Feature-constrained Nonlinear Registration of Lung CT Images

Xiao Han

Elekta Inc., 13723 Riverport Drive, Maryland Heights, MO 63043, USA
Xiao.Han@Elekta.com

Abstract. Deformable image registration is a key enabling technology for advanced treatment of lung cancer patients, as it can facilitate motion estimation, structure segmentation, as well as dose tracking and accumulation. In this work, we developed a hybrid feature-constrained deformable registration method and applied it to tackle the EMPIRE10 (*E*valuation of *M*ethods for *P*ulmonary *I*mage *R*egistration 2010) lung image registration challenge. The method uses the results of feature detection and matching based on robust 3D SURF descriptors to guide an intensity-based deformable image registration. The initial registration result is further refined by a hybrid MI/NSSD deformable registration process. The proposed method is fully automatic and does not require pre-segmentation of any lung structures. Validation results on the EMPIRE10 data showed that our method performed very well among 34 competing algorithms. Future improvement is possible with adaptive parameter selection, site-specific feature detection methods, and better deformation models.

1 Introduction

Lung cancer is the leading cause of cancer death worldwide, with deaths exceeding 1 million cases every year [1]. Radiation therapy (RT) plays an important role in both the curative and palliative treatment of lung cancer patients. More advanced RT techniques, such as 4D or adaptive radiotherapy, offer even better tumor control by taking into account respiration-induced lung and tumor motion as well as physiological and geometric changes of tumors and normal anatomy over the course of treatment.

Deformable or nonlinear image registration has been recognized as a key enabling technology for the successful implementation of adaptive RT [2]. By definition, image registration is the process of establishing spatial correspondences between two images. The correspondence information can be used to characterize patient-specific tumor or lung motion, thus reducing uncertainties related to target volume definition. The estimation of respiratory motion itself is also beneficial for assessing pulmonary functions [3]. The anatomical correspondence can also be used to map daily treatment dose to a reference frame, allowing cumulative dose to be computed and compared against original treatment plan.

A large variety of nonlinear image registration methods have appeared in the literature, and many have been applied for the registration of lung CT images (cf. [3–9] and references therein). In general, existing registration methods can be broadly classified into three categories: feature-based methods, intensity-based methods, and hybrid methods that combine the previous two. Feature-based methods rely on the detection and matching of isolated image features such as points, curves, or surface patches. Such methods are usually more efficient to compute, but an interpolation method is needed to get the image correspondences at non-feature locations. In addition, the automatic detection and matching of a large set of features for the purpose of detailed nonlinear registration is not an easy task. Intensity-based methods directly use the image intensity values to derive image matching, which can be easily made fully automatic and can exploit the full image information. However, such methods tend to have high computational cost, and they typically require a good starting point for the registration optimization due to the large number of local optima in a typical non-linear image matching function. Consequently, hybrid methods have also been proposed, which aim to integrate the merits of both feature- and intensity-based methods [6, 8].

In this work, we have developed a novel feature-constrained hybrid deformable registration method for lung CT images. Incorporating features have been proven beneficial for thoracic image registration due to the presence of rich anatomical landmarks in the lung region, such as the blood vessel trees and the bronchial airways [10, 6, 8]. But instead of trying to segment out these anatomical structures, we use a more generic feature detection method and rely on a robust feature descriptor recently proposed in the computer vision literature for automatic feature point matching. The feature matching result is then used to help constrain an intensity-based nonlinear image registration procedure. In the following, we first present the proposed method and then summarize our results for the EMPIRE10 data.

2 Methods

The overall method consists of four major steps: a global rigid registration, feature point detection and matching, a feature-constrained mutual-information (MI) deformable registration, and finally another dense deformable registration to further refine the results. The major reason for the two separate deformable registration steps is to deal with the large size of the EMPIRE10 data. In particular, most of the EMPIRE10 data have sub-millimeter resolution and very large image sizes. To improve computation speed and reduce memory usage, we first down-sample the input images to an isotropic, lower resolution of 1.5 mm^3. The first three steps are all performed on the low resolution data. The resulting deformation field is then upsampled to the original resolution and refined at the last deformable registration step. In order to handle large image size at the last step, the full image domain is first divided into 8 slightly overlapping sub-regions (10-voxel margin on each side). Each sub-region is computed separately, but the final deformation field in the overlapping area is taken as the average of results from neighboring subregions. In the next, we detail each of the four steps and explain the parameters we used for processing the EMPIRE10 data.

2.1 Rigid Registration

The rigid registration step aims to correct for global differences in position and orientation between the two input images of each subject. To find the optimal transformation parameters (three rotations and three translations), we maximize the global MI between the two images using a multi-resolution stochastic gradient-descent optimization scheme [11]. This scheme exploits the fact that the rigid registration is an over-constrained problem since the number of transformation parameters is far less than the number of voxels in the image domain. Thus, at each iteration, the stochastic scheme uses only a very small subset of random image samples to estimate the image similarity metric and the transformation parameters, which leads to a highly efficient algorithm that usually takes only a few seconds on a modern desktop computer. As mentioned earlier, this step was performed on down-sampled version (1.5 mm^3 isotropic resolution) of the original images. Three resolution levels were used for the multi-resolution computation. At each resolution level, 100 random samples and 32 histogram bins were used to estimate the MI for the stochastic gradient optimization.

2.2 Feature Point Extraction and Matching

Many existing feature-based lung registration methods require the segmentation of the blood vessel trees or airways [10, 6, 8]. Instead, we follow a more generic approach and use a simple 3D landmark detector for feature point extraction. In particular, we use a structure-tensor-based 3D Förstner operator as recommended in [12]. The same feature detector was also used in [5] for lung CT image registration and in [13] for registering liver images. However, we adopt a different feature descriptor in the feature matching stage, and the overall nonlinear registration method proposed here is also different from the previous work, which will become clear later.

The 3D Förstner operator assigns a cornerness measure, denoted by $F(\mathbf{x})$, for each image voxel $\mathbf{x}$, which can be defined by the following equation:

$$F(\mathbf{x}) = \det(C)/\text{trace}(C^{\text{adj}}), \tag{1}$$

with

$$C = \overline{\nabla I(\nabla I)^T},$$

where C^{adj} denotes the adjugate matrix of C. ∇I is the gradient of the input image I, and $\overline{\nabla I(\nabla I)^T}$ represents the average of the structure tensor $\nabla I(\nabla I)^T$ in a local neighborhood of voxel $\mathbf{x}$. We used a $3 \times 3 \times 3$ neighborhood to compute the local average in this EMPIRE10 study.

After the cornerness measure F is computed for every image voxel, we extract feature points as the subset of image voxels whose F-value is a local maximum within a local neighborhood of size 5 × 5 × 5. For illustration, Fig. 1 shows the detected feature points on (approximately) corresponding slices of two separate images of one subject.

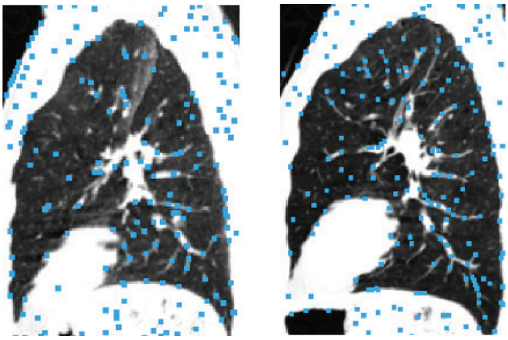

Fig. 1. Illustration of feature point detection. Blue dots show detected feature points on the respective sagittal slices of two images of the same subject.

To perform feature point matching, a feature descriptor is needed in order to assign a unique characteristic for each feature point. The SIFT descriptor introduced by Lowe [14] has become very popular and been widely adopted by the medical imaging community [5,13,15,16]. A major drawback of the SIFT descriptor is its long computation time and large feature size in 3D. In this work we use instead the SURF (*S*peeded *U*p *R*obust *F*eatures) descriptor that is proposed more recently by Bay et al. [17], and extend it to 3D. First, we take a $16 \times 16 \times 16$ neighborhood for each detected feature point, and further divide it into sixty-four $4 \times 4 \times 4$ sub-regions. Second, for each subregion, a six-dimensional description vector $\mathbf{v}$ is computed to characterize the local intensity structure:

$$\mathbf{v} = (\sum g_x, \sum g_y, \sum g_z, \sum |g_x|, \sum |g_y|, \sum |g_z|),$$

where g_x, g_y, and g_z denote the three components of the image gradient vector. The sub-region descriptors are then concatenated together to form a 384-dimensional feature vector, which is further normalized to a unit vector to get the final SURF descriptor for each feature point.

The correspondences between feature points of two input images are established by finding nearest-neighbors (NNs) in the 384-dimensional vector space. Two extra steps are performed to help eliminate false matches as proposed in the original SIFT paper [14]. First, a distance ratio r is computed after the NN search, which is the ratio of the smallest distance value to the second-smallest. A match is rejected if r is greater than a given threshold (0.6 is used in this study). Second, a symmetric criterion is applied, which performs the NN search in two directions and a matched pair $(\mathbf{p}, \mathbf{q})$ is kept only if it satisfies the NN-optimality in both directions.

2.3 Feature-constrained MI Deformable Registration

The matched feature points are usually sparse, thus insufficient to derive a full image correspondence map. Instead, we use the feature-matching results to help constrain an intensity-based image registration method. We adopt the MI as the intensity-based image similarity measure at this step due to its ability to handle intensity variations that often exist between a pair of lung images. The MI of two images I and J measures the degrees of dependence between I and J and can be approximated as (cf. [11]):

$$\mathrm{MI}(I,J) = \frac{1}{N}\sum_{\mathbf{x}} \log\left[\frac{p_{I,J}(I(\mathbf{x}), J(T(\mathbf{x})))}{p_I(I(\mathbf{x}))p_J(J(T(\mathbf{x})))}\right], \tag{2}$$

where N denotes the total number of image points within the overlapped region of the two images and T denotes the image transformation model. The functions $p_{I,J}(\cdot,\cdot)$ and $p_I(\cdot)$, $p_J(\cdot)$ denote the joint and marginal image intensity distributions respectively. A non-parametric transformation model is used at this step where the image transformation is modeled directly as a vectorial displacement field $\mathbf{U}$, such that $T(\mathbf{x}) = \mathbf{x} + \mathbf{U}(\mathbf{x})$ for every image point $\mathbf{x}$.

To incorporate feature-matching constraints, the total objective function for estimating the optimal deformation field is formulated as:

$$\mathcal{J}(I,J,\mathbf{U}) = -\mathrm{MI}(I,J,\mathbf{U}) + \lambda\|\nabla\mathbf{U}\|_2 + \mu\sum_{\mathbf{x}\in S}\|\mathbf{U}(\mathbf{x}) - \mathbf{V}(\mathbf{x})\|^2, \tag{3}$$

where S indicates the set of feature points with established correspondences, and $\mathbf{V}$ denotes the displacements at these points as given by the feature matching results. $\|\nabla\mathbf{U}\|_2$ denotes the L_2 norm of the gradients of $\mathbf{U}$, which enforces a smooth deformation field. The two weighting factors, λ and μ, control the relative strength of each term. Their values are not explicitly set but reflected in the parameter selection of the optimization process as will become clear later.

The optimal deformation field is computed iteratively using a "pair-and-smooth" strategy, which alternates between optimizing the MI similarity metric and satisfying the regularization and feature constraints. In addition, we adopt a compositive update scheme, where the deformation field $\mathbf{U}$ is updated according to the following equation:

$$\mathbf{U}^n = \mathbf{U}^{n-1} \circ (\mathbf{Id} + \mathbf{u}^n) + \mathbf{u}^n; \tag{4}$$

or equivalently,

$$\mathbf{U}^n(\mathbf{x}) = \mathbf{U}^{n-1}(\mathbf{x} + \mathbf{u}^n(\mathbf{x})) + \mathbf{u}^n(\mathbf{x}),\ \forall\mathbf{x}. \tag{5}$$

In Eq. (4), $\mathbf{Id}$ denotes the identity transformation and "$\circ$" denotes transformation composition. $\mathbf{u}^n$ is a local update field at the n-th iteration. A benefit of the compositive update scheme is that the total deformation field $\mathbf{U}$ can be easily guaranteed to be a diffeomorphic mapping. At each iteration, $\mathbf{u}^n$ is computed to maximize the MI between image I and the deformed image $J^n = J \circ (\mathbf{Id} + \mathbf{U}^n)$, i.e., to maximize

$$\begin{aligned}\mathrm{MI}(I, J^n \circ (\mathbf{Id} + \mathbf{u}^n)) &= \frac{1}{N}\sum_{\mathbf{x}} \log\left[\frac{p_{I,J^n}(I(\mathbf{x}), J^n(\mathbf{x}+\mathbf{u}^n(\mathbf{x})))}{p_I(I(\mathbf{x}))p_{J^n}(J^n(\mathbf{x}+\mathbf{u}^n(\mathbf{x})))}\right] \\ &\triangleq \frac{1}{N}\sum_{\mathbf{x}} \mathrm{SMI}(I(\mathbf{x}), J^n(\mathbf{x}+\mathbf{u}^n(\mathbf{x}))). \qquad (6)\end{aligned}$$

In the above equation, $\mathrm{SMI}(\cdot,\cdot)$ denotes a point-wise MI measure (cf. [18]).

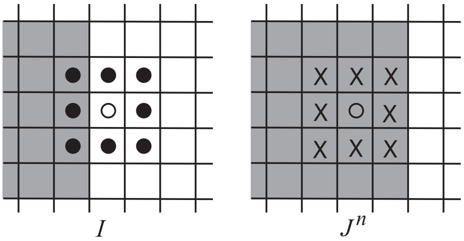

Fig. 2. A 2D illustration of explicit local search. The open circles indicate the image point under consideration and its current correspondence in the other image. The crosses indicate the neighbors involved in the forward search. The solid circles indicate the neighbors involved in the backward search.

As explained in our earlier work [18], the optimal update $\mathbf{u}^n(\mathbf{x})$ at each image location $\mathbf{x}$ can be found by searching the local neighbors of $\mathbf{x}$ in the deformed image J^n and finding the neighbor $\mathbf{x}'$ that maximizes $\mathrm{SMI}(I(\mathbf{x}), J^n(\mathbf{x}' = \mathbf{x} + \mathbf{u}^n(\mathbf{x})))$ (cf. Fig. 2). This simple explicit neighbor search scheme avoids the difficulty in determining the optimal time step size as needed for a traditional gradient-descent optimization approach. It also eliminates the need to compute derivatives of either the images or the image intensity distributions. A symmetric version of the local update scheme was designed in [18] that further improves the convergence rate of the registration method (cf. Fig. 2).

After the update field is found, we smooth it with a spatial Gaussian filter (a fixed kernel width of 1-voxel is used in this study). We then update the total deformation field according to Eq. (4). To satisfy the feature constraint, we simply set $\mathbf{U}^n(\mathbf{x}) = \mathbf{V}(\mathbf{x})$ at the feature point locations. To help eliminate any remaining outliers undetected in the feature matching step, the previous assignment is performed only at points where the original difference $\|\mathbf{U}^n(\mathbf{x}) - \mathbf{V}(\mathbf{x})\|$ is less than a given threshold, for which we have used a fixed value of 50 mm throughout this study. Finally, to ensure smoothness of the deformation field, we perform a spatially weighted Gaussian smoothing of $\mathbf{U}^n$:

$$\mathbf{U}^n \leftarrow G_\sigma * (W \cdot \mathbf{U}^n)/G_\sigma * W, \qquad (7)$$

where W is a spatially-varying weighting function, with high value (1.0) at matched feature point locations and low (0.01) otherwise. We use a relatively large Gaussian smoothing kernel size ($\sigma = 2$ voxels) to help extend the effects of feature-constraint to a larger neighborhood.

To further improve computational speed, a multi-resolution scheme is also applied at this step. For the EMPIRE10 study, we used three resolution levels. The number of histogram bins for the MI computation was set to 32, 64, and 100 respectively for the three levels from coarse to fine.

2.4 Final Dense Deformable Registration

In this last step, we aim to further refine the deformation field estimation from the previous step on the original image resolution. No feature-based constraint is applied at this step since the feature points were detected on the lower resolution images. In order to align structure boundaries better while still being robust to image intensity changes, we have designed a hybrid image matching metric to be used at this step, which is a combination of the MI metric and a new normalized-sum-of-squared-differences (NSSD) metric:

$$\mathcal{J}_H(I, J, \mathbf{U}) = -\mathrm{MI}(I, J, \mathbf{U}) + w \cdot \mathrm{NSSD}(\tilde{I}, \tilde{J}, \mathbf{U}), \tag{8}$$

where w denotes a relative weighting of the two terms, and $\mathbf{U}$ is a dense deformation field that is defined at every voxel of the fixed image. From our experience, this hybrid similarity measure provides better alignment of image edges than using the MI metric alone since the latter cannot account for local image contrast changes. The NSSD metric is an edge-based alignment metric, which is defined as follows:

$$\mathrm{NSSD}(\tilde{I}, \tilde{J}, \mathbf{U}) = \frac{1}{N} \sum_{\mathbf{x}} \|\tilde{I}(\mathbf{x}) - \tilde{J}(T(\mathbf{x}))\|^2, \tag{9}$$

where

$$\tilde{I}(\mathbf{x}) = \frac{I(\mathbf{x}) - \mu_I(\mathbf{x})}{\sigma_I(\mathbf{x})}, \tag{10}$$

and $\mu_I = G_\sigma * I$ and $\sigma_I^2 = G_\sigma * (I - \mu_I)^2$ denote the local intensity mean and local intensity variation for image I. Similar notations hold for image J. We call $\tilde{I}$ and $\tilde{J}$ the *normalized local offset* images as implied by Eq. (10).

A similar iterative and compositive update scheme as in Section 2.3 is applied to compute the optimal deformation field for Eq. (8). But unlike in the previous section, two separate update fields are computed at each iteration step. The first update field optimizes the MI-term and is computed in exactly the same way as explained in Section 2.3. A second update field is computed in a similar fashion to optimize the NSSD term, since this term can also be expressed as the summation of N independent point-wise terms as indicated in Eq. (9). Hence, the explicit search scheme is applied again to find an update field $\mathbf{u}^n$ that optimizes $\mathrm{NSSD}(\tilde{I}, \tilde{J}^n \circ (\mathbf{Id} + \mathbf{u}^n))$. The two update fields are simply averaged to get the final update field, which is then used to update the total deformation field following Eq. (4). To ensure smoothness of the computed deformation field, we regularize it with a uniform Gaussian filter at each iteration. A kernel size of 1-voxel is used in all the experiments.

3 Results and Discussions

To evaluate the performance of the proposed method, we applied it to the 20 pairs of lung CT images provided by the EMPIRE10 workshop. Each pair of scans was taken from a single subject but the whole data sets came from a variety of sources including different institutes with different scanners and protocols. Most of the data have a fine sub-millimeter image resolution (≈ 0.7mm isotropic), but a few have a lower resolution larger than 1 mm. More details of the data can be found in the summary paper by the workshop organizers [19]. In addition to the original CT image data, a pre-segmented binary mask of the lung is also provided for each image. However, our method does not require pre-segmention of the lung and hence these lung masks were ignored when generating the results presented below.

Evaluation of the image registration results was carried out by the workshop organizers independently. Performance measures were computed in four categories: alignment of the lung boundaries, alignment of the major fissures, correspondence of annotated point pairs, and singularities in the deformation field. Details of the evaluation criteria can also be found in [19].

Table 1 summaries the performance of our method and the comparison against other 33 competing algorithms. Overall, our method performed well in each of the four categories, and achieved a final overall ranking of 6 among all 34 algorithms. The use of feature-based constraints clearly helped reducing the landmark error for difficult cases such as Subjects 01, 14, and 18. We note that we have used a fixed set of parameters for all the image data. In particular, deformation field regularization was set relatively low in order to handle large deformations. This low regularization effectively increases the degrees of freedom of the deformation model and renders the method more sensitive to image noise and other artifacts, which is possibly the main reason for the relatively worse performance of our method for the easy cases (i.e., small lung deformation) such as Subject 06.

The feature descriptors applied in this work are robust to local deformations to some extent; large deformations can still prevent the detection of all correct matches. Due to computation time considerations, the feature matching is currently only performed once after the linear registration step. It may be beneficial to repeat the feature matching several times during the iterative estimation of the deformation field.

We are still working on implementing the proposed registration method on GPU in order to improve the computational speed. For the results presented above, only the last hybrid MI/NSSD deformable registration step was computed on GPU using a CUDA-based implementation, and the other three steps were computed on CPU. The computation time was approximately proportional to the input image size. For a pair of images of fine resolution such as Subject 10, the linear registration step took about 3 seconds, the feature detection and matching about 2.5 minutes, the feature-constrained MI deformable registration about 2 minutes, and the last hybrid MI/NSSD deformable registration about 2.3 minutes (for 8 sub-regions in total). The total computation time for this case was about 7 minutes. The time was recorded on a HP xw8400 desktop computer equipped with an Intel Xeon Quad-core 2.66 GHz CPU and a NVIDIA GTX 280 graphics card.

	Lung Boundaries		**Fissures**		**Landmarks**		**Singularities**	
Scan Pair	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**
01	0.01	16.00	0.03	9.00	1.17	1.00	0.00	11.50
02	0.00	11.00	0.00	15.00	0.49	17.00	0.00	12.50
03	0.00	5.50	0.00	12.50	0.49	17.00	0.00	12.00
04	0.00	20.00	0.00	16.50	1.12	14.00	0.00	14.00
05	0.00	13.00	0.00	16.00	0.03	18.00	0.00	13.50
06	0.00	16.00	0.00	15.00	0.49	27.00	0.00	14.00
07	0.00	10.00	1.99	21.00	1.73	6.00	0.00	10.00
08	0.00	8.00	0.22	19.00	0.77	9.00	0.00	12.50
09	0.00	7.00	0.07	30.00	0.59	15.00	0.00	13.00
10	0.00	12.00	0.00	15.00	2.24	17.00	0.00	13.50
11	0.02	14.00	0.07	14.00	0.78	8.00	0.00	11.50
12	0.00	10.00	0.00	13.50	0.05	11.00	0.00	14.50
13	0.00	9.00	0.07	9.00	0.95	15.00	0.00	13.00
14	0.01	12.00	3.21	14.00	1.12	2.00	0.00	9.50
15	0.00	8.00	0.00	7.00	0.70	17.00	0.00	12.50
16	0.00	20.00	0.00	2.50	1.02	11.00	0.00	13.50
17	0.00	6.50	0.05	16.00	0.83	13.00	0.00	14.00
18	0.01	11.00	0.80	8.00	1.49	4.00	0.00	10.50
19	0.00	14.00	0.00	12.00	0.62	22.00	0.00	14.50
20	0.00	9.00	2.86	17.00	1.46	8.00	0.00	10.50
Avg	0.00	11.60	0.47	14.10	0.91	12.60	0.00	12.52
Average Ranking Overall								12.70
Final Placement								6

Table 1. Results for each scan pair, per category and overall. Rankings and final placement are from a total of 34 competing algorithms.

4 Conclusion

We have developed a fully automated, hybrid nonlinear image registration methods for intra-subject registration of lung CT images. Promising results were obtained when tested on the EMPIRE10 data. Future work includes adaptive parameter selection and investigating other feature detection methods. It is also desirable to incorporate proper physical deformation models to further improve the registration accuracy.

References

1. Parkin, D., Bray, F., Ferlay, J., Pisani, P.: Global cancer statistics, 2002. CA Cancer J. Clin. **55**(2) (2005) 74–108
2. Sarrut, D.: Deformable registration for image-guided radiation therapy. Z Med Phys **16** (2006) 285–297
3. Ding, K., Bayouth, J.E., Buatti, J.M., Christensen, G.E., Reinhardt, J.M.: 4DCT-based measurement of changes in pulmonary function following a course of radiation therapy. Medical Physics **37**(3) (2010) 1261–1272
4. Coselmon, M., Balter, J., McShan, D., Kessler, M.: Mutual information based CT registration of the lung at exhale and inhale breathing states using thin-plate splines. Med. Phys. **31** (2004) 2942–8
5. Urschler, M., Bauer, J., Ditt, H., Bischof, H.: Automatic point landmark matching for regularizing nonlinear intensity registration: Application to thoracic CT images. In: MICCAI 2006. LNCS 4191, Spinger-Verlag, Heidelberg (2006) 710–717
6. Li, P., Malsh, U., Bendl, R.: Combination of intensity-based image registration with 3D simulation in radiation therapy. Phys. Med. Biol. **53** (2008) 4621–4637
7. Li, B., Christensen, G.E., McLennan, G., Hoffman, E.A., Reinhardt, J.M.: Pulmonary CT image registration and warping for tracking tissue deformation during the respiratory cycle through 3-D consistent image registration. Med. Phys. **35**(12) (2008) 5575
8. Huang, Y., Feng, H., Zhao, P., Tong, T., Li, C.: Automatic landmark detection and nonrigid registration of intra-subject lung CT images. In: Intl. Conf. Inf. Sci. Eng., Los Alamitos, CA, USA, IEEE Computer Society (2009) 3605–3608
9. Yin, Y., Lin, C., Hoffman, E.: Mass preserving nonrigid registration of CT lung images using cubic B-spline. Med. Phys. **36** (2009) 4213–22
10. Hilsmann, A., Vik, T., Kaus, M., Franks, K., Bissonette, J.P., Purdie, T., Beziak, A., Aach, T.: Deformable 4DCT lung registration with vessel bifurcations. In: CARS 2007, Berlin, Germany (2007)
11. Viola, P., Wells, W.: Alignment by maximization of mutual information. Int. J. Comput. Vision **24** (1997) 137–154
12. Hartkens, T., Rohr, K., Stiehl, H.: Evaluation of 3D operators for the detection of anatomical point landmarks in MR and CT images. Comput. Vision Imag. Under. **86** (2002) 118–136
13. Kwon, D., Yun, I., Lee, K., Lee, S.: Efficient feature-based nonrigid registration of multiphase liver CT volumes. In: BMVC08. (2008)
14. Lowe, D.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60** (2004) 91–110
15. Allaire, S., Kim, J., Breen, S., Haffray, D., Pekar, V.: Full orientation invariance and improved feature selectivity of 3D SIFT with application to medical image analysis. In: MMBIA 2008, IEEE (2008)
16. Cheung, W., Hamarneh, G.: *n*-SIFT: *n*-dimensional scale invariant feature transform. IEEE Trans. Imag. Proc. **18** (2009) 2012–2021
17. Bay, H., Ess, A., Tuytelaars, T., Gool, L.: SURF: Speeded up robust features. Comput. Vision Imag. Under. **110** (2008) 346–359
18. Han, X., Hibbard, L., Willcut, V.: GPU-accelerated, gradient-free MI deformable registration for atlas-based MR brain image segmentation. In: MMBIA 2009: IEEE Workshop on Mathematical Methods in Biomedical Image Analysis. (2009)
19. Murphy, K., van Ginneken, B., Reinhardt, J., Kabus, S., Ding, K., Deng, X., Pluim, J.P.W.: Evaluation of methods for pulmonary image registration: The EMPIRE10 study. In: Grand Challenges in Medical Image Analysis, 2010. (2010)