Image Registration using Tensor Grids for Lung Ventilation Studies

Heike Ruppertshofen[1,2], Sven Kabus[2], Bernd Fischer[1]

[1]Institut für Mathematik, Universität zu Lübeck, Wallstraße 40, 23560 Lübeck
[2]Philips Research Europe - Hamburg, Röntgenstraße 24-26, 22335 Hamburg
heike.ruppertshofen@philips.com

Abstract. In non-parametric image registration it is often not possible to work with the original resolution of the images due to high processing times and lack of memory. However, for some medical applications the information contained in the original resolution is crucial in certain regions of the image while being negligible in others. To adapt to this problem we will present an approach using tensor grids, which provide a sparser image representation and thereby allow the use of the highest image resolution locally. Applying the presented scheme to a lung ventilation estimation shows that one may considerably save on time and memory while preserving the registration quality in the regions of interest.

1 Introduction

One possible way to assess lung ventilation is to evaluate the result of non-parametric image registration of thorax CT scans at different stages of the breathing cycle [1]. For this evaluation a dense deformation vector field is needed that maps one phase of the breathing cycle onto another, providing information about the motion of the lung and therefore about the gas exchange. However, finding this deformation field in a non-parametric registration approach is a challenging task mainly due to the size of the images. For a typical thorax CT scan of size $512^2 \times 141$ voxels a linear system of equations with more than 10 million unknowns needs to be solved in each iteration step inducing a challenge in both time and memory.

A general approach to speed up computations is to embed the registration algorithm into a multi-level setting, which in addition enhances the likelihood of bypassing local minima. By omitting the highest level(s) of resolution the problem of memory can be solved as well. Often, results on a lower level are satisfying. However, Cook et al. [2] stated that the impact of the image resolution on the registration result depends locally on the anatomy. For ventilation studies the information contained in diseased regions (e.g. trapped-air regions) is especially important but often lost after the first downsampling. Thus an approach using adaptive grids should be favored.

An adaptive scheme based on octrees has been presented by Haber et al. [3]. Here the computational grid is automatically chosen based on image gradient

information, leading to a sparse but computationally tricky representation of the problem. Papenberg et al. [4] developed a multi-level approach, where the field of view becomes smaller with every level, hence focusing only on a certain region and not the entire lung.

In this contribution we will introduce an approach using tensor grids, which are straightforward to implement and offer a sparse but meaningful representation of the images. These grids are fine in user chosen regions of interest (ROI), while being coarser in the remaining image domain, yielding a deformation field for the complete field of view. To our best knowledge, tensor grids have not been used in a registration framework before.

2   Materials and methods

2.1   Material

For the evaluation the POPI-Data provided by Vandemeulebroucke et al. [5] were used. The dataset consists of 10 breathing related 3D thorax CT images covering the whole breathing cycle. 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. Furthermore, the CT images are equipped with 41 landmarks, 37 of which are located in the lung and were used for evaluation. The landmarks were annotated by medical experts.

2.2   Elastic registration

The non-parametric registration algorithm [6] applied here uses a standard variational approach. The aim is to minimize a functional $\mathcal{J}$ consisting of a regularizer $\mathcal{S}$ and a similarity measure $\mathcal{D}$. For $\mathcal{D}$ the sum of squared differences is chosen, since monomodal images are used. The regularizer is based on the Navier-Lamé equation. Taking the Gâteaux-derivative of $\mathcal{J}$ results in the following system of non-linear partial differential equations

$$\mu \Delta u + (\mu + \lambda)\nabla \nabla \cdot u = \nabla T_u (R - T_u), \qquad (1)$$

where R and T_u stand for the reference and the deformed template image, u is the deformation vector field and μ and λ are material properties appearing in the Navier-Lamé equation that characterize the elastic behavior.

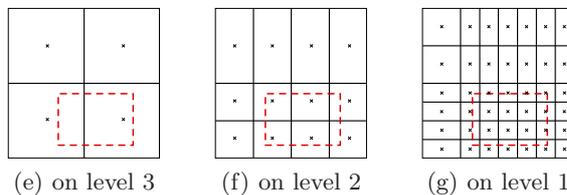(e) on level 3   (f) on level 2   (g) on level 1

Fig. 1. Tensorgrid pyramid with three levels. In this example only two levels of resolution are allowed. The ROI chosen for the refinement is shown as a dashed rectangle.

Eq. 1 is discretized by a finite difference scheme with Neumann boundary conditions, which will be explained in more detail in the next section. The discretized system is solved using a time-marching algorithm and the method of conjugate gradients (CG). The method applies a multi-level strategy combined with an affine pre-registration.

2.3 Tensor grid

The tensor grid of a three dimensional image is defined by vectors x_1, x_2, x_3 containing the coordinates of the voxel centers in the three dimensions. The grid is constructed by forming the tensor product of these vectors.

For the multi-level setting the tensor grid is built by starting on an equidistant grid on the coarsest level and bisecting the grid elements on each level where the grid lines pass through one of the ROIs. Depending on the application the tensor grid can be further refined. Since we are aiming for a rather good deformation of the whole image, only three consecutive levels of resolution are allowed in one image, thus on higher levels the regions outside the ROI will be refined as well. An example of a tensor grid for a multi-level setting can be seen in Fig. 1.

Eq. 1 is discretized on this grid using standard stencils for the elastic regularizer. Note that since the grid is not equidistant the grid spacing has to be included in the stencil, so that neither the stencil nor the resulting matrix is symmetric. However, this matrix can be symmetrized by multiplying each row with a factor that depends on the grid spacing, so that Eq. 1 can still be solved with the CG-method.

3 Results

The maximum-inhale and -exhale phase (i.e. phase1 and phase7) were chosen for a first demonstration of the algorithm. In regions with distinctive structures that can still be well distinguished on lower resolution levels we do not expect the tensor grid approach to be superior to the standard approach. The ROI was therefore placed in the lower left lung enclosing six of the landmarks. A coronal view of phase7 together with the landmarks and the ROI can be seen in Fig. 2.
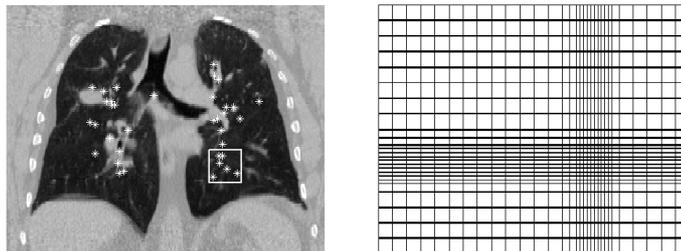

Fig. 2. Coronal view of the lung together with the ROI and the landmarks (left). Same view of the tensor grid used for the computations on level 3 (right).

Table 1. Landmark error in mm before and after registration using the two different grids. The values are in the order: mean±std (max).

landmarks	before registration	using equidistant grids	using tensor grids
all (37)	5.7±2.5 (14.0)	1.8±1.2 (6.0)	1.7±1.3 (5.6)
in ROI (6)	6.4±1.9 (8.1)	1.3±0.5 (1.9)	0.9±0.6 (1.7)
not in ROI (31)	5.5±2.6 (14.0)	1.9±1.3 (6.0)	1.8±1.3 (5.6)

The registration of the two phases was run twice, once with the tensor grid approach and once with the same implementation but using equidistant grids. The computation using the tensor grid was performed up to level 1 (image size: 352×288×128), while the equidistant grid stopped on level 3 (image size: 88×72×64). For the computations an Intel Dual Core with 2.4GHz and 2GB RAM was used. The algorithms were implemented in Matlab and are not yet optimized for speed.

The landmark error before and after registration using the tensor grid (tg) and the equidistant grid (eg) can be seen in Table 1. The error outside of the ROI is approximately the same for both grids since the coarse region of the tensor grid on level 1 has the same resolution as the equidistant grid on level 3. Inside the ROI the landmark error for the tensor grid is smaller due to the higher resolution.

Using the tensor grid approach the image size is reduced by up to 95%. This allows for computations on higher levels as can be seen in Table 2.

The ventilation V is evaluated using the Jacobian of the deformation [1]:

$$V(x) = \det(\nabla(x + u(x))) - 1. \tag{2}$$

A comparison of the ventilation images obtained with the registration using the two grids is given in Fig. 3. The tensor grid yields a result with more details within the ROI and its tails, which have been missed by the equidistant grid.
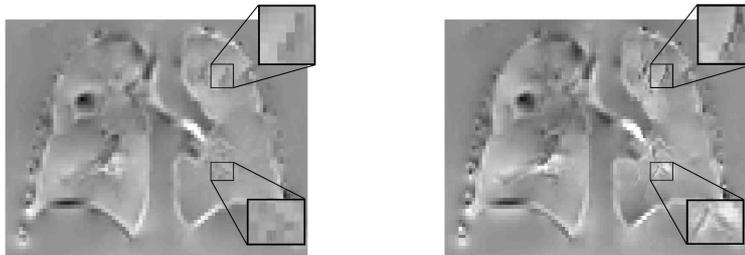

Fig. 3. Ventilation images computed from the registration result using the equidistant grid (left) and the tensor grid (right). Medium gray corresponds to volume preservation, while dark and light gray colors indicate contraction and expansion, respectively.

Table 2. Comparison of the processing time in minutes for each level given the number of voxels and the number of iterations. Level 1 is the original resolution.

	eg:	#voxels	#iterations	time	tg:	#voxels	#iterations	time
level 1		12976128	-	-		943056	100	86.80
level 2		3244032	-	-		164395	100	13.92
level 3		405504	100	9.91		23250	100	1.90
level 4		50688	92	0.65		4046	97	0.26

4 Discussion

In this paper we present an approach using tensor grids for non-parametric image registration. The advantage of this approach over equidistant grids is the possibility to use the highest resolution in user chosen regions of interest 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. Considering the landmark error it turns out that the algorithms using the two different grids perform similar in the regions outside of the ROI while the tensor grid is reaching a better result within the ROI.

The tensor grid is especially useful for ventilation studies where an accurate registration of the complete lung is needed. Due to the size of the images the highest level of resolution is often omitted and thus details in diseased regions are easily lost. Furthermore, the tensor grid approach can be used for any application where one is interested in both a good registration of the complete image and certain regions in detail. It is also possible to define several ROIs in one image, which will be processed at once.

This first result of the new approach is very promising and will be followed by further tests using different datasets to proof its advantage. Moreover, more equally distributed landmarks in the lung need to be identified to allow a profound testing of the algorithm using different ROIs.

References

1. Kabus S, von Berg J, Yamamoto T, et al. Lung ventilation estimation based on 4D-CT imaging. Proc MICCAI Workshop: Pulmonary Image Analysis. 2008; p. 73–81.
2. Cook TS, Tustison N, Biederer J, et al. How do registration parameters affect quantitation of lung kinematics. In: Proc MICCAI; 2007. p. 817–824.
3. Haber E, Heldmann S, Modersitzki J. Adaptive mesh refinement for non-parametric image registration. SIAM J Sci Comp. 2007;Submitted.
4. Papenberg N, Modersitzki J, Fischer B. Registrierung im Fokus. Proc BVM. 2008; p. 138–142.
5. Vandemeulebroucke J, Sarrut D, Clarysse P. The POPI-model, a point-validated pixel-based breathing thorax model; 2007. ICCR, Toronto, Canada.
6. Modersitzki J. Numerical Methods for Image Registration. OUP; 2004.