SIFT and Shape Context for Feature-Based Nonlinear Registration of Thoracic CT Images*

Martin Urschler[1], Joachim Bauer[2], Hendrik Ditt[3], and Horst Bischof[1]

[1] Institute for Computer Graphics & Vision, Graz University of Technology, Austria
{urschler, bischof}@icg.tu-graz.ac.at
[2] VRVis Research Centre, Graz, Austria
bauer@vrvis.at
[3] Siemens Medical Solutions, CTE PA, Forchheim, Germany
hendrik.ditt@siemens.com

Abstract. Nonlinear image registration is a prerequisite for various medical image analysis applications. Many data acquisition protocols suffer from problems due to breathing motion which has to be taken into account for further analysis. Intensity based nonlinear registration is often used to align differing images, however this requires a large computational effort, is sensitive to intensity variations and has problems with matching small structures. In this work a feature-based image registration method is proposed that combines runtime efficiency with good registration accuracy by making use of a fully automatic feature matching and registration approach. The algorithm stages are 3D corner detection, calculation of local (*SIFT*) and global (*Shape Context*) 3D descriptors, robust feature matching and calculation of a dense displacement field. An evaluation of the algorithm on seven synthetic and four clinical data sets is presented. The quantitative and qualitative evaluations show lower runtime and superior results when compared to the Demons algorithm.

1 Introduction

Many medical image analysis applications require a nonlinear (deformable) registration of data sets acquired at different points in time. Especially when dealing with soft tissue organs, like lung or liver during breathing, there are almost always motion differences that have to be compensated for further analysis. The contributions in this paper focus on thoracic CT images coming from CT angiography (CTA) studies for clinical diagnosis of pulmonary embolism [1]. Nonlinear registration is necessary to guarantee that the same anatomical regions are subtracted from each other since patients in bad condition often have problems holding their breath. Despite the focus on thoracic CTA studies the developed algorithm is in principle suitable for other applications as well.

The registration literature distinguishes intensity- and feature-based nonlinear registration methods. Surveys on nonlinear registration methods in medical imaging can be found in Maintz and Viergever [2] or Zitova and Flusser [3]. Often

* This work was funded by Siemens MED CT, Forchheim and Siemens PSE AS, Graz.

R.R. Beichel and M. Sonka (Eds.): CVAMIA 2006, LNCS 4241, pp. 73–84, 2006.

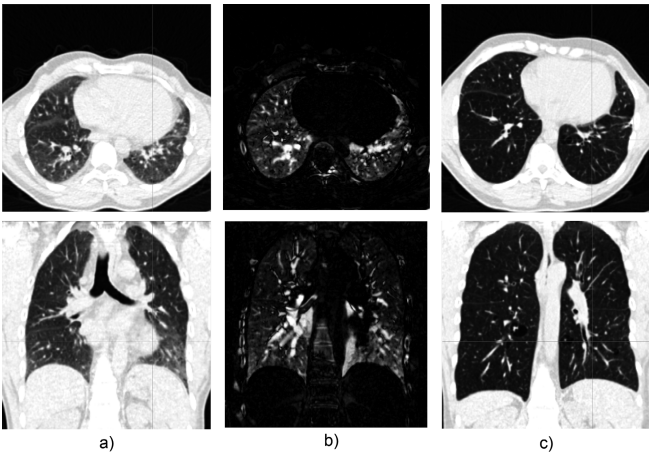

Fig. 1. Evaluation data set F with axial slices in top and sagittal slices in bottom row. a) and c) show the differences in inspiration and exspiration. b) gives the difference image after Demons registration. Note the misregistered vessel structures.

feature-based methods are more accurate than intensity-based methods as long as the feature extraction or segmentation steps are reliable and accurate. Due to the reduction of the problem space, feature-based methods are also significantly faster to compute. On the other hand, segmentation of the organs of interest is not always an easy task and inaccuracies in the segmentation or feature extraction process have severe effects on a subsequent registration step, making the intensity based methods perform better in many practical applications.

This paper presents a novel nonlinear registration approach based on automatically extracted and matched feature points. Although intensity-based approaches are getting more attention by the research community, they face two kinds of practical problems when applied to large thoracic data sets. First, due to their mathematical complexity they require large computational effort. Second, those approaches that are computationally feasible often tend to misregister small structures in the lung like vessels and airways. Further, intensity variations that occur when comparing inhaled and exhaled lungs are not modeled due to implicit brightness constancy assumptions. Fig. 1 shows a nonlinear registration example on lung CT data where the widely-used Demons algorithm [4] leads to misregistrations of vessel structures. Starting from this problem, an automatic feature matching and registration pipeline was established using state of the art techniques from the computer vision community. This pipeline contains Foerstner corner detection [5], forward-backward matching using a 3D scale invariant feature transform (SIFT) descriptor [6] and a global descriptor similar to shape context [7] and dense displacement field estimation in the thin-plate spline (TPS) framework [8]. Especially SIFT and shape context have proved to be very powerful approaches in traditional computer vision applications like wide-baseline matching or object recognition. The main contributions of this work are the 3D extension and the runtime optimization of these stages and

their application to medical images. Related work on feature-based registration was presented in Rohr [9] showing an approximating thin-plate spline registration using manually defined landmarks. In Johnson and Christensen [10] a combined landmark and intensity based approach that establishes a consistent deformation field was proposed. Chui et al. [11] have shown a unified nonlinear feature registration approach using a joint clustering and matching framework. Note that none of these works addresses the problem of fully automatic feature extraction, matching and registration.

2 Methods

Breathing motion mainly stems from two sources, the diaphragm and the rib cage muscles. Expected tissue deformations are not extremely large even in the case of matching full exhalation to full inhalation and they change smoothly over the image domain. These considerations imply that a robust and reproducible feature extraction step producing large numbers of feature candidates followed by the automatic matching of feature descriptors is a valid approach to find corresponding structures in the images (see Fig. 2 for the matching and registration pipeline). Due to the large similarity of local neighborhoods in lung images it is important to not only look at local feature descriptors but also add a notion of global correspondence. Mortensen et al. [12] have recently proposed a combined local and global descriptor for the matching of repetitive patterns. Their ideas were adapted to solve the ambiguities with locally similar structures. After establishing sparse corresponding features a dense displacement field has to be calculated. Bookstein [8] motivated the thin-plate spline (TPS) framework as the appropriate way of displacement field interpolation. However, the interpolating behavior of TPS is not desirable since it may lead to foldings. To overcome this problem the decision to use TPS approximation [9] was made.

The large size of current routinely acquired CT volume data poses runtime and memory restrictions on practically useful algorithms. Acquired CT data

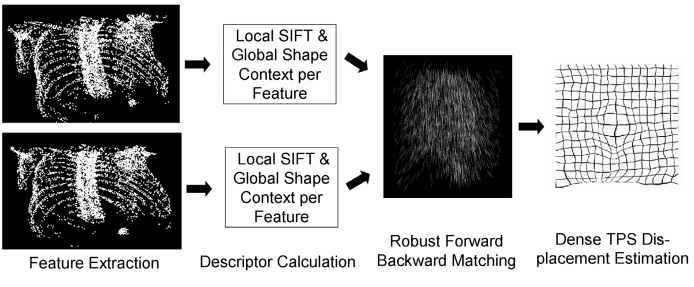

Fig. 2. Nonlinear matching and registration pipeline. The feature extraction stage only shows extracted bone corners, while the method also extracts lung and tissue features.

sets easily require hundreds of MB in memory, so it is necessary to consider computational and memory effort when designing algorithms. In the presented pipeline two performance bottlenecks were identified due to the large number of detected feature points. First, the calculation of the shape context descriptor is critical due to its internally used log-polar histogram bin structure. Therefore an approximation using axis-aligned histogram bins was developed. Second, the calculation of the final dense displacement field is very expensive when confronted with a large number of matched points. Consequently, the global TPS transform was replaced with a k-d tree based locally restricted TPS (LRTPS). The following subsections describe the different pipeline stages in more detail.

2.1 Feature Extraction

The first step in the nonlinear registration pipeline is fast and reproducible feature extraction. This was already extensively investigated by Rohr [9]. His evaluations of several different 3D anatomical feature detection operators resulted in the recommendation of the structure tensor based 3D Foerstner operator [5].

2.2 Local SIFT Feature Descriptor

For each detected feature a distinctive local SIFT feature descriptor [6] is built. SIFT descriptors are robust to local deformations and to errors in feature detection. Performance evaluations show its excellent matching behavior on various kinds of transformations [13]. In this work only the SIFT descriptor representation is used, since keypoint localization is performed using Foerstner corners. To apply the SIFT descriptor representation on volume data an extension to 3D is necessary. The 3D SIFT descriptor quantizes gradient locations in a 2x2x2 grid while gradient orientations are quantized into two 8-bin orientation histograms. Each 512-dimensional descriptor is normalized by its L_2-norm. In contrast to the 2D SIFT formulation, the proposed 3D SIFT descriptor is not rotation-invariant, since this saves computation time and breathing motion is assumed not to lead to strong rotation-like local deformations.

2.3 Global Shape Context Feature Descriptor

The 3D shape context descriptor [7] assumes that objects are captured by point sets $P = \{p_1, ..., p_n\}$ obtained from a feature detector or as locations of edges from an edge detector. If one looks at the set of vectors emitted from one point p_k to all other points p_i of a shape with $i \neq k$, this set can be interpreted as a rich description of the shape configuration relative to p_i. The *relative distribution* of this set of vectors is used as a compact, yet highly discriminative histogram descriptor. This histogram uses bins that are uniform in a 3D spherical coordinate system (θ,ϕ,r). The r coordinate axis is logarithmically scaled, so positions of nearby sample points have stronger influence on the descriptor. The log-polar histogram binning of this method was identified as a performance problem when

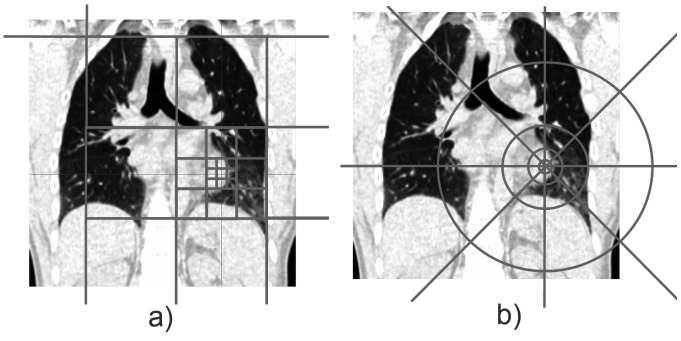

Fig. 3. Approximated global context (a) vs shape context (b) histogram bin structure

applied to a large number of extracted feature points. Therefore an approximation of the shape context descriptor was used, which offers good matching results compared to the classical approach. Here the log-polar histogram bin structure is replaced with a bin structure based on rectangular, axis-aligned image patches. Fig. 3 compares the equivalent 2D histogram bin structures of the classical and the approximated descriptors. The patch size increases exponentially with growing distance from the feature point. This strategy preserves local information close to the feature point and generalizes at larger distance to a coarser quantization. If one uses a 3D integral image representation to store feature point locations, this descriptor is extremely efficient to compute due to its axis-aligned bin structure. The integral image allows to count feature points per rectangular patch in constant time, so for each feature point the histogram is calculated in logarithmic instead of linear time.

2.4 Robust Feature Matching

To find corresponding feature points a matching algorithm has to be used. The previous stages have established a local and a global descriptor for each feature point, the task of the matching stage now is to find those point pairs from two volumes that minimize a cost function derived from the descriptors. Forward-backward matching, a simple but robust approach in terms of consistency, occlusions and erroneous feature extraction, was presented by Fua [14] on stereo matching problems. The basic idea is to perform the matching step twice by reversing the roles of the two volumes V_1, V_2 and considering only those matches as valid for which the corresponding points $P_{i,1}$ and $P_{i,2}$ are identical when matching from V_1 to V_2 and from V_2 to V_1. The cost function used in the two matching steps is a weighted linear combination of distance metrics. The distance metric d_{SIFT} for the SIFT feature descriptor is the Euclidean distance in the 512-dimensional feature space. Shape context descriptors SC_i and SC_j are compared using a χ^2 statistic $d_{SC} = \chi^2 = \frac{1}{2}\sum_k \frac{(SC_{i,k}-SC_{j,k})^2}{SC_{i,k}+SC_{j,k}}$. Both distance metrics are normalized between 0 and 1. The total cost function is given by

$d = \omega d_{SIFT} + (1 - \omega) d_{SC}$ where ω is a weighting factor. Matches with a cost function value above some user-defined threshold T_d are discarded.

2.5 Dense Displacement Field Interpolation

The final step in the registration pipeline is the estimation of a dense displacement field from the sparse matching result. For this purpose a TPS interpolation is used [8,9]. In its original formulation the interpolating behavior of the TPS often is too restrictive and might lead to overfitting to the correspondences or folding of the displacement field in case of erroneous correspondences. In this work the findings of Rohr [9] for approximating TPS mappings were considered. A regularization term is added to the formulation, which is steered by a parameter λ, weighting the tradeoff between interpolation and smoothness of the solution. Further, a landmark error term is introduced to give each pair of corresponding features an uncertainty measure, which is directly derived from the matching costs of the feature matching stage.

It is desirable that the matching stage produces a large number of feature correspondences n_{match}. On typical volume data sets thousands of correspondences might be achieved. In this case the warping of large volume data sets is very costly, since for each voxel a multiplication with $O(n_{match})$ weighted landmarks including the calculation of $O(n_{match})$ vector norms is involved. Therefore a locally restricted version (LRTPS) of the normally global TPS transform is used. The source points of the feature correspondences are put into a k-d tree structure to give efficient access to its neighbor features. For each feature point a TPS transform restricted to a pre-defined number of neighbors is calculated. So only a subset of the total set of correspondences in the volume is taken for locally estimating the transform. The dense displacement field approximation step always looks for the nearest feature correspondence in the k-d tree and takes the stored local TPS transform to compute a displacement.

3 Experiments and Results

To assess the validity of the feature-based registration approach qualitative and quantitative evaluations were performed on synthetically transformed and clinical thorax CT data sets. For the synthetic deformation experiments two different kinds of deformations were used. The first deformation model is a *Simulated Breathing Transformation* simulating rib-cage and diaphragm muscle behavior. The second synthetic transformation makes use of evenly distributed landmarks that are moved in random directions. The synthetic experiments give numbers on the RMS of the intensity differences before and after registration, compares the registered and the synthetic displacement fields and compares the method with the Demons algorithm. Real data experiments show the decrease in the RMS of the intensity differences, compare the RMS with the Demons algorithm and give qualitative difference images. All experiments were performed on a dual 2GHz AMD Opteron system with 8GB RAM running Linux.

3.1 Synthetic Deformation Experiments

Synthetic experiments were performed on seven test data sets (A,B,C,D,E,F,G) taken at inspiration, each of them having a volume size of 512 by 512 by 256 voxels. The first synthetic transformation intends to model breathing behavior. The nonlinear transformation $\boldsymbol{d}(x, y, z) : \mathcal{R}^3 \rightarrow \mathcal{R}^3$ simulates diaphragm and rib cage movement. Diaphragm movement is applied as a translational force in the data sets negative z direction. Nonlinearity is introduced by weighting the constant vertical translation $t_{vertical}$ with a two-dimensional Gaussian distribution that depends on the x and y coordinates of the data set. Mathematically a displacement vector $\boldsymbol{d_1} = (0, 0, z')^T$ is applied to each point $(x, y, z)^T$ that maps it to $(x, y, z')^T$ with

$$z' = z - t_{vertical} e^{-\frac{(x-\mu_x)^2+(y-\mu_y)^2}{2\sigma^2}}$$

where (μ_x, μ_y) is the center of gravity of the diaphragm points and σ is chosen such that points lying at the exterior of the diaphragm surface nearly remain fixed. In a similar fashion, simulation of rib cage behavior during breathing leads to a radial, center-directed translation $t_{inplane}$. It is used to form a second displacement $\boldsymbol{d_2} = (x', y', 0)^T$ that maps points $(x, y, z)^T$ to $(x', y', z)^T$ with

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \mu_x \\ \mu_y \end{pmatrix} + \left(\left| \begin{pmatrix} x-\mu_x \\ y-\mu_y \end{pmatrix} \right| - t_{inplane} * (1 - e^{-\frac{(x-\mu_x)^2+(y-\mu_y)^2}{2\sigma^2}}) \right) * \frac{\begin{pmatrix} x-\mu_x \\ y-\mu_y \end{pmatrix}}{\left| \begin{pmatrix} x-\mu_x \\ y-\mu_y \end{pmatrix} \right|}$$

Combining displacements $\boldsymbol{d_1}$ and $\boldsymbol{d_2}$ results in transformation $\boldsymbol{d}$. Finally, to simulate a change in lung gray-values due to inhalation, all gray values smaller than -800 Hounsfield Units (HU) are increased by a random number drawn from a normal distribution centered at 25 HU with a standard deviation of 3 HU.

The seven test data sets were synthetically transformed with a small and a large deformation. The small deformation is defined by the translations $t_{vertical} =$ 25mm and $t_{inplane} =$ 10mm while the large deformation is defined by $t_{vertical} =$ 55mm and $t_{inplane} =$ 25mm. Fig. 4 a)-c) shows the effects of these deformations on data set A. First, the feature matching produces corresponding points which can be compared to the ground truth simulated breathing transformation in terms of the RMS of the displacement difference vectors (RMS_{disp}) and the maximum of the lengths of the displacement difference vectors (MAX_{disp}) over all correspondences. Evaluations showed that the RMS_{disp} varies between 0.265mm and 0.314mm for the small and between 0.558mm and 2.479mm for the large deformation over the data sets. Accordingly MAX_{DISP} varies between 2.62mm and 8.59mm for small and between 14.95mm and 28.37mm for large deformations respectively.

Table 1 gives the results of the synthetic registration experiments. All comparisons are always performed only on those regions which are present in both registered data sets. The RMS of the intensity differences before ($RMSD_{initial}$) and after ($RMSD_{feature}$) registration are calculated, as well as the difference of

Fig. 4. Synthetic transformations. a) original data set A, b) the small and c) the large simulated breathing deformation. d) original data set B, e) randomly displaced landmark transformation -8...+8, f) displacement -24...+24, g) displacement -48...+48.

Table 1. Simulated breathing transformation. Registration results in terms of RMS intensity differences and displacement difference vectors.

	Measure		A	B	C	D	E	F	G	Mean
SimBreath 25-10	$RMSD_{initial}$	[HU]	385.99	327.25	303.92	359.49	318.73	316.58	316.45	332.63
	$RMSD_{demons}$	[HU]	114.44	115.56	92.05	100.06	96.33	90.21	105.95	102.08
	$RMSD_{feature}$	[HU]	45.18	45.71	43.43	50.78	46.09	41.64	47.91	45.82
	$RMS_{disp,demons}$	[mm]	4.564	5.961	5.412	4.842	4.756	4.931	5.741	5.172
	$RMS_{disp,feature}$	[mm]	0.662	0.769	0.59	0.66	0.601	1.338	1.025	0.806
	$MAX_{disp,demons}$	[mm]	34.55	38.56	38.06	29.95	28.45	35.26	37.59	34.63
	$MAX_{disp,feature}$	[mm]	8.59	9.56	8.02	8.52	8.87	12.64	11.79	9.71
	# matches		2678	2121	2330	2022	1805	5204	3714	2825.9
SimBreath 55-25	$RMSD_{initial}$	[HU]	549.02	477.81	446.84	521.65	477.07	472.65	442.95	484
	$RMSD_{demons}$	[HU]	134.46	135.69	128.57	181.32	145.77	153.67	153.11	147.51
	$RMSD_{feature}$	[HU]	69.35	82.24	71.09	96.98	70.11	74.78	78.93	77.64
	$RMS_{disp,demons}$	[mm]	6.844	7.384	6.912	5.822	5.113	7.012	6.992	6.583
	$RMS_{disp,feature}$	[mm]	1.059	1.382	1.327	1.331	1.252	2.256	2.03	1.519
	$MAX_{disp,demons}$	[mm]	39.45	38.12	39.99	41.72	38.09	43.01	42.95	40.48
	$MAX_{disp,feature}$	[mm]	15.22	21.44	19.84	23.49	18.24	24.14	21.66	20.58
	# matches		1940	1424	1709	1287	1277	2778	2856	1895.9

Table 2. Synthetic TPS transformation results. Registration results in terms of RMS intensity differences and displacement difference vectors.

Measure		-8...+8	-16...+16	-24...+24	-32...+32	-48...+48
$RMSD_{initial}$	[HU]	164.89	233.48	289.76	333.94	397.89
$RMSD_{demons}$	[HU]	161.13	188.67	193.75	184.12	253.81
$RMSD_{feature}$	[HU]	90.58	169.89	236.83	271.113	360.72
$RMS_{disp,demons}$	[mm]	5.834	9.374	13.874	17.099	19.933
$RMS_{disp,feature}$	[mm]	4.802	8.355	11.172	15.792	18.562
# matches		1729	1101	426	312	87

the resulting and the synthetic displacement fields in terms of the RMS of the displacement vector field RMS_{disp} and the maximum of the lengths of the displacement difference vectors MAX_{disp}. The algorithms performance is compared to the widely used Demons [4] algorithm. Its implementation was taken from the Insight Segmentation and Registration Toolkit [1]. The Demons algorithm uses a five level multi-resolution framework to calculate a smooth displacement field with a fixed number of iterations per multi-resolution level (between 100 and 15 from coarse to fine) in a gradient-descent scheme. Fig. 5 a)-c) shows difference images of data set D which had the worst behavior in terms of decreasing the RMS of the intensity differences using the large simulated breathing deformation.

The second synthetic transformation is calculated using a number of evenly distributed landmark points and randomly assigning displacements to these landmarks. The amount of the displacement is increased up to the sampling size of the landmark distribution. These displacements are not physically motivated and the larger the assigned displacements are the harder it is to correct them. The dense synthetic displacement field is calculated using a TPS interpolation. Note that we use TPS interpolation here not TPS approximation. For data sets size 512x512x256 every 64 voxels a landmark is placed in the original image. This leads to a grid of 7x7x3 landmarks. Now a random displacement is calculated for each landmark coordinate in the range from -8 to +8 voxels. This is repeated five times while always doubling the displacement range. Evaluations are solely performed on data set B, Fig. 4 d)-g) shows the effect of these synthetic displacements. The main motivation of these experiments is to determine the degree of deformation where the algorithms are not capable to register the data anymore. Table 2 gives the results of these experiments.

3.2 Clinical Data

The algorithm was also evaluated on four clinical thoracic data sets consisting of two scans at different breathing states. The data sets show different problem characteristics. Data sets B and G differ by a small breathing deformation and intensity variations due to contrast agent application. Data set G additionally shows a lung disease in the upper lobe region. Data sets E and F differ by large breathing deformations and the images have intensity differences due to a lung disease making them very hard to register. For the clinical data no gold standard displacement was available for comparison, therefore solely the decrease in the RMS of the intensity differences before and after registration are calculated. Again the novel feature-based algorithm is compared with the Demons algorithm. The RMS of data set B was decreased from 201.57HU to 129HU (Demons) and to 104.83HU (feature-based). Data set E decreased from 403.49HU to 235.88HU and to 197.45HU, data set F from 413.62HU to 288.31HU and to 294.98HU and finally data set G from 367.66HU to 274.14HU and to 241.43HU. The numbers of found correspondences lies between 685 and 1632. For qualitative results difference images are shown in Fig. 5 d)-f) for data set B and g)-i) for data set E.

[1] http://www.itk.org

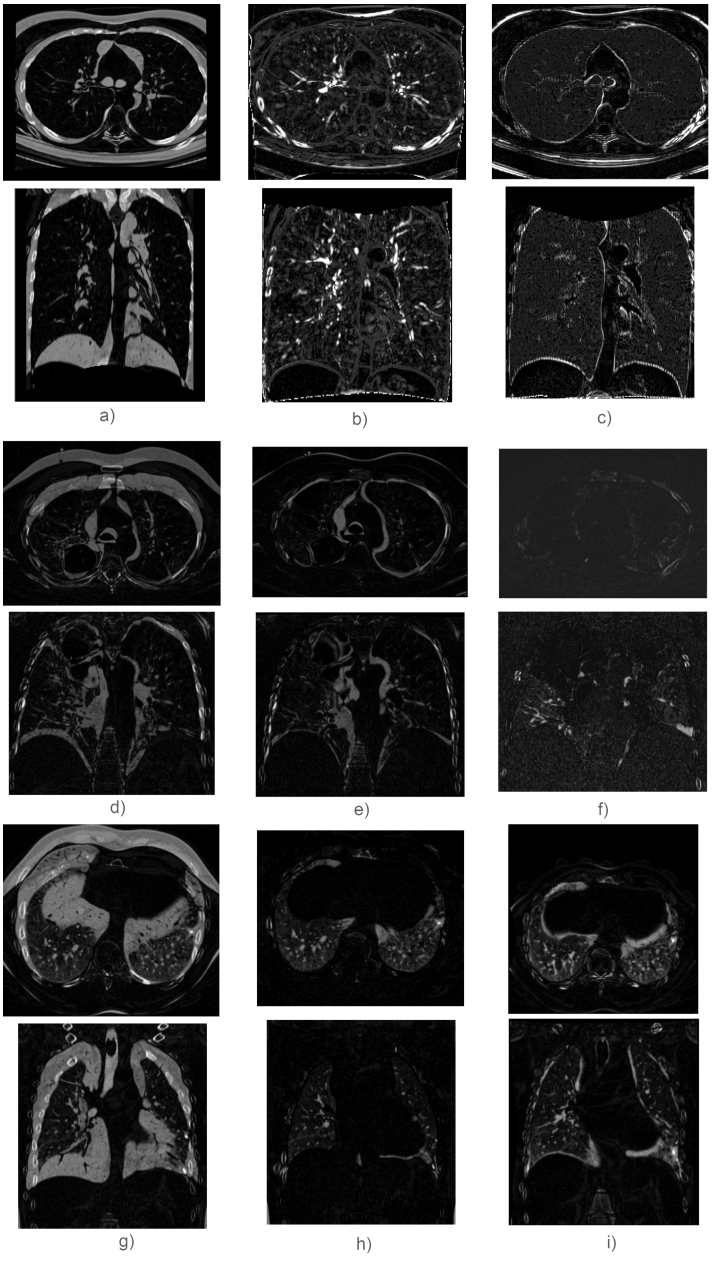

Fig. 5. Selected Results. Top row always axial, bottom row always sagittal slices. Left image shows difference image before, middle image after Demons and right image after feature-based registration. a)-c) shows synthetic results of data set D with a simulated breathing transformation of $t_{vertical} = 55$mm and $t_{inplane} = 25$mm. d)-f) shows results on real data set B, g)-i) on real data set E.

4 Discussion and Outlook

The different stages of the proposed feature-based algorithm require some parameters to be chosen. In the matching stage normalized local and global descriptors are used to form the matching cost function. A meaningful threshold T_d has to be found to exclude bad matches which was empirically chosen between 0.25 and 0.35. The matching stage also produces some outliers. The MAX_{disp} measures reflect this fact, especially in the matching evaluation. However, the registration stage with the approximating TPS framework takes the magnitude of the matching cost into account, such that outlier matches have a low influence on the final dense displacement field. Parameter λ of the TPS displacement field approximation was selected as $\lambda = 0.005$ after experimenting with several data sets. The LRTPS implementation needs a choice on the number of neighboring points that defines a local TPS, this parameter was set to 150.

The two goals of the proposed algorithm, to be faster and at least as accurate as a state-of-the-art nonlinear intensity-based registration algorithm, were both met. Computation time of the feature-based algorithm on the 512x512x256 data sets ranges between 1632s and 2282s, depending on the number of identified correspondences. The largest part (more than 50%) of the algorithm runtime still goes into the calculation of the dense displacement field. The Demons algorithm takes on average around 2540s for registration of two 512x512x256 data sets. If one further increases data sets size a feature-based algorithm will be even more efficient due to its inherent reduction of matching complexity.

Registration accuracy of the feature-based algorithm exceeds the Demons algorithm in most of the synthetic examples. Demons only performs better on the evenly distributed landmark TPS experiments with a high degree of random deformation where the feature-based approach is not able to find enough correspondences. However, these deformations are physically implausible and not representative. Especially the simulated breathing transformation was very accurately registered using the feature-based approach. This is reflected in the substantial decrease of the RMS of the intensity difference and the RMS of the displacement difference vector fields (in the order of 0.5mm to 2.0mm). The difference images of the simulated breathing experiment (Fig. 5) illustrate the problems of the Demons approach with the vascular structures. The real data experiments show that the performance of the feature-based algorithm is comparable to Demons. Performance on data sets B and G was better, while the performance on the very difficult data sets E and F is similar. Although Demons shows lower RMS values, the difference images of the feature-based approach have the same quality. However, the feature-based approach also has some problems with registration of vessel structures on the difficult data sets. The largest disadvantage of the feature-based approach is to get a large number of robust, evenly distributed feature matches. This can not be guaranteed in the current implementation, which explains the registration problems and can also lead to artifacts at the edges of the local TPS regions.

Future work will investigate methods to further speed-up the algorithm in its time-consuming stages. Another important point is to find a way to guarantee

a good distribution of matching features. For evaluation a comparison with a more elaborate intensity-based registration approach will be performed. The fusion of this algorithm with a suitable intensity-based method seems to be a very promising direction, since the fast feature-based matching should provide a very good initial condition.

References

1. Wildberger, J.E., Klotz, E., Ditt, H., Mahnken, A.H., Spüntrup, E., Günther, R.W.: Multi-slice CT for Visualization of Acute Pulmonary Embolism: Single Breath-hold Subtraction Technique. Fortschritte auf dem Gebiet der Roentgenstrahlen und der bildgebenden Verfahren **177** (2005) 17–23
2. Maintz, J., Viergever, M.: A Survey of Medical Image Registration. Medical Image Analysis **2**(1) (1998) 1–36
3. Zitova, B., Flusser, J.: Image registration methods: A survey. Image and Vision Computing **21**(11) (2003) 977–1000
4. Thirion, J.P.: Image matching as a diffusion process: An analogy with Maxwell's demons. Medical Image Analysis **2**(3) (1998) 243–260
5. Förstner, W.: A Feature Based Correspondence Algorithm for Image Matching. Int. Arch. of Photogrammetry and Remote Sensing **26**(3) (1986) 150–166
6. Lowe, D.G.: Distinctive Image Features from Scale-Invariant Keypoints. International Journal on Computer Vision **60**(2) (2004) 91–110
7. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. IEEE PAMI **24**(4) (2002) 509–522
8. Bookstein, F.: Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. IEEE PAMI **11**(6) (1989) 567–585
9. Rohr, K.: Landmark-Based Image Analysis Using Geometric and Intensity Models. Computational Imaging and Vision. Kluwer Academic Publishers (2001)
10. Johnson, H.J., Christensen, G.E.: Consistent Landmark and Intensity-Based Image Registration. IEEE Transactions on Medical Imaging **21**(5) (2002) 450–461
11. Chui, H., Win, L., Schultz, R., Duncan, J., Rangarajan, A.: A unified non-rigid feature registration method for brain mapping. Medical Image Analysis **7** (2003) 113–130
12. Mortensen, E.N., Deng, H., Shapiro, L.: A SIFT Descriptor with Global Context. In: IEEE CVPR 2005. (2005) 184–190
13. Mikolajczyk, K., Schmid, C.: Performance Evaluation of Local Descriptors. IEEE PAMI **27**(10) (2005) 1615–1630
14. Fua, P.: A Parallel Stereo Algorithm that Produces Dense Depth Maps and Preserves Image Features. Machine Vision and Applications **6** (1993) 35–49