam, P.O. Box 2040, 3000 CA, Rotterdam, The Netherlands
[b] Faculty of Applied Sciences, Delft University of Technology, Lorentzweg 1, 2628 CJ, Delft, The Netherlands

ARTICLE INFO

Article history:
Received 12 February 2010
Received in revised form 12 October 2010
Accepted 18 October 2010
Available online 28 October 2010

Keywords:
Dynamic
Nonrigid registration
Motion estimation
nD + t
4D

ABSTRACT

A registration method for motion estimation in dynamic medical imaging data is proposed. Registration is performed directly on the dynamic image, thus avoiding a bias towards a specifically chosen reference time point. Both spatial and temporal smoothness of the transformations are taken into account. Optionally, cyclic motion can be imposed, which can be useful for visualization (viewing the segmentation sequentially) or model building purposes. The method is based on a 3D (2D + time) or 4D (3D + time) free-form B-spline deformation model, a similarity metric that minimizes the intensity variances over time and constrained optimization using a stochastic gradient descent method with adaptive step size estimation. The method was quantitatively compared with existing registration techniques on synthetic data and 3D + t computed tomography data of the lungs. This showed subvoxel accuracy while delivering smooth transformations, and high consistency of the registration results. Furthermore, the accuracy of semi-automatic derivation of left ventricular volume curves from 3D + t computed tomography angiography data of the heart was evaluated. On average, the deviation from the curves derived from the manual annotations was approximately 3%. The potential of the method for other imaging modalities was shown on 2D + t ultrasound and 2D + t magnetic resonance images. The software is publicly available as an extension to the registration package `elastix`.

1. Introduction

1.1. Background

Dynamic imaging data are increasingly available due to ongoing advancements in medical imaging techniques (Li et al., 2008a). Motion estimation of the anatomy of interest from these images is often desirable, e.g. to quantify motion-related markers of disease, to construct motion and deformation models for therapeutic or surgical planning and guidance, or to remove motion to allow the analysis of intensity features at corresponding anatomical locations over time. Examples of motion quantification are the measurement of the distensibility of blood vessels or aneurysms (e.g. Li et al., 2008b; Ganten et al., 2008), the quantification of lung function (e.g. Reinhardt et al., 2008; Boldea et al., 2008) and the quantification of left ventricular function of the heart (e.g. Mahnken et al., 2009). The application of motion and deformation models in image-guided interventions was, for example, discussed by Hawkes et al. (2005). Motion removal has been applied, for example, in the analysis of perfusion CT or perfusion MRI images (e.g. Xue et al., 2008; Milles et al., 2008).

Manual motion estimation from a time series of images is a tedious task. Corresponding landmark positions in time need to be determined and depending on the application of interest the number of required landmarks may be very large. Image registration methods are often applied to automate this process. In these methods, the correspondence between the anatomy at different time points is found by minimizing a landmark based, segmentation based or intensity based similarity measure (Maintz and Viergever, 1998; Hill et al., 2001). These registration procedures must be sufficiently robust to handle the challenges inherent to dynamic imaging, such as fast moving anatomy, motion artifacts (Li et al., 2008a), and varying contrast-to-noise ratio over time, e.g. due to the application of dose reduction techniques such as ECG-derived pulsing windows in CT coronary angiography (Weustink et al., 2009).

1.2. Previous work on motion estimation

In this work we focus on intensity based registration approaches, which work directly on the input images without the need for preprocessing techniques to extract features from the images. There is a vast amount of work on the application of intensity-based image registration techniques for motion estimation. Next to these approaches, groupwise registration techniques have

* Corresponding author. Tel.: + 31 107044078; fax: + 31 107044722.
E-mail address: c.metz@erasmusmc.nl (C.T. Metz).

1361-8415/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2010.10.003

been proposed for the simultaneous alignment of multiple images from different patients (e.g. for atlas building), which is closely related to the alignment of different time point images without taking the temporal continuity of the data into account. We distinguish the existing techniques by the basic components of a registration approach: the transformation model, cost function and optimization strategy. Details about these categorizations are outlined below. An overview is given in Table 1. Besides these three categories, the table also reports the support for constraints on cyclic motion.

Several models can be used to describe the transformation that aligns the images. We discriminate between methods using an Eulerian approach, in which all deformations are described with respect to the neighboring time point, and methods using a Lagrangian approach, in which deformations are described with respect to a chosen reference frame. In the latter approach the reference frame is often chosen to be directly related to one of the time points of the input image, but sometimes also defined implicitly, e.g. as the mean of the population. Most existing methods use a Lagrangian transformation model which can either take or not take into account the temporal smoothness of the deformations (respectively denoted with nD + t and nD in Table 1; see also Fig. 1 for an illustration). However, the majority of these methods only force the deformations to be smooth, viz. continuous and differentiable, in the spatial domain. Note that nD + t transformation models have not only been applied in motion estimation methods presented in Table 1, but also in inter-patient and intra-patient alignment of dynamic imaging sequences (Peyrat et al., 2010; Perperidis et al., 2005; Lopez et al., 2008; Schreibmann et al., 2008).

The cost function, or dissimilarity metric, computes the dissimilarity between the images to measure the quality of the current transformation estimate. We distinguish three different approaches, often related to the chosen transformation model. The first is a consecutive approach in which the similarity is determined between the images of consecutive time points. The second is a reference approach in which the similarity is determined between the image to be registered and a chosen reference image. The last one is a global approach, in which the imaging data of all time points are taken into account in the computation of the cost function. A disadvantage of the first two approaches is that a limited amount of available image

Table 1
Existing dynamic and groupwise registration approaches. For every method, the transformation model, cost function, optimization strategy and inclusion of constraints on cyclic motion are listed. More than one check mark for a certain method in a certain category means that both approaches are applied in the same work. The different categories are explained in the introduction (Section 1.2). Methods are sorted on transformation model, cost function and optimization strategy subsequently. Cyclic motion can either be implemented in the cost function (c) or in the transformation model (t). The last three methods are used to compare the proposed method with (Section 3.1).

	Transformation model			Cost function			Optimization		Cyclic
Method	Eulerian	Lagrangian (nD)	Lagrangian (nD + t)	Consecutive	Reference	Global	Pairwise	Global	
Boldea et al. (2008)	√	√		√	√		√		
Reinhardt et al. (2008)	√	√		√	√		√		
De Craene et al. (2009)	√					√		√	
Bidaut and VallTe (2001)		√			√		√		
Kaus et al. (2004)		√			√		√		
Lorenzo-Valdés et al. (2002)		√			√		√		
Marsland et al. (2003)		√			√		√		
Marsland et al. (2008)		√			√		√		
Rao et al. (2002)		√			√		√		
Rietzel and Chen (2006)		√			√		√		
Wierzbicki et al. (2004)		√			√		√		
Joshi et al. (2004)		√				√	√		
Balci et al. (2007a,b)		√				√		√	
Bhatia et al. (2004)		√				√		√	
Miller et al. (2000)		√				√		√	
Zöllei et al. (2005)		√				√		√	
Sundar et al. (2009)			√		√			√	√ (c)
Castillo et al. (2010)			√		√			√	
Proposed method			√			√		√	√ (t)
Reference time point		√			√		√		
Consecutive time points	√			√			√		
Groupwise method		√				√		√	

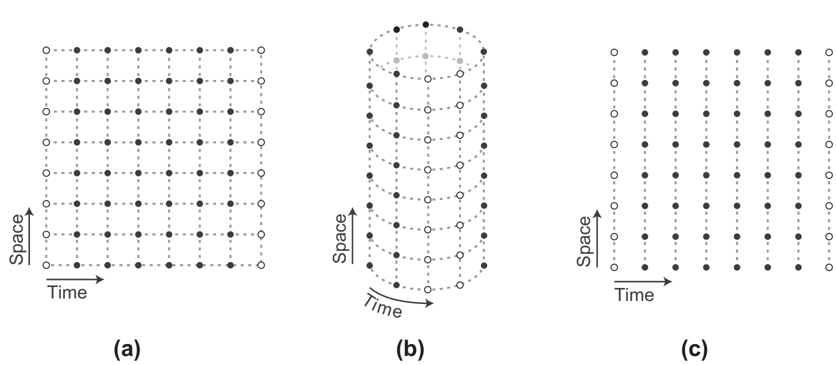

Fig. 1. (a) nD + t B-spline grid, (b) cyclic nD + t B-spline grid and (c) nD B-spline grid used in reference time point, consecutive time point and groupwise registration approaches (see Section 3.1). In the cyclic version (b), the grid points at the temporal border (open nodes) are direct neighbors.

information is used during the registration procedure. The individual registrations only exploit the information present in the reference image and the image to be registered, whereas the other images may also contain valuable information. Moreover, by choosing a single reference time point, the registration result can be biased towards this image. In the global cost functions all image information is taken into account simultaneously, potentially leading to more robust and consistent registration results, without a bias towards a certain reference image.

Finally, we distinguish two kinds of optimization approaches for finding the optimal transformation. While the first approach optimizes the cost function for every time point separately, the second approach performs this optimization for all time points simultaneously, which we call a global approach. The optimization approach used is often related to the chosen cost function and transformation model. For global cost functions, a global optimization approach is needed. The same holds for the Lagrangian *n*D + t transformation model that takes temporal smoothness into account. When a consecutive or reference cost function is applied, the optimization is most often performed in a pairwise manner.

In certain cases it is known a priori that the anatomical motion has a cyclic nature. When this knowledge is taken into account during the acquisition procedure, e.g. by ECG-gating or respiratory gating, one might want to incorporate this into the registration procedure. Two different approaches can be distinguished. In the first approach, a term is added to the cost function to penalize non-cyclic transformations. In the other approach, cyclic motion is enforced by adapting the transformation model. To the best of our knowledge, only the first approach has been used in previous work.

1.3. Proposed method

In this work, we focus on the estimation of anatomical motion from dynamic medical imaging data. For this, we assume that physiologically motion is smooth (continuously differentiable) over time. Finding this smooth motion is, for example, useful for the construction of statistical motion models (Metz et al., 2010) or motion visualization. The amount of smoothness depends on the expected motion of the anatomy of interest and the expected distortion of the motion due to pathology. Whereas the motion of the anatomy is expected to be smooth, the *appearance* of the moving anatomy in the reconstructed image may be non-smooth because of imaging artifacts. In our registration approach, we use a Lagrangian *n*D + t transformation model parametrized by B-splines. The search space for the transformation that minimizes the dissimilarity metric is thereby reduced to those transformations that are both spatially and temporally smooth.

With respect to the cost function, we choose the global approach to eliminate a bias towards a chosen reference frame and use as much image information as possible. The use of a global cost function automatically leads to the choice for a global optimization routine.

A Lagrangian *n*D + t transformation model, a global cost function, and a global optimization routine have previously been addressed in literature for motion estimation in 4D medical imaging data (see Table 1), but never jointly in one framework. We additionally propose a cyclic version of the B-spline transformation model and investigate its influence on the registration results.

The method is evaluated quantitatively on a 2D + t synthetic image, 3D + t computed tomography (CT) images of the lungs and 3D + t computed tomography angiography (CTA) images of the heart. Further examples are presented on 2D + t ultrasound (US) images of the carotid artery and 2D + t magnetic resonance (MR) images of the lungs. To summarize, the main contributions of this work are:

- The development and evaluation of a registration method for motion estimation combining a Lagrangian *n*D + t B-spline transformation model, a global cost function and global optimization strategy.
- The possibility to include a cyclic motion constraint that is strictly enforced by the transformation model.
- The quantitative comparison of the proposed technique with three well-known techniques.

Furthermore, the software developed for this publication is publicly available.

2. Method

The proposed method is based on a 3D (2D + time) or 4D (3D + time) free-form B-spline deformation model, incorporating both the spatial and time dimensions. It aims to minimize the image intensity changes over time. An implicit reference frame is used to eliminate the need to choose a reference time point image. The following subsections describe the different components of the approach.

2.1. Transformation

A B-spline transformation model is used (Rueckert et al., 1999) because the compact support of B-splines keeps the running time reasonably low for higher dimensional imaging data. We restrict the deformations to only take place in the spatial domain and thereby search for those deformations that spatially align the different time point images. The deformation is regularized by assuming smoothness of the deformation in both the spatial and temporal direction of the data.

The D-dimensional input image is denoted with $I(\boldsymbol{y})$, where $y = (x^T, t)^T \in \mathbb{R}^s \times \mathbb{R}$ denotes a coordinate in I which consists of a spatial location $x \in \mathbb{R}^s$ and temporal location $t \in \mathbb{R}$. $D = s + 1$ equals the dimension of the spatiotemporal image data. The B-spline based coordinate transformation $\boldsymbol{T}_{\boldsymbol{\mu}}$ is defined as follows:

$$\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{y}) = \boldsymbol{y} + \sum_{\boldsymbol{y}_k \in \mathcal{N}_y} \boldsymbol{p}_k \beta^r(\boldsymbol{y} - \boldsymbol{y}_k) \tag{1}$$

with $\boldsymbol{y}_k$ the control points, $\beta^r(\boldsymbol{y})$ the rth order multidimensional B-spline polynomial (Unser, 1999), $\boldsymbol{p}_k$ the B-spline coefficient vectors, and $\mathcal{N}_y$ the set of all control points within the compact support of the B-spline at $\boldsymbol{y}$. The control points $\boldsymbol{y}_k$ are defined on a D-dimensional regular grid, overlaid on the image. The parameter vector $\boldsymbol{\mu}$ consists of the collection of the first $D - 1$ elements of each $\boldsymbol{p}_k$. The last element of every $\boldsymbol{p}_k$ is fixed to zero making sure that only deformations in the spatial domain are allowed. This is in contrast to the work of Perperidis et al. (2005) and Peyrat et al. (2010), where deformations in the temporal directions are allowed.

Optionally, cyclic motion can be enforced by letting the B-spline polynomials wrap around in the temporal direction (see Fig. 1a and b). This is achieved by adapting the definition of the control point neigbourhood $\mathcal{N}_y$. A prerequisite for cyclic motion is that the number of time points of the image should be a multiple of the temporal B-spline control point spacing.

In the remainder of the paper the notation $\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{y})$ is interchanged with $\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x}, t)$ for convenience of notation.

2.2. Dissimilarity metric

Because we are working with monomodal dynamic imaging data, our method is based on the assumption that after correct registration the intensity values at corresponding spatial locations over time are equal. This can effectively be measured by computing the variance of intensity values at corresponding spatial locations over time (Bhatia et al., 2007). The dissimilarity metric, or cost function, is therefore defined as:

$$C(\boldsymbol{\mu}) = \frac{1}{|\mathcal{S}||\mathcal{T}|} \sum_{\boldsymbol{x}\in\mathcal{S}} \sum_{t\in\mathcal{T}} \left(I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t)) - \bar{I}_{\boldsymbol{\mu}}(\boldsymbol{x})\right)^2 \tag{2}$$

with $\bar{I}_{\boldsymbol{\mu}}(x)$ the average intensity value over time after applying transformation $\boldsymbol{T}_{\mu}$:

$$\bar{I}_{\boldsymbol{\mu}}(x) = \frac{1}{|\mathcal{T}|} \sum_{t\in\mathcal{T}} I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t)) \tag{3}$$

and $\mathcal{S}$ and $\mathcal{T}$ the set of spatial and temporal voxel coordinates respectively.

2.3. Zero average displacement constraint

The registration is performed directly on the D-dimensional input image, and does not require a reference image. This results in an underconstrained optimization problem, because multiple solutions exist for the minimization of the dissimilarity metric (Eq. (2)). A translation of the image volume will, for example, not change the metric value. We therefore constrain the average deformation in time to be the identity transform, like (Bhatia et al., 2004 and Balci et al., 2007a) did for groupwise registration:

$$\frac{1}{|\mathcal{T}|} \sum_{t\in\mathcal{T}} \boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t) = \boldsymbol{x}. \tag{4}$$

The next subsection explains how this constraint is enforced in the optimization procedure.

2.4. Optimization

For the final solution we need to determine those transform parameters that minimize the dissimilarity metric:

$$\hat{\boldsymbol{\mu}} = \arg\min_{\boldsymbol{\mu}} C(\boldsymbol{\mu}) \text{ subject to } (4). \tag{5}$$

Hereto, we use an adaptive stochastic gradient descent optimizer (ASGD) (Klein et al., 2009). The main advantage of this optimizer compared with conventional gradient-based optimizers is that it applies random sampling of the data in the computation of the derivatives, which causes a significant reduction in computation time. This sampling strategy is applied to select the voxel locations in $\mathcal{S}$ and the temporal indices in $\mathcal{T}$. Note that new samples are drawn at each iteration of the optimization.

The ASGD optimizer requires that the derivative of the cost function with respect to $\boldsymbol{\mu}$ is known, which follows from differentiating Eq. (2):

$$\frac{\partial C}{\partial \boldsymbol{\mu}} = \frac{2}{|\mathcal{S}||\mathcal{T}|} \sum_{\boldsymbol{x}\in\mathcal{S}} \sum_{t\in\mathcal{T}} (I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t)) - \bar{I}_{\boldsymbol{\mu}}(\boldsymbol{x})) \cdot \left(\frac{\partial I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t))}{\partial \boldsymbol{\mu}} - \frac{\partial \bar{I}_{\boldsymbol{\mu}}(\boldsymbol{x})}{\partial \boldsymbol{\mu}}\right) \tag{6}$$

$$= \frac{2}{|\mathcal{S}||\mathcal{T}|} \sum_{\boldsymbol{x}\in\mathcal{S}} \left[\sum_{t\in\mathcal{T}} (I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t)) - \bar{I}_{\boldsymbol{\mu}}(\boldsymbol{x})) \frac{\partial I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t))}{\partial \boldsymbol{\mu}} - \frac{\partial \bar{I}_{\boldsymbol{\mu}}(\boldsymbol{x})}{\partial \boldsymbol{\mu}} \sum_{t\in\mathcal{T}} (I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t)) - \bar{I}_{\boldsymbol{\mu}}(\boldsymbol{x}))\right] \tag{7}$$

Substituting Eq. (3) in the last term of (7) results in:

$$\frac{\partial C}{\partial \boldsymbol{\mu}} = \frac{2}{|\mathcal{S}||\mathcal{T}|} \sum_{\boldsymbol{x}\in\mathcal{S}} \sum_{t\in\mathcal{T}} (I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t)) - \bar{I}_{\boldsymbol{\mu}}(\boldsymbol{x})) \frac{\partial I(\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x},t))}{\partial \boldsymbol{\mu}} \tag{8}$$

To apply the constraint that the average deformation over the time series is zero (see Section 2.3), we follow the approach of (Balci et al., 2007a): we subtract the mean from each derivative vector, causing the sum of B-spline coefficients to be zero. We therefore use for every element i of $\partial C(\mu)/\partial\boldsymbol{\mu}$ the following equation to determine the constrained update:

$$\frac{\partial C'}{\partial \boldsymbol{\mu}_i} = \frac{\partial C}{\partial \boldsymbol{\mu}_i} - \frac{1}{|Q_i|} \sum_{q\in Q_i} \frac{\partial C}{\partial \boldsymbol{\mu}_q} \tag{9}$$

where Q_i denotes the collection of all elements of $\boldsymbol{\mu}$ over time that correspond to the same spatial grid point location and direction as element i.

2.5. Inverse transformation

The zero average displacement constraint described in Section 2.3 implicitly defines a reference frame that lies in the center of the dynamics described by the image. After registration all time point images are aligned in this reference frame. Depending on the type of application the registration procedure is used for, it might be useful to know transformation $\boldsymbol{T}^{ij}_{\boldsymbol{\mu}}$ which maps coordinates from time point i to time point j. To be able to define this transformation, the inverse mapping $\boldsymbol{T}^{-1}_{\boldsymbol{\mu}}$, which maps coordinates from the input image coordinate frame to the reference frame, needs to be known.

Because the inverse of a B-spline transformation cannot be derived in closed-form, an additional subsequent optimization procedure is applied, formulated in a way similar to the registration procedure. The inverse transformation $\boldsymbol{T}^{-1}_{\hat{\mu}}$ is derived by searching for a B-spline transformation $\boldsymbol{T}_{\boldsymbol{v}}$ that cancels $\boldsymbol{T}_{\hat{\mu}}$, by minimizing the following cost function:

$$F(\boldsymbol{v}) = \frac{1}{|Y|} \sum_{\boldsymbol{y}\in Y} \|\boldsymbol{T}_{\boldsymbol{v}}(\boldsymbol{T}_{\hat{\mu}}(\boldsymbol{y})) - \boldsymbol{y}\|^2. \tag{10}$$

with Y the set of voxel locations. The result $\boldsymbol{T}_{\hat{v}}$ of this minimization is used as an estimate of $\boldsymbol{T}^{-1}_{\hat{\mu}}$. To make sure an accurate inverse can be estimated one should prevent foldings in the transformations resulting from the forward registration procedure. In this paper we do this by choosing appropriate grid spacings, but one could also consider adding a penalty term which incorporates constraints on the Jacobian of the transformations (Chun et al., 2010; Sdika, 2008). As the inverse of a B-spline transformation cannot be modelled exactly with another B-spline transformation, we choose a smaller grid spacing for the inverse transform than was used for the forward transform that aligns all time point images to yield more accurate results.

The $(D-1)$-dimensional transformation $\boldsymbol{T}^{ij}$ that aligns time point image i with time point image j can now be derived by combining the forward transform at time point j and the inverse transformation at time point i:

$$\boldsymbol{T}^{ij}_{\hat{\mu}}(\boldsymbol{x}) = \left[\boldsymbol{T}_{\hat{\mu}}\left(\left[\boldsymbol{T}^{-1}_{\hat{\mu}}(\boldsymbol{x},t_i)\right]_x, t_j\right)\right]_x \tag{11}$$

where $[\cdot]_x$ selects the $(D-1)$-dimensional part of the D-dimensional transformation $\boldsymbol{T}$.

2.6. Implementation details

Linear interpolation in the spatial domain is used for the derivation of intensity values at non grid-point positions in the images.

A multi-resolution strategy is employed to improve the capture range and robustness of the registration. In the lower resolutions the image is convolved with a Gaussian kernel. The standard deviation of this kernel and the spacing of the B-spline grid in the spatial directions of the image are reduced with a factor two in the next resolution level. This multi-resolution approach is used for both the registration procedure to align the time point images

Table 2
Image dimensions, average voxel sizes and parameter settings used in the experiments. The voxel size in the temporal direction was always set to 1.0 'mm'. The spatial grid spacing for the inverse transformation in the synthetic experiments was chosen to be 1.0 mm smaller than the grid spacing for the forward transformation. The spatial grid spacing for the inverse transformation in the other experiments was the listed grid spacing − 3.0 mm.

	Dimensions	Voxelsize (mm)	Grid spacing (mm)	r	$\|\mathcal{S}\|$	$\|\mathcal{T}\|$	Iterations	Res. levels	Cyclic
Synthetic tube example	64 × 64 × 64	0.5 × 0.5	4 × 4 × 4	3	500	5	2000	2	(√)
2D + t US carotids	512 × 512 × 60	0.06 × 0.06	1 × 1 × 3	3	1000	5	2000	3	
2D + t MR lungs	256 × 256 × 60	1.5 × 1.5	30 × 30 × 2	3	1000	5	2000	3	
3D + t CT heart	256 × 256 × 160 × 20	0.7 × 0.7 × 0.8	15 × 15 × 1	2	2000	5	2000	4	(√)
3D + t CT lungs (POPI)	347 × 274 × 141 × 10	0.98 × 0.98 × 2.0	13 × 13 × 1/2/3	2	2000	5	2000	4	(√)
3D + t CT lungs (DIR-lab)	256 × 256 × 100 × 10	1.1 × 1.1 × 2.5	13 × 13 × 1/2/3	2	2000	5	2000	4	(√)

and the optimization procedure to find the inverse B-spline transformation.

The method has been implemented as an extension to the open source registration package `elastix` (Klein et al., 2010) and is freely available for download.[1]

3. Experiments and results

Three types of experiments were conducted to evaluate the proposed method. First, we applied our approach to a synthetic image to compare the results with existing registration approaches (Section 3.3). Second, the performance of the method was quantitatively evaluated using publicly available 3D + t CT data of the lungs (Section 3.4) and 3D + t CTA data of the heart (Section 3.5). And third, further examples are presented on a 2D + t ultrasound image of the carotid artery and a pediatric 2D + t MR image of the lungs (Section 3.6).

For the experiments on the synthetic image, the cardiac CTA data and the examples on ultrasound and MRI images, parameter settings were empirically determined. For the experiments on the CT data of the lungs, parameter settings were tuned on the publicly available POPI-model (Vandemeulebroucke et al., 2007). Resulting parameter settings and image dimensions are listed in Table 2. Parameter files are online available in the parameter file database of the `elastix` website.[1]

3.1. Registration approaches

In the experiments described in Sections 3.3 and 3.4, the proposed registration method is compared with existing registration approaches. The details of these approaches are outlined in the following paragraphs. The B-spline control point spacing, number of resolution levels, and number of iterations were chosen to be the same as the settings used for the proposed method to make a fair comparison possible.

3.1.1. Reference time point registration method

In the reference time point registration method the individual time point images are independently registered to the image of a chosen reference time point. The method uses a Lagrangian (nD) transformation model, a reference cost function and a pairwise optimization strategy. Furthermore, a mean squared difference metric was used, which is strongly related to the proposed variance metric.

3.1.2. Consecutive time point registration method

The consecutive time point registration method registers all individual time point images to the image of the neighbouring time point. It uses a Eulerian transformation model, a consecutive cost function and a pairwise optimization strategy. When the time point at which the registration is started is not equal to zero, registration is performed in two directions, to minimize the propagation of registration errors. A mean squares metric was used as the cost function.

3.1.3. Groupwise registration method

The groupwise registration approach simultaneously aligns the individual time point images. The method uses a Lagrangian (nD) transformation model, a global cost function and a global optimization strategy. The variance metric (Eq. (2)) is used as a cost function and the zero average displacement constraint (Eq. (4)) is applied. This approach is most similar to the proposed method, but does not impose smoothness of the deformations in the temporal direction of the image nor cyclic motion.

3.2. Evaluation measures

We used two evaluation measures for the synthetic experiment and the experiment on the lung data: the primary measure is the accuracy of the registration results and the secondary measure is the temporal smoothness of the transformations. Whereas smoothness on its own does not reflect the quality of the registration, the relation between registration accuracy and smoothness is relevant, as it can be useful when deciding on the right registration strategy for the application of interest. With equal, or slightly worse accuracy, a smoother result is often preferred.

The accuracy for transformation parameters $\boldsymbol{\mu}$ was defined as the mean target registration error (mTRE) (van de Kraats et al., 2005) between a set of landmark collections $P = \{P_1, P_2, \ldots, P_T\}$ for time points $\{1, \ldots, T\}$ and landmarks transformed from a reference time point r to all time points for which the landmarks are available:

$$\mathrm{mTRE}(\boldsymbol{\mu}) = \frac{1}{T|P_1|}\sum_{t\neq r}\sum_{\boldsymbol{p}_{t,i}\in P_t}\left\|\boldsymbol{T}_{\boldsymbol{\mu}}^{rt}(\boldsymbol{p}_{r,i}) - \boldsymbol{p}_{t,i}\right\|, \tag{12}$$

with $\boldsymbol{p}_{t,i}$ landmark i in time point t and r the reference time point.

The smoothness of transformation $\boldsymbol{T}_{\boldsymbol{\mu}}$ was measured as the irregularity of the landmark trajectories:

$$\mathrm{mIrr}(\boldsymbol{\mu}) = \frac{1}{T|P_1|}\sum_{t}^{T}\sum_{\boldsymbol{p}_{t,i}\in P_t}\left\|\frac{\partial^2 \boldsymbol{T}_{\boldsymbol{\mu}}^{rt}(\boldsymbol{p}_{r,i})}{\partial t^2}\right\|^2, \tag{13}$$

with $\boldsymbol{p}_{t,i}$ landmark i at time point t and r the reference time point. Higher values mean more irregular/less smooth trajectories. We computed the derivatives using finite differences to be able to compute the irregularity for all considered registration procedures.

The standard deviation of TRE and irregularity values were also derived.

3.3. Quantitative evaluation on synthetic data

The 2D + t synthetic example consists of a 64 × 64 × 64 pixel image containing a circle with a Gaussian profile with a standard deviation of three voxels. The circle follows a cosine shaped trajec-

[1] http://elastix.isi.uu.nl.

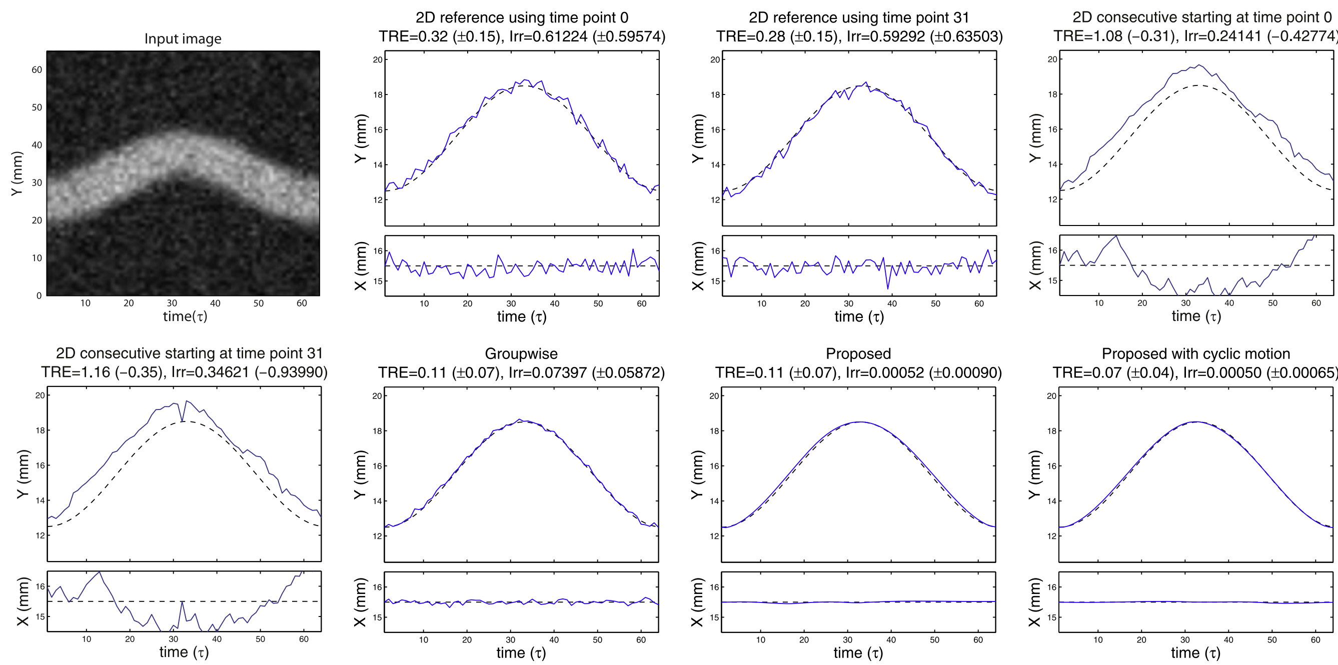

Fig. 2. Registration results for the experiment on a synthetic image. A cross section of the input image is shown in the left top image. The other images show the trajectory resulting from the registration procedure as a solid line. The reference standard is plotted with a dashed line and the title of the plot lists the mean and standard deviation of the accuracy in mm and irregularity in mm/τ^2 of the registration results.

tory over time in the *Y*-direction of the image, i.e. the center of the circle over time is $(x_c, y_c + \alpha\cos(2\pi t/w))$, with $\alpha = 3$, $x_c = y_c = 15.5$ mm and $w = 32$ mm. The cosine is positioned in such a way that the deformation over time is cyclic. The contrast between the tube and the background is 1000 and Gaussian noise was added with a standard deviation of 150 resulting in a contrast-to-noise ratio of $\frac{20}{6}$. A cross section of the resulting 2D + t image is shown in Fig. 2. Seven registration procedures were tested. The first four are the consecutive and reference approach (Section 3.1) using both times 0 and 31 (halfway the time dimension of the data) as reference time point. The fifth is the groupwise registration method (Section 3.1). The last two approaches are the non-cyclic and cyclic version of the proposed registration method. After registration both the accuracy (Eq. (12)) and irregularity (Eq. (13)) of the transformations were determined.

The *X*-displacements and *Y*-displacements found by the different registration procedures are plotted in Fig. 2. The dashed line shows the reference curve. The titles of the plots show both the accuracy and irregularity. It can be noticed that the 2D reference method works reasonably well with respect to the accuracy, but the resulting trajectory is not smooth. The 2D consecutive method is distracted by the image noise and delivers inaccurate results. Among the non-temporally smooth methods the groupwise registration method performs best with respect to both accuracy and smoothness, which may be caused by the more robust global dissimilarity metric. The proposed temporally smooth method has the highest accuracy and delivers the most smooth trajectories of all methods. Accuracy is even slightly improved by imposing cyclic motion.

3.4. Quantitative evaluation on 3D + t CT data of the lungs

Four quantitative experiments were performed on clinical 3D + t CT data of the lungs. The first two experiments assess the accuracy, smoothness, and consistency of the registration results. In the third experiment, the influence of the spatial grid spacing used to obtain the inverse transformation was investigated. In the last experiment, the transitive consistency of the method is evaluated and compared with the reference registration method. The publicly available POPI-model (Vandemeulebroucke et al., 2007) and DIR-lab data were used (Castillo et al., 2009). Both consist of 3D + t CT scans of the lungs and corresponding landmarks in two or more time points of the image. For the DIR-lab data 300 landmarks in time points 0 (inspiration) and 5 (expiration) and 75 landmarks for each time point between time points 0 and 5 were available. For the POPI-model 37 landmarks for all ten time points were available. Parameter settings were tuned on the POPI image and subsequently used for all six images.

3.4.1. Accuracy

Registration accuracy was evaluated by computation of the TRE (Section 3.2) between the reference landmark positions and the landmark positions propagated from time point 0 (DIR-lab data) or time point 1 (POPI-model) to all other time points. The reference time point for these propagations was chosen according to previously published work on this data. The results of the proposed method are compared with the reference time point, consecutive time point and groupwise registration methods described in Section 3.1. For the proposed method, registration was always performed on the complete 4D image and using a lung mask. For the POPI-model the provided mask was used. For the DIR-lab images the masks were created by thesholding, connected component analysis and morphological closing using a spherical structuring element with a diameter of nine voxels. During the registration procedure the sample locations $\boldsymbol{x}$ are drawn from a dilated version of these masks (kernel radius of 13 voxels). Moreover, $\boldsymbol{T}_{\mu}(\boldsymbol{y})$ should lie within the non-dilated mask to be taken into account for the computation of the variance metric. Registration was performed for grid point spacings of one, two and three time points in the time-dimension and with and without applying the cyclic motion constraint to quantify the effect of the temporal and cyclic smoothness on the results.

Results for the 300 two-time-point landmarks of the DIR-lab data are presented in Table 3. Results for the 75 landmarks in six

Table 3
Average and standard deviation of target registration errors in mm for the 300 landmarks in two-time points of the DIR-lab images using the proposed method with three different temporal control point spacings and with and without the cyclic motion constraint. Results are compared with the reference time point, consecutive time point, and groupwise registration approaches. The best previously published results are listed in the last row of the table.

	DIR-lab case 1	DIR-lab case 2	DIR-lab case 3	DIR-lab case 4	DIR-lab case 5
Initial	3.89 (2.78)	4.34 (3.90)	6.94 (4.05)	9.83 (4.85)	7.48 (5.50)
Temp. spacing of three time points (non-cyclic)	1.09 (0.53)	1.12 (0.61)	1.29 (0.70)	1.86 (1.34)	2.06 (1.97)
Temp. spacing of two time points (non-cyclic)	1.05 (0.50)	1.09 (0.60)	1.28 (0.82)	1.73 (1.34)	1.82 (1.56)
Temp. spacing of one time point (non-cyclic)	1.02 (0.47)	1.06 (0.55)	1.21 (0.68)	1.57 (1.20)	1.70 (1.48)
Temp. spacing of three time points (cyclic)	1.05 (0.49)	1.30 (0.83)	1.57 (1.01)	2.52 (2.23)	2.49 (2.22)
Temp. spacing of two time points (cyclic)	1.04 (0.48)	1.16 (0.70)	1.35 (0.77)	1.86 (1.44)	2.12 (1.85)
Temp. spacing of one time point (cyclic)	1.02 (0.50)	1.06 (0.56)	1.19 (0.66)	1.57 (1.20)	1.73 (1.49)
Groupwise	1.02 (0.49)	1.07 (0.56)	1.22 (0.68)	1.56 (1.19)	1.74 (1.47)
3D reference time point	0.99 (0.48)	0.96 (0.49)	1.11 (0.62)	1.49 (1.08)	1.37 (1.21)
3D consecutive time points	1.15 (0.60)	1.06 (0.62)	1.27 (0.68)	1.55 (1.17)	1.82 (1.59)
Castillo et al. (2010)	0.97 (1.02)	0.86 (1.08)	1.01 (1.17)	1.40 (1.57)	1.67 (1.79)

Table 4
Average and standard deviation of target registration errors in mm for the 75 landmarks in six time points of the DIR-lab images and for the 37 landmarks in ten time points of the POPI-model using three different temporal control point spacings and with and without the cyclic motion constraint. Results are compared with the reference time point, consecutive time point, and groupwise registration approaches. The best previously published results are listed in the last row of the table.

	POPI	DIR-lab case 1	DIR-lab case 2	DIR-lab case 3	DIR-lab case 4	DIR-lab case 5
Initial	3.68 (2.97)	2.18 (2.54)	3.78 (3.69)	5.05 (3.81)	6.69 (4.72)	5.22 (4.61)
Temp. spacing of three time points (non-cyclic)	1.13 (0.59)	1.12 (0.75)	1.09 (0.73)	1.24 (0.69)	1.65 (1.12)	1.89 (1.74)
Temp. spacing of two time points (non-cyclic)	1.03 (0.56)	1.02 (0.70)	1.05 (0.70)	1.20 (0.68)	1.47 (1.03)	1.68 (1.49)
Temp. spacing of one time point (non-cyclic)	1.02 (0.58)	0.95 (0.66)	1.00 (0.62)	1.15 (0.61)	1.39 (1.02)	1.50 (1.32)
Temp. spacing of three time points (cyclic)	1.22 (0.79)	0.97 (0.71)	1.54 (1.30)	1.73 (1.28)	2.82 (2.29)	2.33 (1.99)
Temp. spacing of two time points (cyclic)	1.07 (0.58)	0.96 (0.71)	1.11 (0.77)	1.22 (0.71)	1.62 (1.15)	1.84 (1.64)
Temp. spacing of one time point (cyclic)	1.02 (0.58)	0.95 (0.66)	1.00 (0.62)	1.14 (0.61)	1.40 (1.02)	1.50 (1.31)
Groupwise	1.00 (0.56)	0.94 (0.65)	1.01 (0.61)	1.14 (0.63)	1.41 (1.04)	1.49 (1.30)
3D reference time point	0.95 (0.56)	0.93 (0.65)	0.89 (0.51)	1.05 (0.56)	1.40 (1.10)	1.27 (1.10)
3D consecutive time points	1.47 (1.08)	0.97 (0.71)	0.98 (0.61)	1.17 (0.66)	1.37 (0.97)	1.46 (1.40)
Kabus et al. (2009)	0.96 (0.56)	n.a.	n.a.	n.a.	n.a.	n.a.

time points of the DIR-lab data and for the 37 landmarks in 10 time points of the POPI-model are listed in Table 4. For comparison, the initial TRE and the best published results on the POPI-model (Kabus et al., 2009) using 3D B-spline registration and on the DIR-lab data (Castillo et al., 2010) using a four-dimensional optical flow method based on trajectory modeling are included in the tables. Both tables show that the proposed method can achieve subvoxel accuracy, yielding TRE values that are similar to the best published results on the same data. The results for the proposed method without the use of the cyclic motion constraint show that temporal smooth deformations can be achieved by compromising only slightly on registration accuracy. Enabling the cyclic motion constraint helps for DIR-lab case 1, but decreases the accuracy for the other cases.

The average running time on the DIR-lab data is around 40 min for the 3D reference, consecutive and groupwise registration approaches and around 1 h, 1 h and 15 min, and 1 h and 30 min for the proposed method while using a temporal spacing of 3, 2 and 1 time points respectively (AMD Opteron® 2216 2400 MHz). The increase in computation time while using smaller temporal grid spacings is mainly due to the larger size of $\boldsymbol{\mu}$. A visualization of the imaging data before and after registration is shown in Fig. 3.

3.4.2. Smoothness

The registration results of Section 3.4.1 were subsequently used to determine the smoothness of the landmark trajectories by computing their irregularity (Eq. (13)). The results are shown in a bar chart in Fig. 4. Every bar represents one of the registration approaches and every group of bars represents a certain test image. The proposed method results in the most smooth trajectories. It can be seen that increasing the temporal spacing of the B-spline grid improves the temporal smoothness. The inclusion of the constraint on cyclic motion reduces the irregularity even further for larger temporal control point spacings. Furthermore, it can be noticed that the 3D reference registration method performs worst in this sense.

3.4.3. Inverse transformation

An experiment was conducted in which different spatial control point spacings for computation of the inverse transformation were tested. The result of the proposed method using a temporal control point spacing of 2.0 time points and the cyclic motion constraint from the previous sections was used as the forward transformation. Evaluation was performed on the POPI-model. We tested spatial control point spacing for the inverse computation ranging from 13.0 mm (the spacing of the forward transform) to 7.0 mm. Subsequently, we computed the accuracy and irregularity of the results (Section 3.2), and the inverse errors. The multi-resolution strategy was the same as was used for the forward registration procedure. The inverse error was defined as the average magnitude of the transformation vector after subsequent transformation with the forward and inverse transform. Errors were computed for all voxel positions within the mask used in the registration. Results are shown in Fig. 5. The dashed lines show the mean ± the standard deviation of the TRE and irregularity values.

3.4.4. Transitive consistency

The registration results of the reference registration approach (Section 3.1) will depend on the chosen reference image. The choice of different reference images may thus lead to inconsistent results, where we define registrations transitive consistent when for all $i,j \in \mathcal{T}$ and all $\boldsymbol{x} \in \mathcal{S}$:

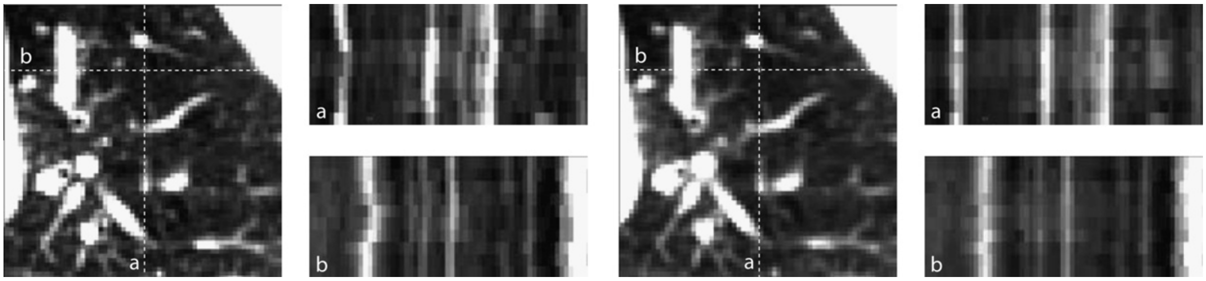

Fig. 3. Registration result for a 3D + t CT image of the lungs. Left: input image, right: registration result. Image (a) and (b) show the dotted lines in the left image over time.

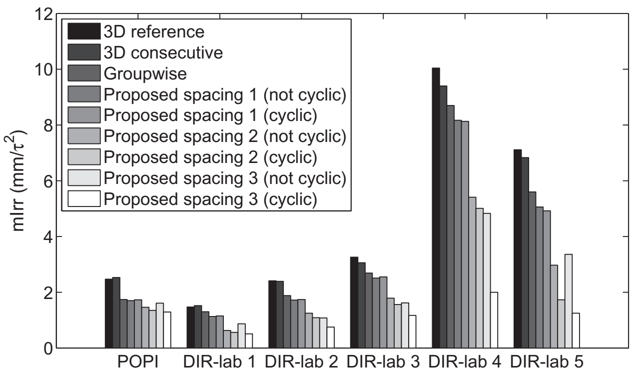

Fig. 4. Irregularity of landmark trajectories using different registration approaches. Values are averaged over 75 landmarks for the DIR-lab images and over 37 landmarks for the POPI-model. Lower values mean smoother results.

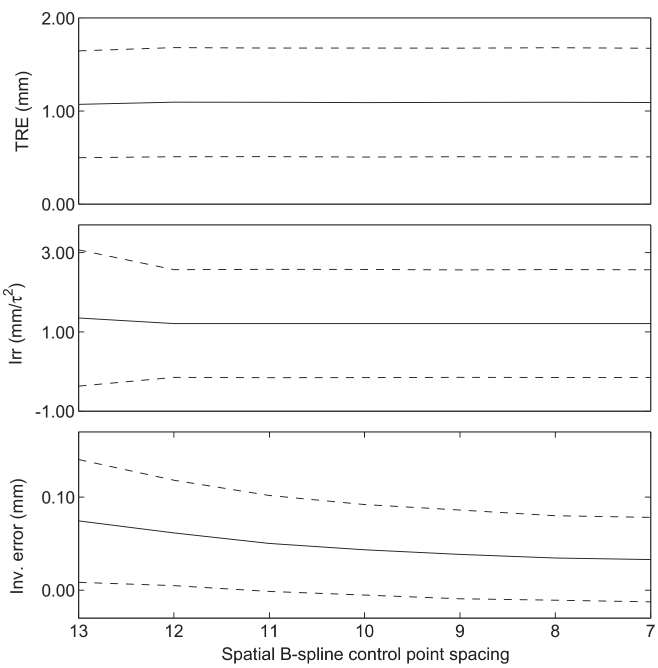

Fig. 5. Results of the inverse experiment. The figures show the influence of the spatial control point spacing on the accuracy (TRE), irregularity (Irr) and inverse error. The solid lines represent the mean and the dashed lines show the mean ± the standard deviation.

$$\boldsymbol{T}^{ij}(\boldsymbol{T}^{ki}(\boldsymbol{x})) = \boldsymbol{T}^{kj}(\boldsymbol{x}). \tag{14}$$

In the proposed method all time points are aligned simultaneously without the use of a reference time point, but there still remains an inconsistency, which is caused by errors in the approximation of the inverse transform.

To assess the inconsistency of both the reference time point and proposed cyclic approach, we performed an experiment on the POPI image. We computed the inconsistency errors as:

$$E^{ij}(\boldsymbol{x}) = \|\boldsymbol{T}^{ij}(\boldsymbol{T}^{ki}(\boldsymbol{x})) - \boldsymbol{T}^{kj}(\boldsymbol{x})\| \tag{15}$$

for all $k \in \{1, \ldots, T\}$, all voxel positions and all $i \neq j$. For the 3D reference approach, $\boldsymbol{T}^{ij}$ were computed by pairwise 3D registration for all i, j. For the proposed method the $\boldsymbol{T}^{ij}$ were computed according to Eq. (11). Similar settings were used for both registration approaches. Results of these experiments are shown in a histogram in Fig. 6 with a solid line for the consecutive registration approach and a dashed line for the proposed registration method. The dotted vertical lines indicate the average inconsistency error for both approaches. The inconsistencies for the proposed registration approach are in general smaller, which is apparent from the peak in the left of the histogram and the smaller average error value. The consistency errors for the 3D reference method could decrease with methods such as proposed by Christensen et al. (2001) and Geng et al. (2005).

3.5. Quantitative evaluation on 3D + t CTA data of the heart

An experiment was performed to assess the accuracy of semi-automatic derivation of left ventricular volume curves from 3D + t CTA data of the heart. To this end, the left ventricle was manually annotated for five patients at 10 time points in the cardiac cycle. The curves describing the left ventricular volume over the cardiac cycle were determined from these manual annotations. Subsequently, these curves were also generated by propagating the end-diastolic manual annotation to all other time points using the transformation resulting from the registration procedures. The proposed method was used both with and without imposing cyclic motion. Registration was performed in a two-step approach. First the registration was performed on the whole 4D image. Subsequently, an atlas based segmentation of the heart surface at end-

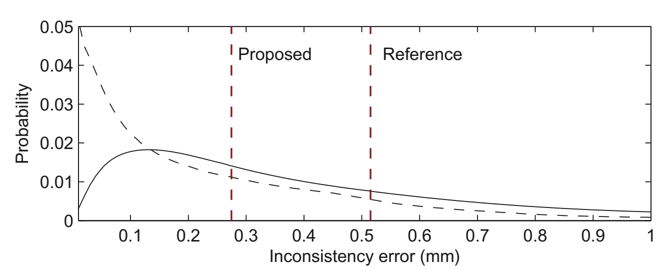

Fig. 6. Histogram of the inconsistency results for the consecutive registration approach (solid line) and for the proposed registration method (dashed line). For visualization purposes the maximum value at the x-axis was set to 1.0 and the maximum value at the y-axis was set to 0.05. The dashed vertical lines indicate the average error for both approaches.

diastole (Kirisli et al., 2010) was propagated to the whole sequence using the resulting transformation. In the second step the 4D registration was performed while using the 4D heart mask for computation of the dissimilarity metric, to be able to handle the non-smooth sliding motion of the heart along the lung surface.

The left ventricular volume curves derived from semi-automatically determined left ventricle surfaces and the manual measurements can be found in Fig. 7. The average and standard deviation of the volume error was 3.02% (±2.46%) and 3.02% (±2.49%) for the cyclic and non-cyclic registration approach respectively. An example of the imaging data before and after registration is shown in Fig. 8.

3.6. Further examples on clinical data

To show the potential of the proposed registration method on other imaging modalities, registration was performed on a 2D + t ultrasound (US) image of the carotid artery and a pediatric 2D + t magnetic resonance (MR) image of the lungs.

The US image of the carotid artery was acquired to measure the distensibility of the carotid artery (Gamble et al., 1994). This requires the accurate estimation of the vessel wall deformation. Images before and after registration are shown in Fig. 9.

The MR image of the lungs was acquired to analyze lung function in cystic fibrosis patients, which requires measuring the compression and decompression of the lungs over the respiratory cycle (Failo et al., 2009). The input image and the results after registration are shown in Fig. 10. The resulting images still show some misalignment, visible on the right side in image (a), caused by the anatomy moving in and out the field of view.

4. Discussion

A registration method for motion estimation in dynamic medical imaging data combining a Lagrangian nD + t B-spline

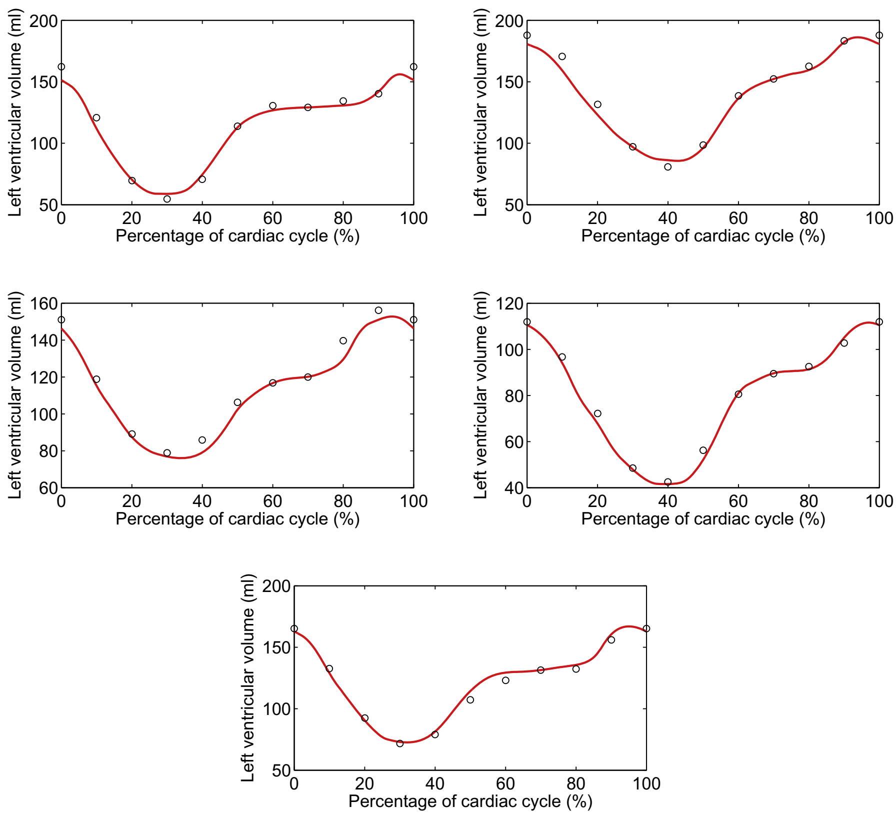

Fig. 7. Left ventricular volume curves derived from 3D + t CTA data. Circles: manual measurements, solid line: proposed method with cyclic motion constraint.

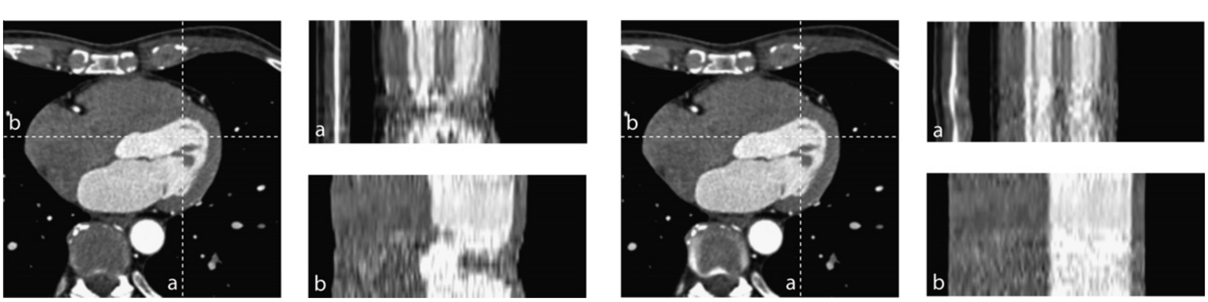

Fig. 8. Registration result for a 3D + t CTA image of the heart. Left: input image, right: registration result. Image (a) and (b) show the dotted lines in the left image over time. The time varying noise levels, which are caused by variation of dose over the cardiac cycle (ECG-pulsing) (Weustink et al., 2009), are clearly visible and make the registration more challenging. Note that the left side of images (b) falls outside the heart mask used in the registration.

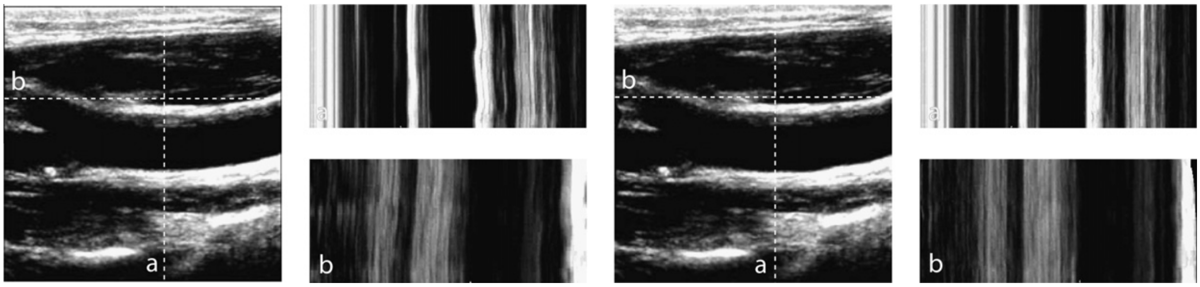

Fig. 9. Registration result for a 2D + t ultrasound image of the carotid artery. Left: input image, right: registration result. Image (a) and (b) show the dotted lines in the left image over time.

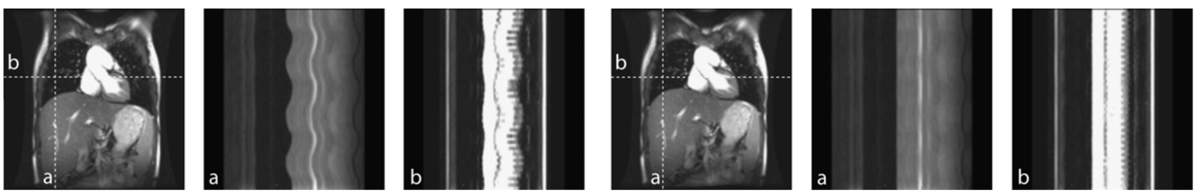

Fig. 10. Registration result for a pediatric 2D + t MRI image of the lungs. Left: input image, right: registration result. Image (a) and (b) show the dotted lines in the left image over time.

transformation model, a global cost function and global optimization strategy has been proposed and quantitatively evaluated.

In a synthetic experiment we compared the proposed approach with other existing registration approaches. It is demonstrated that the temporal smoothness and the constraint on cyclic motion help the registration when images are distorted by noise. Furthermore, the proposed method yields temporally smooth (continuously differentiable) transformations.

Quantitative experiments on 3D + t CT data of the lungs showed that the method was able to derive the dynamics from the images with subvoxel accuracy, which is comparable to previously published results of other state of the art image registration methods on the same data. Furthermore, it shows that temporally smooth results can be achieved by only compromising slightly on registration accuracy. The temporal smoothness of the results can be regulated by adapting the temporal control point spacing, which can be chosen in such a way that it takes into account the expected smoothness of the motion of the anatomy. This prevents the method to fit the transformation to errors in the imaging data, such as acquisition artefacts. Furthermore, cyclic motion can be enforced. We showed that this was only beneficial for one of the DIR-lab cases. A possible explanation for the decrease in accuracy for the other cases can be that the data is not as cyclic as expected, which is also suggested by the relatively low temporal smoothness for the non-cyclic results of DIR-lab case 5 (see Fig. 4). The use of cyclic motion constraints should therefore be considered carefully for the application of interest. It should be noted that previously published results on the DIR-lab data are not directly comparable to the values derived in this paper, because only a subset of 300 from the approximately 1200 3D landmarks is publicly available. We do, however, expect these 300 landmarks to be a representative subset and the results therefore to be representative as well. Also, we computed all TRE values in world coordinates, while (Castillo et al., 2010) round the transformed landmark coordinates to the closest voxel coordinate first. If we would follow their approach, the average values for our results would be the same and the standard deviations would be around 0.4 mm larger. In an experiment on the POPI-model, we showed that the proposed method outperforms the often used reference time point method with respect to registration consistency. This is caused by the use of a global cost function, avoiding a bias towards a specifically chosen reference time point.

Further quantitative experiments on 3D + t CTA data of the heart showed the ability of the method to semi-automatically determine left ventricular volume curves with only a deviation from manually derived curves of approximately 3%. These deviations may be partly explained by errors in the manual annotations, because outlining the left ventricle is especially challenging at the location of the atrioventricular valves during fast moving phases of the cardiac cycle. Furthermore, contrary to the registration method, temporal continuity is not taken into account during manual annotation.

Examples on 2D + t ultrasound and MR images showed the potential of the method for other imaging modalities. Because the method is publicly available as an extension to `elastix`, it can readily be applied by other researchers to various types of dynamic medical imaging data.

We also described an approach to approximate an inverse B-spline transformation, which was used to find the relation between time points in the input image. Based on the application of interest one can choose to perform analyses on the registered image directly or relatively to a chosen reference time point. The errors of the inverse transformation were shown to be very small and even smaller when choosing a somewhat smaller spatial grid spacing for obtaining the inverse transformation than was used in the forward registration procedure. The accuracy was almost not affected by this smaller spacing, but a small effect on the smoothness of the results was noticed.

The method builds upon groupwise registration approaches, with the main difference that the smoothness in the temporal direction of the data is incorporated and that the transformations can be constrained to be cyclic. The advantage of this approach is that the alignment of the data does not depend on a chosen reference time point. Although the deformations between the time points and the used implicit reference frame might be larger than the deformations between consecutive time frames, this does not seem to affect registration robustness. This is most probably an effect of the temporal smoothness which helps to find these larger deformations. Additionally, these deformations are still relatively small compared with the inter-patient differences in groupwise registration. The error introduced when establishing the relation between time points in the input image through application of the inverse transformation was shown to be much smaller than the consistency errors made using the 3D reference registration approach.

Additional constraints on the transformation can be taken into account by adding extra penalty terms to the cost function (Bistoquet et al., 2008; Mansi et al., 2009; Sdika, 2008). These penalty terms could, for example, ensure inverse consistent transformations or impose biomechanical constraints on the transformations. This depends on the application, and is not pursued in our work.

It should be mentioned that the assumption of constant intensity over time does not hold in perfusion imaging, where contrast flow can cause the same anatomy to have a different appearance over time. The development of a similarity metric that accounts for contrast influx would therefore be an interesting future research direction as our groupwise nD + t B-spline framework can accommodate different similarity measures, which may be selected based on different assumptions.

5. Conclusions

A registration method combining a Lagrangian nD + t B-spline transformation model, a global cost function and global optimization strategy for motion analysis in dynamic medical imaging data was proposed. It takes smoothness into account in both the spatial and temporal direction of the data. Moreover, it can enforce the transformations to be cyclic. Registration accuracy and smoothness were assessed using a synthetic image, publicly available imaging data of the lungs and imaging data of the heart. On the synthetic image, the best results with respect to accuracy and smoothness were achieved using the proposed method while imposing cyclic motion. On the lung data, the accuracy was found to be comparable to previously reported results and the smoothness was found to be best when using the proposed approach. Furthermore, it was shown that the proposed method performs better than the reference time point registration method with respect to the consistency of the registration results. Regarding the cardiac CTA data, semi-automatically derived left ventricular volume curves showed a deviation of approximately 3% with respect to the curves derived from manual annotations. Further examples were shown on a 2D + t US image of the carotid artery and a 2D + t MR image of the lungs. The software is publicly available as an extension to the registration package `elastix`.

Acknowledgements

C.T. Metz, M. Schaap and W.J. Niessen are supported by the Dutch Foundation for Scientific Research, NWO, STW. S. Klein is supported by the Dutch Foundation for Scientific Research, NWO, Exact Sciences (VENI grant). The authors would like to thank W.B. Vletter, J.G. Bosch, and K.Y. Leung, Thoraxcenter, Department of Cardiology, Erasmus MC, Rotterdam, The Netherlands, for providing the ultrasound images, R. Failo, H.A.W.M. Tiddens, Departments of Pediatric Radiology and Pulmonology, Erasmus MC, Rotterdam, The Netherlands and M. de Bruijne, Departments of Radiology and Medical Informatics for providing the MR images and L.A. Neefjes and N.R. Mollet, Departments of Radiology and Cardiology, Erasmus MC, Rotterdam, The Netherlands for providing the CTA images of the heart.

References

Balci, S., Golland, P., Shenton, M., Wells, W., 2007a. Free-form B-spline deformation model for groupwise registration. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. – Stat. Regist. Workshop, vol. 2007. Citeseer.

Balci, S., Golland, P., Wells, W., 2007b. Non-rigid groupwise registration using B-spline deformation model. In: MICCAI Open Science Workshop Proceedings, pp. 106–122.

Bhatia, K.K., Hajnal, J.V., Puri, B.K., Edwards, A.D., Rueckert, D., 2004. Consistent groupwise non-rigid registration for atlas construction. In: Proc. IEEE Int. Sympos. Biomed. Imaging: Nano to Macro, pp. 908–911.

Bhatia, K.K., Hajnal, J., Hammers, A., Rueckert, D., 2007. Similarity metrics for groupwise non-rigid registration. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. vol. 10, pp. 544–552.

Bidaut, L.M., VallTe, J.P., 2001. Automated registration of dynamic MR images for the quantification of myocardial perfusion. J. Magn. Reson. Imaging 13 (4), 648–655.

Bistoquet, A., Oshinski, J., Skrinjar, O., 2008. Myocardial deformation recovery from cine MRI using a nearly incompressible biventricular model. Med. Image Anal. 12 (1), 69–85.

Boldea, V., Sharp, G.C., Jiang, S.B., Sarrut, D., 2008. 4D-CT lung motion estimation with deformable registration: quantification of motion nonlinearity and hysteresis. Med. Phys. 35 (3), 1008–1018.

Castillo, R., Castillo, E., Guerra, R., Johnson, V.E., McPhail, T., Garg, A.K., Guerrero, T., 2009. A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Phys. Med. Biol. 54 (7), 1849–1870.

Castillo, E., Castillo, R., Martinez, J., Shenoy, M., Guerrero, T., 2010. Four-dimensional deformable image registration using trajectory modeling. Phys. Med. Biol. 55 (1), 305–327.

Christensen, G.E., Johnson, H.J., 2001. Consistent image registration. IEEE Trans. Med. Imaging 20 (7), 568–582.

Chun, S., Schretter, C., Fessler, J., 2010. Sufficient condition for local invertibility of spatio-temporal 4D B-spline deformations. In: Proc. IEEE Int. Sympos. Biomed. Imaging: Nano to Macro.

De Craene, M., Camara, O., Bijnens, B.H., Frangi, A.F., 2009. Large diffeomorphic FFD registration for motion and strain quantification from 3D-US sequences. In: FIMH '09: Proc. 5th Int. Conf. Funct. Imaging Model. Heart. Springer-Verlag, Berlin, Heidelberg, pp. 437–446.

Failo, R., Wielopolski, P.A., Tiddens, H.A.W.M., Hop, W.C.J., Mucelli, R.P., Lequin, M.H., 2009. Lung morphology assessment using MRI: a robust ultra-short TR/TE 2D steady state free precession sequence used in cystic fibrosis patients. Magn. Reson. Med. 61 (2), 299–306.

Gamble, G., Zorn, J., Sanders, G., MacMahon, S., Sharpe, N., 1994. Estimation of arterial stiffness, compliance, and distensibility from M-mode ultrasound measurements of the common carotid artery. Stroke 25 (1), 11–16.

Ganten, M.-K., Krautter, U., von Tengg-Kobligk, H., Böckler, D., Schumacher, H., Stiller, W., Delorme, S., Kauczor, H.-U., Kauffmann, G.W., Bock, M., 2008. Quantification of aortic distensibility in abdominal aortic aneurysm using ECG-gated multi-detector computed tomography. Eur. Radiol. 18 (5), 966–973.

Geng, X., Kumar, D., Christensen, G.E., 2005. Transitive inverse-consistent manifold registration. Inf. Process Med. Imaging 19, 468–479.

Hawkes, D.J., Barratt, D., Blackall, J.M., Chan, C., Edwards, P.J., Rhode, K., Penney, G.P., McClelland, J., Hill, D.L.G., 2005. Tissue deformation and shape models in image-guided interventions: a discussion paper. Med. Image Anal. 9 (2), 163–175.

Hill, D., Batchelor, P., Holden, M., Hawkes, D., 2001. Medical image registration. Phys. Med. Biol. 46 (3), R1–R45.

Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. Neuroimage 23 (Suppl. 1), S151–S160.

Kabus, S., Klinder, T., Murphy, K., van Ginneken, B., Lorenz, C., Pluim, J., 2009. Evaluation of 4D-CT lung registration. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. Lect. Notes Comput. Sci., vol. 5761, pp. 747–757.

Kaus, M., Netsch, T., Kabus, S., Pekar, V., McNutt, T., Fischer, B., 2004. Estimation of organ motion from 4D CT for 4D radiation therapy planning of lung cancer. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. Lect Notes Comput. Sci., vol. 3217, pp. 1017–1024.

Kirisli, H., Schaap, M., Klein, S., Neefjes, L., Weustink, A., van Walsum, T., Niessen, W., 2010. Fully automatic cardiac segmentation from 3D CTA data: a multiatlas based approach. In: Proc. SPIE Med. Imaging.

Klein, S., Pluim, J.P., Staring, M., Viergever, M.A., 2009. Adaptive stochastic gradient descent optimisation for image registration. Int. J. Comput. Vis. 81 (3), 227–239.

Klein, S., Staring, M., Murphy, K., Viergever, M., Pluim, J., 2010. `elastix`: a toolbox for intensity based medical image registration. IEEE Trans. Med. Imaging 29 (1), 196–205.

Li, G., Citrin, D., Camphausen, K., Mueller, B., Burman, C., Mychalczak, B., Miller, R.W., Song, Y., 2008a. A dvances in 4D medical imaging and 4D radiation therapy. Technol. Cancer Res. Treat. 7 (1), 67–81.

Li, Z.-Y., Howarth, S., Graves, M., Gillard, J.H., 2008b. Measurement of stenotic carotid arterial compliance with MRI. In: Conf. Proc. IEEE Eng. Med. Biol. Soc. 2008, pp. 1403–1406.

Lopez, A., Fritscher, K.D., Trieb, T., Schubert, R., Mattes, J., 2008. 3D inters-subject cardiac registration using 4D information. In: Reinhardt, J.M., Pluim, J.P.W. (Eds.), Medical Imaging 2008: Image Processing, vol. 6914. SPIE, p. 691413.

Lorenzo-Valdés, M., Sanchez-Ortiza, G.I., Mohiaddin, R., Rueckert, D., 2002. Atlas-based segmentation and tracking of 3D cardiac MR images using non-rigid registration. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. Lect. Notes Comput. Sci., vol. 2488, pp. 642–650.

Mahnken, A.H., Bruners, P., Schmidt, B., Bornikoel, C., Flohr, T., Gnnther, R.W., 2009. Left ventricular function can reliably be assessed from dual-source CT using ECG-gated tube current modulation. Invest. Radiol. 44 (7), 384–389.

Maintz, J.B., Viergever, M.A., 1998. A survey of medical image registration. Med. Image Anal. 2 (1), 1–36.

Mansi, T., Peyrat, J.-M., Sermesant, M., Delingette, H., Blanc, J., Boudjemline, Y., Ayache, N., 2009. Physically-constrained diffeomorphic demons for the estimation of 3d myocardium strain from cine-mri. In: FIMH '09: Proc. 5th Int. Conf. Funct. Imaging Model. Heart. Springer-Verlag, Berlin, Heidelberg, pp. 201–210.
Marsland, S., Twining, C., Taylor, C., 2003. Groupwise non-rigid registration using polyharmonic clamped-plate splines. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. Springer, pp. 771–779.
Marsland, S., Twining, C., Taylor, C., 2008. A minimum description length objective function for groupwise non-rigid image registration. Image Vis. Comput. 26 (3), 333–346.
Metz, C., Baka, N., Kirisli, H., Schaap, M., van Walsum, T., Klein, S., Neefjes, L., Mollet, N., Lelieveldt, B., de Bruijne, M., et al., 2010. Conditional shape models for cardiac motion estimation. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv., pp. 452–459.
Miller, E., Matsakis, N., Viola, P., 2000. Learning from one example through shared densities on transforms. In: CVPR. IEEE Computer Society, p. 1464.
Milles, J., van der Geest, R.J., Jerosch-Herold, M., Reiber, J.H.C., Lelieveldt, B.P.F., 2008. Fully automated motion correction in first-pass myocardial perfusion MR image sequences. IEEE Trans. Med. Imaging 27 (11), 1611–1621.
Perperidis, D., Mohiaddin, R., Rueckert, D., 2005. Spatio-temporal free-form registration of cardiac MR image sequences. Med. Image Anal. 9, 441–456.
Peyrat, J., Delingette, H., Sermesant, M., Xu, C., Ayache, N., 2010. Registration of 4D cardiac CT sequences under trajectory constraints with multichannel diffeomorphic demons. IEEE Trans. Med. Imaging 29 (7), 1351–1368.
Rao, A., Sanchez-Ortiz, G.I., Chandrashekara, R., Lorenzo-Valdés, M., Mohiaddin, R., Rueckert, D., 2002. Comparison of cardiac motion across subjects using non-rigid registration. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. Lect Notes Comput Sci., vol. 2488, pp. 722–729.
Reinhardt, J.M., Ding, K., Cao, K., Christensen, G.E., Hoffman, E.A., Bodas, S.V., 2008. Registration-based estimates of local lung tissue expansion compared to xenon CT measures of specific ventilation. Med. Image Anal. 12 (6), 752–763.
Rietzel, E., Chen, G.T.Y., 2006. Deformable registration of 4D computed tomography data. Med. Phys. 33 (11), 4423–4430.
Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. IEEE Trans. Med. Imaging 18 (8), 712–721.
Schreibmann, E., Thorndyke, B., Li, T., Wang, J., Xing, L., 2008. Four-dimensional image registration for image-guided radiotherapy. Int. J. Radiat. Oncol. Biol. Phys. 71 (2), 578–586. Jun.
Sdika, M., 2008. A fast nonrigid image registration with constraints on the Jacobian using large scale constrained optimization. IEEE Trans. Med. Imaging 27 (2), 271–281.
Sundar, H., Litt, H., Shen, D., 2009. Estimating myocardial motion by 4D image warping. Pattern Recognit 42 (11), 2514–2526.
Unser, M., 1999. Splines: a perfect fit for signal and image processing. IEEE Signal Process Mag. 16 (6), 22–38.
van de Kraats, E., Penney, G., Tomazevic, D., van Walsum, T., Niessen, W., 2005. Standardized evaluation methodology for 2-D–3-D registration. IEEE Trans. Med. Imaging 24 (9), 1177–1189.
Vandemeulebroucke, J., Sarrut, D., Clarysse, P., 2007. The POPI-model, a point-validated pixel-based breathing thorax model. In: XVth Int. Conf. Use Comput. Radiat. Therapy (ICCR), Toronto, Canada.
Weustink, A.C., Mollet, N.R., Neefjes, L.A., van Straten, M., Neoh, E., Kyrzopoulos, S., Meijboom, B.W., van Mieghem, C., Cademartiri, F., de Feyter, P.J., Krestin, G.P., 2009. Preserved diagnostic performance of dual-source CT coronary angiography with reduced radiation exposure and cancer risk. Radiology 252 (1), 53–60.
Wierzbicki, M., Drangova, M., Guiraudon, G., Peters, T., 2004. Validation of dynamic heart models obtained using non-linear registration for virtual reality training, planning, and guidance of minimally invasive cardiac surgeries. Med. Image Anal. 8 (3), 387–401.
Xue, H., Guehring, J., Srinivasan, L., Zuehlsdorff, S., Saddi, K., Chefdhotel, C., Hajnal, J.V., Rueckert, D., 2008. Evaluation of rigid and non-rigid motion compensation of cardiac perfusion MRI. Int. Conf. Med. Image Comput. Comput. Assist. Interv. 11 (Pt 2), 35–43.
Zöllei, L., Learned-Miller, E., Grimson, E., Wells, W., 2005. Efficient population registration of 3D data. In: Comput. Vis. Biomed. Image Appl. Springer, pp. 291–301.