

Tensor Grid Based Image Registration with Application to Ventilation Estimation on 4D CT Lung Data

Heike Ruppertshofen · Sven Kabus · Bernd Fischer

Received: date / Accepted: date

Abstract

Purpose For many image registration tasks the information contained in the original resolution of the image data is crucial for a subsequent medical analysis, e.g. accurate assessment of local pulmonary ventilation. However, the complexity of a non-parametric registration scheme is directly connected to the resolution of the images. Therefore, the registration is often performed on a downsampled version in order to meet runtime demands and thereby producing sub-optimal results. To enable the application of the highest resolution at least in regions of high clinical importance, an approach is presented replacing the usually taken equidistant grids by tensor grids for image representation.

Methods We employ a non-parametric approach for the registration of a respiratory-gated 4D CT thorax scan. Tensor grids are introduced for the registration setting and compared to equidistant grids. For ventilation assessment the Jacobian metric is explored.

Results The application of the tensor grid approach makes the local usage of the original resolution feasible; thereby a smaller registration error is achieved in regions of higher resolution using the tensor grids, while the two types of grids perform similar in regions of equal resolution. Concerning the ventilation assessment, the Jacobian metric yields rea-

sonable results, showing more detail using the tensor grids due to the higher resolution.

Conclusions The proposed approach using tensor grids preserves registration accuracy, while reducing computational demands. The application of the Jacobian metric for ventilation assessment in conjunction with tensor grids is promising; however, due to a missing ground-truth the medical relevance could not be established for the ventilation estimation so far.

Keywords Image Registration · Tensor Grid · Pulmonary Ventilation · 4D-CT

1 Introduction

The usefulness of the new registration technique to be introduced is best explained in conjunction with a real example. Therefore, we start out by introducing the application. The assessment of lung function has become increasingly interesting in recent years, especially in diagnosis of pulmonary diseases, e.g. emphysema [1, 2], placement of radiation beams in radiation therapy [3, 4] or drug delivery via inhalation [5]. The two main parameters evaluated to analyze lung function are perfusion and ventilation [6, 7]. The latter is a measure of the gas exchange of the lung and the ambient air, while pulmonary perfusion quantifies the amount of blood flow inside the vessels supplying the alveoli. For the lung to work efficiently, these two parameters have to correlate, meaning that, if the amount of air inside the lung is high, there has to be an increase in blood flow, such that the exchange of oxygen and carbon dioxide can take place.

In this work, we will confine to the assessment of ventilation. Almost all medical imaging modalities can be used for this task [5], with nuclear imaging like PET (Positron Emission Tomography) and SPECT (Single Photon Emission Computed Tomography) being considered the current

Heike Ruppertshofen
University of Applied Sciences Kiel, Institute of Applied Computer
Science, Kiel, Germany
E-mail: heike.ruppertshofen@fh-kiel.de

Sven Kabus
Philips Research Europe - Hamburg, Department Digital Imaging,
Hamburg, Germany

Bernd Fischer
University of Lübeck, Institute of Mathematics, Lübeck, Germany

gold-standard [6]. However, most of these modalities have major drawbacks like low spatial or temporal resolution or the need of special equipment like suitable contrast agents.

To resolve these drawbacks, research has been undertaken on the possibility of evaluating ventilation from 4D-CT (Computed Tomography) images, which consist of several 3D images taken at different phases of the breathing cycle. In radiation therapy planning these images are acquired routinely and thus no further acquisition is necessary to determine the ventilation in this case, yielding a further advantage for 4D-CT usage in ventilation assessment. The ventilation information is extracted from the images by determining the local change in volume relative to a reference phase or at each phase transition [8–11]. To assess the ventilation using the proposed methods, a registration of the different phase images is needed to establish a correspondence between the images and to be able to quantify the change in volume.

There are various ways to handle the registration problem. In this work, we will use a popular non-parametric approach, where a dense deformation vector field is searched for that maps one image onto the other. The advantage of this approach is that a physical model is included in the registration process, which globally models the physical nature of the image and yields a dense deformation field after registration that can be used directly to evaluate the ventilation. However, the problem with non-parametric image registration is that it is subject to the curse of dimensionality. The resolution of medical images has become higher over the last years such that more detail is visible in the images. From a medical point of view the improvement in resolution is remarkable, conversely the increase in data size that goes along with the improved resolution makes computations much more demanding. For a thorax CT scan with a slice thickness of 2.5 mm and a size of about $512^2 \times 141$ voxel, a linear system of equations with more than 100 million unknowns needs to be solved in each iteration step, inducing a challenge in both time and memory.

A popular approach to speed up computations is to embed the registration algorithm into a multi-level setting. By omitting the highest level(s) of resolution the problem of memory is solved as well. This approach is feasible and often results at lower levels are satisfactory. However, Cook et al. [12] stated that the impact of image resolution on the registration result depends locally on the anatomy. This proves right for various medical applications, like e.g. the above mentioned ventilation studies. In this case the information contained in the highest resolution is crucial in some regions of the images and should not be neglected, e.g. trapped air regions (emphysema), where lung texture and thus image contrast is low. The little information contained in these areas will be lost after downsampling. Yet in other regions with distinctive structures, which can still be well distinguished on lower resolution levels, the original resolution is

negligible. Another example is the planning of beam placement for radiation therapy. Here the beams should be placed such that high functional regions in the lung are spared to reduce side effects induced by the radiation. A precise ventilation estimation and hence registration of the image in regions close to the tumor is therefore of higher interest than in regions distant to the tumor.

Thus an approach using adaptive grids should be favored as has been introduced by Papenberg et al. [13] or Haber et al. [14]. Haber et al. propose to use octrees for non-parametric registration. Here the adaptive grid is chosen automatically based on image gradient information; thereby a sparse but also computationally fiddly representation of the problem is obtained. Papenberg et al. initiate a more straightforward approach using a multi-level setting, where the registration starts on a low level with a coarse resolution of the image. The field of view is diminished on each level, while the resolution becomes finer, such that the same number of voxel is processed on every level. Hence the registration focuses on a certain region of interest and does not yield a dense deformation of the entire image.

In this note we will present another method for arriving at a computationally feasible but still satisfying registration method for large data sets. Its main advantage with respect to the above mentioned approaches is its simplicity. The idea is to apply so called tensor grids in the registration framework [15, 16]. These grids achieve to adapt to the problem of memory and processing time by being fine in user chosen regions of interest, while being coarser in the remaining image domain, such that the highest resolution can be applied at least locally. This results in a sparser but still meaningful representation of the image and a reduction in the size of the data.

2 Registration using Tensor Grids

For the task of registration two images are considered, one is called the reference image R and the other the template image T . Both images are defined on a domain Ω . The task is now to find a deformation φ such that the deformed image $T \circ \varphi$ is similar to R .

The task of image registration [17, 18] can be solved using different approaches, which can be subdivided into the two large categories of parametric and non-parametric registration. The easiest parametric approach is to use rigid or affine registration, where a small set of parameters is searched for to globally transform the image. The other big category is non-parametric image registration [19], which will be used within this work. The advantage of non-parametric registration over the parametric methods is the possibility to model the physical behavior of the underlying anatomy in the images.

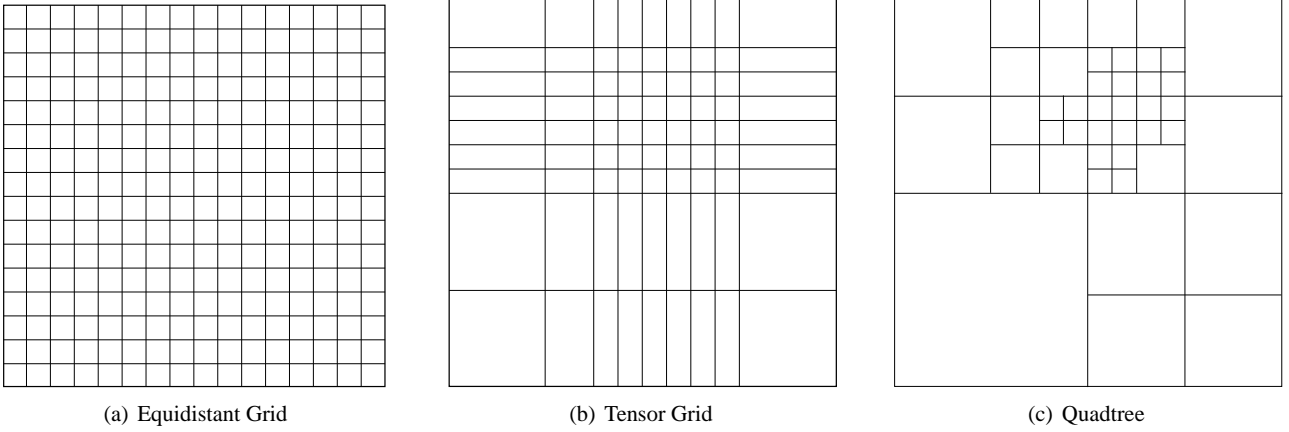


Fig. 1 Comparison of the different types of grids.

However, one disadvantage of non-parametric registration approaches is the high computational demand concerning run time and memory need. To reduce the computational cost we introduce a new type of grid for the registration in the next section, which yields a sparser data representation.

2.1 Tensor Grids

Tensor grids [15] are an adaptive grid structure that can be ranked in-between equidistant grids and quadtrees (2D) or octrees (3D). As example an illustration of each type of grid is given in Fig. 1.

As can be seen there, quadtrees are the sparsest data representation. Starting on a coarse equidistant grid the quadtree is refined locally by subdividing a grid element into four smaller grid elements, until the desired precision is reached. This approach leads to a regular but non-uniform grid, where each grid element may have more than one neighbor on each side. The non-uniformness introduces a significant overhead. It might be cumbersome to compute derivatives and most important, the structure of the resulting linear systems is no longer regular, ruling out the direct application of fast solvers.

Tensor grids have a regular and non-equidistant spacing as well but here every grid element has exactly one neighbor on each side. Therefore its data structure is much simpler compared to quadtrees, resulting in less overhead. Furthermore, the regular spacing of the tensor grid still allows for the usage of the fast and efficient solvers, which typically constitute the computational bottleneck in a registration scheme.

A three dimensional tensor grid is represented by three vectors \mathbf{x}_i , $i = 1, 2, 3$ containing the center coordinates of the grid elements in each dimension. By taking the tensor product of these vectors the tensor grid G is formed:

$$G = \left\{ (x_1^{j_1}, x_2^{j_2}, x_3^{j_3}) \mid j_i = 1, \dots, n_i, i = 1, 2, 3 \right\}.$$

Note that in the following $i = 1, 2, 3$ will represent the dimension, while $j_i = 1, \dots, n_i$ indexes the grid elements.

Opposed to equidistant grids, the grid spacing and the size of the grid elements are not equal for the tensor grid. The resolution of the grid is defined by a coarsest and a finest spacing, noted by $\mathbf{h}_c \in \mathbb{R}^3$ and $\mathbf{h}_f \in \mathbb{R}^3$, respectively, which bound the values of the grid spacing:

$$h_{c_i} \geq h_i^{j_i+1/2} \geq h_{f_i} \quad \forall j_i \forall i,$$

where $h_i^{j_i+1/2}$ is the spacing in-between $x_i^{j_i}$ and $x_i^{j_i+1}$.

To obtain a regular grid, we will constrain the spacing of the tensor grid to values given by a multiplication of the finest spacing by a power of two, such that the coarsest and finest spacing are linked through the following relationship:

$$\mathbf{h}_c = \mathbf{h}_f \cdot 2^{l_c-1},$$

thereby l_c is the number of different resolution levels.

The finest grid spacing will later be given by the image resolution, the coarsest will be user chosen. Depending on the application, the user may chose one or more rectangular regions of interest (ROI), where the grid has the finest resolution h_f .

The tensor grid is built by starting on a coarse equidistant grid with spacing \mathbf{h}_c . In each iteration it is further refined by bisecting the grid elements, where they pass through one of the ROIs. The algorithm stops after l_c iterations when the finest spacing \mathbf{h}_f is reached inside the ROIs. An example of the different steps of the construction can be seen in Fig. 2. The intermediate grids that are generated as a by-product in each iteration of the construction procedure can be saved and later be used for the multi-level setting of the image registration. The resolution inside the ROIs differs by factor two for subsequent intermediate grids; therefore basically leading to a usual image pyramid representation.

Depending on the application and the desired accuracy, the grid can be further constrained or refined by

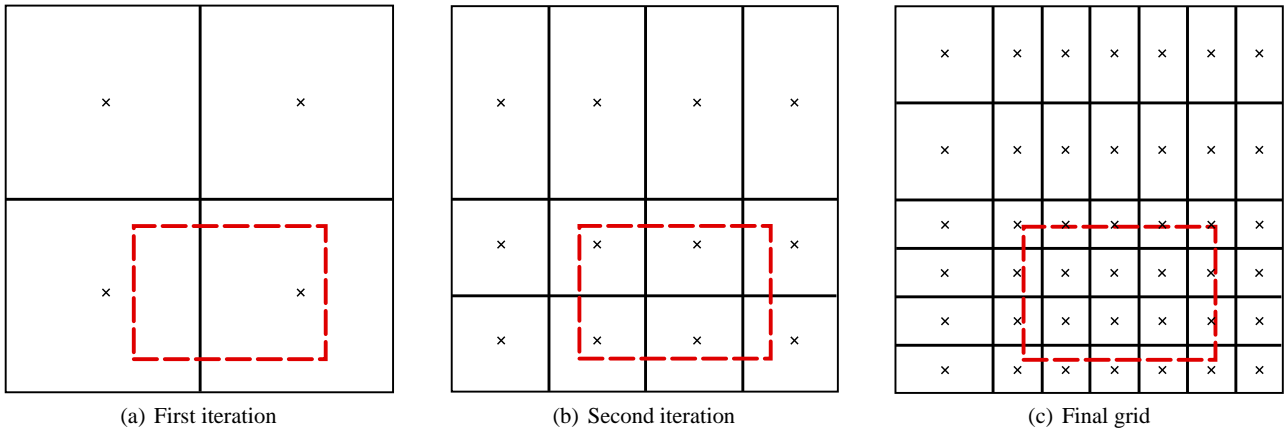


Fig. 2 Intermediate and final grids of the construction. In this example only two levels of resolution are allowed. The ROI chosen for the refinement is shown as a dashed rectangle.

- allowing only a certain number of consecutive resolution levels l_r on each level l , i.e.

$$h_i^{j_i+1/2} \in \left\{ h_{f_i} \cdot 2^k \mid k \in \{l + l_r - 1, \dots, l\}, k \leq l_c \right\} \quad \forall j_i \forall i,$$

- restricting the ratio of the sizes of neighboring grid elements by q , i.e.

$$\frac{h_i^{j_i+1/2}}{h_i^{j_i-1/2}} \leq q \quad \text{and} \quad \frac{h_i^{j_i-1/2}}{h_i^{j_i+1/2}} \leq q \quad \forall j_i \forall i.$$

The tensor grids are applied in the registration approach as described in the next section. For the evaluation of the results in Sect. 5.1 we will use three partitions of the tensor grid, which are defined based on the resolution. Naturally one partition of the grid is the ROI, which has the highest resolution. In addition, we will subdivide the area outside the ROI into the tails of the ROI and the coarse regions. The tails are defined as those regions beside the ROI that have a higher resolution in at least one dimension. The coarse regions are the ones with the lowest resolution.

2.2 Elastic Registration

We will use the tensor grids in a non-parametric image registration approach as described in [19] to determine the transformation φ that maps one image onto another. In the following φ will be separated into:

$$\varphi(x) = x + u(x),$$

where x is the identity and u the deformation that we want to find.

For the registration two elements are needed, one is a distance measure \mathcal{D} and the other one a regularizer \mathcal{S} . A combination of these two elements yields the formulation of

the non-parametric images registration problem in a variational setting [19], where given two smooth images

$$R, T : \Omega \rightarrow \mathbb{R}$$

a deformation vector field $u \in \mathcal{C}^2(\overline{\Omega}, \mathbb{R}^3)$ is searched for such that the functional \mathcal{J} ,

$$\mathcal{J}[u] := \mathcal{D}[R, T; u] + \alpha \cdot \mathcal{S}[u] \xrightarrow{u} \min, \quad (1)$$

is minimized. At this $\mathcal{C}^2(\overline{\Omega}, \mathbb{R}^3)$ is the class of twice continuously differentiable functions, $u : \overline{\Omega} \rightarrow \mathbb{R}^3$ and $\alpha \in \mathbb{R}^+$ is a weighting parameter that regulates the impact of the regularizer and the distance measure on the registration. A larger α will thereby yield a smoother deformation field.

The distance measure \mathcal{D} determines the similarity of the two images. It should be a function that, given the two images and the deformation, will return a scalar value. Depending on the application this similarity measure can be based on intensity, correlation or mutual information. Since we will work with monomodal images, the intensity based sum of squared differences (SSD) will be used, which compares the images pointwise:

$$\mathcal{D}[R, T; u] := \frac{1}{2} \int_{\Omega} [R(x) - T_u(x)]^2 dx,$$

with T_u being an abbreviated form of the deformed image $T(x + u(x))$.

The SSD becomes smaller if the difference of the intensity values of the two images becomes smaller.

The regularizer \mathcal{S} is used to privilege certain deformations, which are more likely or reasonable than others. The reasonability of the deformation is given in most cases by a physical model. For this reason the regularizer can also be used to model the physical behavior of the underlying anatomy in the image by e.g. presuming that the template

image emerged from the reference image through a diffusion process or by making assumptions about the material properties of the image as is done e.g. in elastic or fluid registration. For the registration of the lung images the elastic regularizer will be used, which has first been introduced by Broit [20]:

$$\mathcal{S}[u] := \int_{\Omega} \frac{\mu}{4} \sum_{i,j=1}^3 (\partial_{x_j} u_i(x) + \partial_{x_i} u_j(x))^2 + \frac{\lambda}{2} (\nabla \cdot u(x))^2 dx.$$

λ and μ are the so called Lamé-parameters, which are used to model the elastic behavior.

A possible physical interpretation of the elastic regularizer is to imagine the image to be drawn upon an elastic canvas. The elasticity of the canvas is determined by λ and μ . If a force is applied to the canvas, it will stretch. How strong it stretches and how strong the neighboring material will be affected depends on the Lamé-parameters.

Setting the associated Gâteaux derivative of (1) to zero, results in a system of non-linear partial differential equations for u , which constitute a necessary condition for a minimizer of (1). Given appropriate boundary conditions, the system can be discretized and finally solved with e.g. a conjugate gradient method [21].

To accelerate the algorithm, it will be embedded into a multi-level setting by performing the registration on subsequent levels of an image pyramid; thereby the images are first aligned on a coarse level where the size of the images is small and the registration is fast. The resulting deformation vector field is upsampled and used as an input for the next level of the pyramid, which thus has a better starting point and therefore converges after fewer iterations. Furthermore, this approach does not only speed up computations but also helps to bypass local minima of the registration problem, since the images are first aligned on a coarser level where less detail is visible.

In the next section, an application for the registration in form of lung ventilation assessment is presented.

3 Lung Ventilation

Pulmonary ventilation is the exchange of gas between the lungs and the ambient air. The amount of ventilation provides information about the function of the lung.

The assessment of ventilation is needed for the distinction of high- and low-functional regions of the lung. This is important e.g. in the diagnosis of emphysema or the planning of the placement of beams during radiation therapy [3, 4]. For the latter the ventilation information should be used to avoid high-functional areas of the lung and therefore reduce the damage imposed upon healthy tissue.

Pulmonary ventilation can be measured globally with a spirometer or locally using different imaging modalities. A

spirometer measures the amount of air that is taken in and out during each breath. However, due to the ability of the lung to compensate for mal-functioning areas, this measurement is not very precise and differences in lung ventilation will only be apparent in late stages of pulmonary diseases like lung tumors or emphysema; therefore a local analysis of the ventilation has a higher significance. An overview of the imaging modalities used for lung function analysis can be found in Hoffman et al. [5]. Among these are nuclear medicine imaging [22], currently also considered as the gold-standard, magnetic resonance imaging [6] and Xenon-CT [11]. However, all of these modalities have major drawbacks like the need of special contrast agents, lack of anatomical information, susceptible artifacts or low axial coverage.

Another imaging modality that can be used for the assessment of ventilation and that resolves some of the just mentioned drawbacks is 4D-CT [9,2], which we will use for the evaluation of the ventilation. Advantages of this approach include the high resolution of the images, the availability of the equipment in most hospitals and the combination of functional and anatomical information in the images. In addition, the evaluation of the 4D-CT images yields a distribution of the ventilation over the breathing cycle, instead of one average value as is obtained from the above mentioned static imaging modalities.

For the evaluation of the ventilation using the 4D-CT images a registration of the different phases of the dataset is needed as has been realized in the previous section. In literature different metrics¹ have been proposed [8–11] that analyze the ventilation based on the registration results and the change in volume that can be observed therein. All of these metrics assume that the change in volume is related to the amount of ventilation.

In the following, we will introduce one of the proposed ventilation metrics, which from our point of view yields the most reasonable result [15,9].

3.1 Jacobian Metric

The Jacobian metric, which will be applied here to evaluate the ventilation, is given by the Jacobian determinant of the deformation:

$$\Delta V_{Jac}(x) := \det(\nabla(x + u(x))) - 1.$$

The absolute value of the Jacobian at a point x yields the factor by which the deformation vector field u expands or contracts a given volume or, in our case, more precisely the neighborhood of a voxel. The metric is centered around

¹ Note that in conjunction with the ventilation analysis the term metric is not to be understood in the mathematical sense of a metric.




Fig. 3 Coronal view of the lung in the end-exhale phase with the tumor visible in the superior right lobe.

zero, such that negative values indicate a volume contraction, while positive values correspond to an expansion. Assuming that the ventilation and the change in volume within the lungs are correlated, the Jacobian metric gives a physiologically meaningful estimate of the ventilation. Furthermore, the metric is little sensitive to noise and artifacts due to the smoothness of the underlying vector field which is used for the evaluation.

Several research groups [9, 23] have found a good global agreement of the known volume change of the entire lung with the volume change computed with the Jacobian analysis as well as a good local correlation to Xenon-CT [11], supporting our decision to favor this metric for the ventilation assessment.

4 Material

For the evaluation of the registration algorithm and the subsequent ventilation assessment the POPI-Data provided by Vandemeulebroucke et al. [24] was used. The dataset is freely available on the internet² and contains a breathing gated 4D-CT image with 10 phases covering the whole breathing cycle, with phase 7 being the end-exhale and phase 1 being the end-inhale state. A sample slice of the end-exhale image with a tumor visible in the superior lobe of the right lung can be seen in Fig. 3. The images were acquired on a Philips Brilliance CT Big Bore Oncology with an isotropic in-slice resolution of 0.98 mm and a slice thickness of 2 mm. After reducing the field of view to the lungs, the images had a size of $352 \times 288 \times 128$ voxel.





Fig. 4 Coronal (top) and transversal (bottom) view of the lung together with the chosen ROIs. The landmarks that lie within a ROI are shown in red.

To assess the registration accuracy, each of the 10 phase images was equipped with a set of 100 landmarks. The landmarks were first set in phase 7 using the Distinctive Point Finder (DPF) by Murphy et al. [25]. The advantage of the DPF is that it does not necessarily place the landmarks on all the bifurcations of the bronchial tree as a medical expert would most likely do, but instead searches for distinctive points that are the furthest possible apart from each other. Thereby a set of landmarks is created that is equally distributed across the whole lung and yields a better ground-truth since also regions with little texture are included in the

² <http://www.creatis.insa-lyon.fr/rio/popii-model>




Fig. 5 Boxplot of the distribution of the inter-observer error (left) together with an explanation of the boxplot (right).

evaluation. In Fig. 4 a projection of the landmarks on a coronal and transversal slice of the end-exhale image is shown.

The landmarks were propagated to the remaining phase images by two independent observers using the Matching Points Annotator also supplied by Murphy et al. Dependent on the phase a mean inter-observer error of 0.59 mm to 1.32 mm was noted. The distribution of the inter-observer error can be seen in form of a boxplot in Fig. 5. The almost constant value of the median in most phases is outstanding, but can be explained by the fact that more than 50% of the landmarks annotated by the two observers differ only by one voxel or less. A closer look at the outliers visible in the plots showed that these landmarks lie within artifact regions showing duplicated or blurred structures, which hamper the right positioning of the landmarks.

Six different ROIs were chosen for the evaluation, which can also be seen in Fig. 4. The ROIs are named after their position in the lung (upper, middle, lower and right, left) meaning that ROI_{lr} is located in the lower part of the right lung.

It should be noted that for the ventilation assessment no ground-truth was available. Therefore an evaluation of the medical significance of the computed ventilation was not possible.

5 Results

5.1 Registration

For a comparison of the tensor and the equidistant grid approach, the registration was run for both types of grids using the same set of parameters in the registration. Thereby the

Lamé-parameters were set to $\mu = 0.0025$ and $\lambda = 0$; these values were chosen based on previous analyses. The computations were performed in an all-to-phase setting, where one reference phase is chosen and the remaining phases are registered to this reference phase. As reference image the end-exhale phase was chosen since it is supposed to have the least artifacts [26]. As we are only interested in a precise registration of the lungs, the images underwent a preprocessing, where gray-values above zero HU were clamped by convolution with a smooth function; thereby the strong influence of surrounding brighter structures, in particular the ribs, is eliminated. The tensor grids were constrained with the number of allowed resolution levels in one grid $l_r = 3$ and the ratio of neighboring grid elements $q = 2$. For the multi-level setting six levels were chosen such that the registration starts on an image with $11 \times 9 \times 8$ voxel. Due to the size of the images, we were only able to run the algorithm based on full equidistant grids up to level 3, while the algorithm using tensor grids does not cause any trouble up to level 1. Furthermore, to speed up the registration, it was preceded by an affine pre-registration. In addition, the pre-registration is needed to incorporate the affine motion of the lungs, which would otherwise be penalized by the elastic regularizer [19].

In Tab. 2 the processing time and image size applying the different grids are displayed. As can be seen there the reduction in image size for the images on the equidistant grid on level 3 and the tensor grids on level 1 is about the same. Yet, the tensor grid has a much higher resolution in certain areas of the image, namely in the ROIs and its tails, while still covering the whole field of view. The difference in processing time compared to the number of voxel is striking but can be explained by a small overhead that is introduced by the tensor grid approach and a higher number of iterations since

Table 1 Comparison of the target registration error using the equidistant grid and the tensor grid evaluated for the different regions of the tensor grids. Results are shown for all ROI's, also stating the number of landmarks for each region. The error is displayed as 'mean \pm std (max)'.

ROI	region	#landmarks	before registration	using equidistant grid	using tensor grid	improvement
upper left	all	100	6.28 \pm 3.32 (20.00)	1.80 \pm 1.21 (6.40)	1.85 \pm 1.44 (7.11)	-0.05
	roi	4	2.41 \pm 1.05 (3.50)	0.97 \pm 0.20 (1.23)	0.77 \pm 0.18 (1.01)	0.20
	tail	52	5.74 \pm 2.99 (14.04)	1.60 \pm 0.98 (4.88)	1.41 \pm 1.00 (4.71)	0.19
	coarse	44	7.26 \pm 3.46 (20.00)	2.12 \pm 1.46 (6.40)	2.47 \pm 1.68 (7.11)	-0.35
upper right	all	100	6.28 \pm 3.32 (20.00)	1.80 \pm 1.23 (6.40)	1.89 \pm 1.37 (7.10)	-0.09
	roi	2	4.96 \pm 0.31 (5.18)	1.89 \pm 0.25 (2.07)	1.18 \pm 0.54 (1.56)	0.71
	tail	36	6.35 \pm 3.77 (14.53)	1.86 \pm 1.11 (4.88)	1.72 \pm 1.06 (4.60)	0.14
	coarse	62	6.27 \pm 3.11 (20.00)	1.77 \pm 1.32 (6.40)	2.01 \pm 1.53 (7.10)	-0.24
middle left	all	100	6.28 \pm 3.32 (20.00)	1.80 \pm 1.23 (6.40)	1.82 \pm 1.34 (6.95)	-0.02
	roi	3	4.43 \pm 0.45 (4.95)	0.78 \pm 0.47 (1.27)	0.74 \pm 0.21 (0.88)	0.04
	tail	47	5.11 \pm 2.18 (10.29)	1.44 \pm 0.87 (3.93)	1.19 \pm 0.81 (3.85)	0.25
	coarse	50	7.49 \pm 3.85 (20.00)	2.21 \pm 1.41 (6.40)	2.48 \pm 1.45 (6.95)	-0.27
middle right	all	100	6.28 \pm 3.32 (20.00)	1.80 \pm 1.23 (6.40)	1.83 \pm 1.25 (6.51)	-0.03
	roi	2	6.51 \pm 0.50 (6.86)	0.91 \pm 0.43 (1.21)	0.61 \pm 0.00 (0.61)	0.30
	tail	34	6.08 \pm 2.75 (13.80)	1.82 \pm 1.11 (4.88)	1.73 \pm 0.94 (4.41)	0.09
	coarse	64	6.37 \pm 3.65 (20.00)	1.83 \pm 1.30 (6.40)	1.92 \pm 1.38 (6.51)	-0.09
lower left	all	100	6.28 \pm 3.32 (20.00)	1.80 \pm 1.23 (6.40)	1.78 \pm 1.21 (6.31)	0.02
	roi	2	7.54 \pm 0.03 (7.56)	1.44 \pm 0.94 (2.11)	0.79 \pm 0.63 (1.24)	0.65
	tail	30	8.36 \pm 3.23 (15.02)	2.30 \pm 1.38 (5.49)	2.04 \pm 1.35 (5.69)	0.26
	coarse	68	5.32 \pm 2.98 (20.00)	1.60 \pm 1.11 (6.40)	1.69 \pm 1.14 (6.31)	-0.09
lower right	all	100	6.28 \pm 3.32 (20.00)	1.80 \pm 1.23 (6.40)	1.80 \pm 1.23 (6.84)	0.00
	roi	2	10.49 \pm 0.62 (10.92)	4.71 \pm 0.62 (5.15)	4.62 \pm 1.85 (5.93)	0.09
	tail	18	9.63 \pm 3.33 (15.02)	2.47 \pm 1.55 (5.49)	2.31 \pm 1.49 (5.53)	0.16
	coarse	80	5.42 \pm 2.77 (20.00)	1.58 \pm 1.00 (6.40)	1.61 \pm 1.04 (6.84)	-0.03

Table 2 The table shows the number of voxel, reduction in image size in % and average total processing time in minutes for the equidistant grid on level 1 and 3 and the tensor grids on level 1. Note that the processing time for the equidistant grid on level 1 is only a prediction.

grid	#voxel	reduc. [%]	proc. time [min]
eq. grid, level 1	12,976,128	0	\approx 300
eq. grid, level 3	405,504	96.88	12.70
tensor grid, ROI _{ul}	784,875	93.95	78.34
tensor grid, ROI _{ur}	731,745	94.36	65.91
tensor grid, ROI _{ml}	728,910	94.38	67.72
tensor grid, ROI _{mr}	707,940	94.54	62.96
tensor grid, ROI _{ll}	728,886	94.38	56.51
tensor grid, ROI _{lr}	679,470	94.76	52.94

the tensor grid approach is run on all six levels of the image pyramid unlike the approach applying the equidistant grids. If the registration on the equidistant grid had been performed up to the first level of the image pyramid a total processing time of up to 5 hours would have been expected.

Considering the landmark error of the registered images using the two different types of grids, we would expect a lower error inside the ROIs and possibly in the tails using the tensor grids due to the higher resolution. Since the resolution of the images in the coarse regions is equal for both types of grids, we would expect a similar behavior there. As a reminder, the resolution inside the ROIs is equal to level 1,

in the tails it is approximately level 2, while the resolution in the coarse regions is the same as for the equidistant grid namely level 3. Exemplary, the landmark error for the registration of the end-inhale to the end-exhale phase is shown in Tab. 1 for all proposed ROIs. From there it becomes obvious that globally no significant change in registration error can be observed. However, locally in the ROIs and the tails an improvement in registration accuracy by up to 1 mm is achieved.

To show that, apart from increasing registration accuracy in some regions of the image, the tensor grid approach is also beneficial for subsequent medical evaluations, the obtained registration results are applied to assess the ventilation in the next section.

5.2 Ventilation

In conjunction with the radiotherapy planning, the registration result using the tensor grid with the ROI in the superior right lobe, which is located close to the tumor, is evaluated for the ventilation assessment. To this end the Jacobian metric will only be computed inside the lungs, where the ventilation takes place. As segmentation of the lungs the masks provided within the POPI-data were used, which were ob-

tained by simple thresholding and connected labelling techniques [24].

The resulting ventilation maps for the different stages of the breathing cycle can be seen in Fig. 6. The images exhibit mainly a blue colouring, which indicates an expansion of the lungs. It can be observed that the blue coloring grows stronger towards the end-inhale phase and subsides again thereafter. Since the end-exhale phase was chosen as reference phase this result is physiologically reasonable and would have been expected. The first and last images show some light brown areas which suggest that the inhale and exhale phase do not start at the same point in time in all regions of the lungs. Note that the brownish spots close to the diaphragm can be attributed to slight motion artifacts, which may occur in 4D imaging.

To justify the usage of the tensor grids a coronal view of the ventilation images assessed with both approaches is given in Fig. 7. It is clearly visible that the approach using the tensor grids yields more detail, especially in the ROI, than the one applying equidistant grid. However, the medical relevance of these details still needs to be proven.

6 Discussion

An approach using tensor grids in non-parametric image registration was presented. The advantage of this type of grid over the common equidistant grids is the possibility to use a higher resolution in user chosen ROIs, while keeping the original field of view. The sparser image representation results in a lower memory demand and a faster algorithm on each level of the multi-level setting. Although the size of the images used in this work is not very impressive, they were still too large to carry out the registration on equidistant grids with the original image resolution on a desktop PC due to lack of memory. By applying the tensor grid approach it was possible to at least locally take advantage of the original resolution and the information contained therein.

The registration of a 4D-CT thorax scan using different ROIs for the tensor grid showed that both algorithms work equally well on the coarse regions, while the tensor grid approach yields better results within the ROIs and its tails thanks to the higher resolution, as would have been expected. A decrease in landmark error of up to 1 mm was observed inside the ROIs, while preserving registration accuracy on the remaining part of the image. In this analysis the same registration parameters, which were determined for the equidistant approach, were used for both types of grids for a better comparability. By optimizing the parameters for the tensor grid approach a further improvement in accuracy is conceivable.

Concerning ventilation studies an accurate registration of the complete lungs is needed for a precise analysis of

lung function. Using equidistant grids the highest level of resolution needs to be omitted due to the size of the images and thus details in diseased regions are easily lost. The tensor grid approach enabled the registration of the complete lungs, putting focus on user chosen regions, which could be of high interest for the analysis.

Using the results of the registration, it turned out that the ventilation assessed inside the ROIs showed more detail when computed on the results using the tensor grid approach relative to the registration performed on the equidistant grid. The images indicate an expansion for all phases compared to the end-exhale reference phase. From our point of view, the Jacobian metric therefore yields a physiologically meaningful result. Yet the significance of the ventilation map and the higher level of detail inside the ROIs still needs to be determined since no ground-truth was given for this evaluation. Nevertheless, the higher resolution in the ventilation images should be useful for a precise classification of the severity of emphysema or for the placement of radiation beams in radiation therapy planning to avoid high-functional regions of the lung.

Moreover, the tensor grid approach is not only useful for the assessment of ventilation, but can be applied to any problem, where a rather good registration of a complete image is needed, with a high interest in certain regions. If a focus in multiple areas of the image is of high interest, it is also possible to include several ROIs in one single grid and process them at once.

Acknowledgements We thank the Léon Bérard Cancer Center & CRE-ATIS lab, Lyon, France for the distribution of the POPI dataset, Keelin Murphy for providing the landmark tools and Tobias Klinder for carrying out the landmark propagation as one observer.

References

1. T Yamamoto, S Kabus, J von Berg, T Klinder, T Blaffert, C Lorenz, and P Keall. Physiological validation of 4D-CT-based ventilation imaging in patients with chronic obstructive pulmonary disease (COPD). *Med. Phys.*, 36(6):2821–2821, 2009.
2. M Haraguchi, S Shimura, W Hida, and K Shirato. Pulmonary function and regional distribution of emphysema as determined by high-resolution computed tomography. *Respiration*, 65:125–129, 1998.
3. B P Yaremko, T M Guerrero, J Noyola-Martinez, R Guerra, D G Lege, L T Nguyen, P A Balter, J D Cox, and R Komaki. Reduction of normal lung irradiation in locally advanced non-small-cell lung cancer patients, using ventilation images for functional avoidance. *Int J Radiat Oncol Biol Phys*, 68(2):562–571, 2007.
4. T Yamamoto, S Kabus, J von Berg, C Lorenz, and P J Keall. Impact of four-dimensional CT-derived pulmonary ventilation images on radiotherapy treatment planning for lung cancer. *Int J Radiat Oncol Biol Phys*, 75(3):S443, 2009. Supplement.
5. E A Hoffman and D Chon. Computed tomography studies of lung ventilation and perfusion. In *Proc Am Thorac Soc*, volume 2, pages 492–498, 2005.

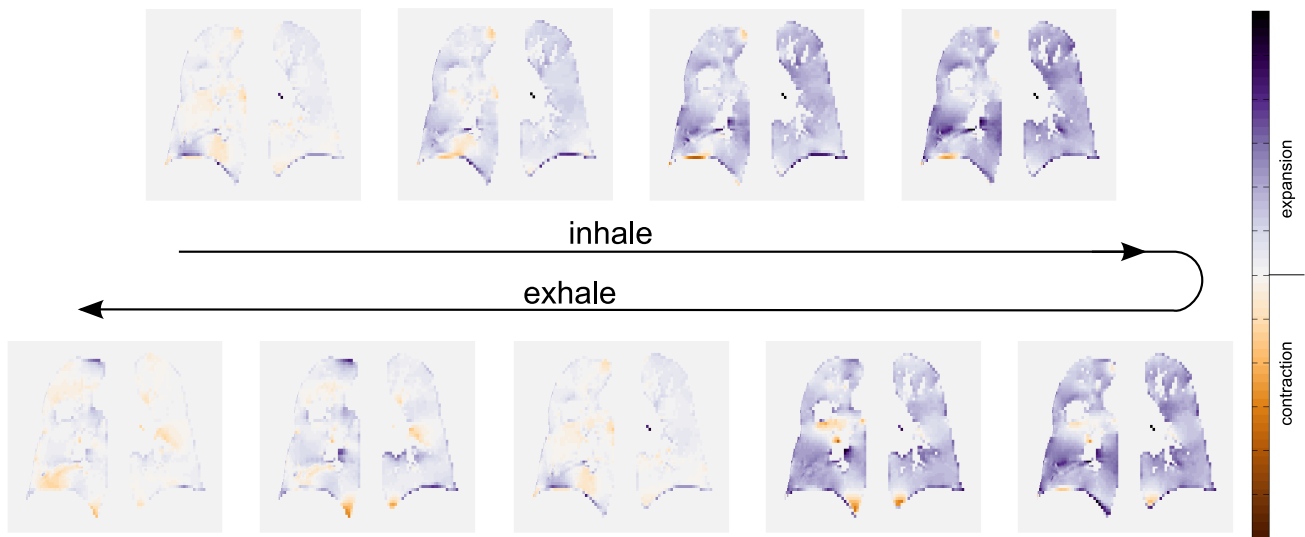


Fig. 6 Ventilation maps computed for the different phases using the registration result with the ROI in the middle of the left lung close to the

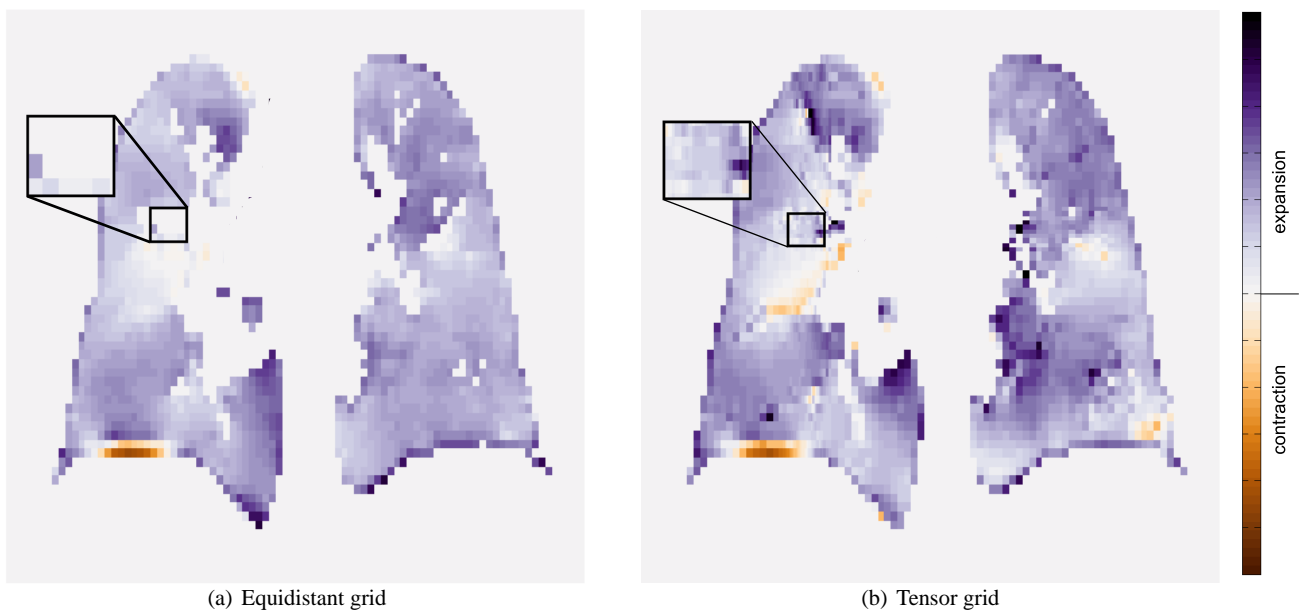


Fig. 7 Comparison of the ventilation map computed from the registration results applying the two different types of grids.

6. H-U Kauczor, A Hanke, and E J R van Beek. Assesment of lung ventilation by MR imaging: Current status and future perspectives. *Eur Radiol*, 12(8):1962–1970, 2002.
7. H-U Kauczor, editor. *Functional Imaging of the Chest*. Springer, 2004.
8. K Cao, K Ding, M L Raghavan, G E Christensen, E A Hoffman, and J M Reinhardt. Registration-based estimates of lung tissue strain in supine sheep. In *Proc International Summit on the Future of Quantitative and Functional Lung Imaging*, 2008.
9. S Kabus, J von Berg, T Yamamoto, R Opfer, and P Keall. Lung ventilation estimation based on 4D-CT imaging. In *Proc First international workshop on pulmonary image analysis, MICCAI*, pages 73–81, 2008.
10. M K Fuld, R B Easley, O I Saba, D Chon, J M Reinhardt, E A Hoffman, and B A Simon. CT-measured regional specific volume change reflects regional ventilation in supine sheep. *J Appl Physiol*, 104:1177–1184, 2008.
11. J M Reinhardt, K Ding, K Cao, G E Christensen, E A Hoffman, and S V Bodas. Registration-based estimation of local lung tissue expansion compared to xenon CT measures of specific ventilation. *Med Image Anal*, 12:752–763, 2008.
12. T S Cook, N Tustison, J Biederer, R Tetzlaff, and J Gee. How do registration parameters affect quantitation of lung kinematics. In *Proc Medical Imaging and Computer Computing and Computer Assisted Intervention (MICCAI)*, pages 817–824, 2007.
13. N Papenberg, J Modersitzki, and B Fischer. Registrierung im Fokus. In *Proc Bildverarbeitung für die Medizin (BVM)*, pages 138–142, 2008.
14. E Haber, S Heldmann, and J Modersitzki. Adaptive mesh refinement for nonparametric image registration. *SIAM J Sci Comput*,

-
- 30(6):3012–3027, 2008.
15. Heike Ruppertshofen. Using tensor grids for non-parametric image registration of 4D CT lung data for ventilation studies. Master's thesis, Universität zu Lübeck, 2008.
 16. Heike Ruppertshofen, Sven Kabus, and Bernd Fischer. Image registration using tensor grids for lung ventilation studies. In *Proc Bildverarbeitung für die Medizin (BVM)*, pages 117–121, 2009.
 17. Bernd Fischer and Jan Modersitzki. Ill-posed medicine – an introduction to image registration. *Inverse Probl*, 24(3):1–16, 2008.
 18. B Zitová and Jan Flusser. Image registration methods: a survey. *Image Vis Comput*, 21:977–1000, 2003.
 19. J Modersitzki. *Numerical Methods for Image Registration*. Oxford University Press, 2004.
 20. C Broit. *Optimal Registration of Deformed Images*. PhD thesis, University of Pennsylvania, 1981.
 21. C Geiger and C Kanzow. *Numerische Verfahren zur Lösung unrestringierter Optimierungsverfahren*. Springer, 1999.
 22. R S Harris and D P Schuster. Visualizing lung function with positron emission tomography. *J Appl Physiol*, 102:448–458, 2007.
 23. B A Simon, G E Christensen, D A Low, and J M Reinhardt. Computed tomography studies of lung mechanics. In *Proc Am Thorac Soc*, volume 2, pages 517–521, 2005.
 24. J Vandemeulebroucke, D Sarrut, and P Clarysse. The POPI-Model, a point-validated pixel-based breathing thorax model. In *Proc International Conference on the Use of Computers in Radiation Therapy (ICCR)*, 2007.
 25. K Murphy, B van Ginneken, J P W Pluim, S Klein, and M Staring. Semi-automatic reference standard construction for quantitative evaluation of lung CT registration. In *Proc Medical Imaging and Computer Computing and Computer Assisted Intervention (MICCAI)*, pages 1006–1013, 2008.
 26. S S Vedam, P J Keall, V R Kini, H Mostafavi, H P Shukla, and R Mohan. Acquiring a four-dimensional computed tomography dataset using an external respiratory signal. *Phys Med Biol*, 48(1):45–62, 2003.