MRF-based deformable registration and ventilation estimation of lung CT

Mattias P. Heinrich, Mark Jenkinson, *Member, IEEE*, Sir Michael Brady, Julia A. Schnabel, *Member, IEEE*

***Abstract*—Deformable image registration is an important tool in medical image analysis. In the case of lung CT registration there are three major challenges: large motion of small features, sliding motions between organs, and changing image contrast due to compression. Recently, MRF-based discrete optimisation strategies have been proposed to overcome problems involved with continuous optimisation for registration, in particular its susceptibility to local minima. However, to date the simplifications made to obtain tractable computational complexity reduced the registration accuracy. We address these challenges and preserve the potentially higher quality of discrete approaches with three novel contributions. First, we use an image-derived minimum spanning tree as a simplified graph structure, which copes well with the complex sliding motion and allows us to find the global optimum very efficiently. Second, a stochastic sampling approach for the similarity cost between images is introduced within a symmetric, diffeomorphic B-spline transformation model with diffusion regularisation. The complexity is reduced by orders of magnitude and enables the minimisation of much larger label spaces. In addition to the geometric transform labels, hyper-labels are introduced, which represent local intensity variations in this task, and allow for the direct estimation of lung ventilation. We validate the improvements in accuracy and performance on exhale-inhale CT volume pairs using a large number of expert landmarks.**

***Index Terms*—Non-rigid registration, stochastic optimisation, Markov random fields, minimum-spanning-tree, lung ventilation, sliding motion, discrete optimisation**

I. Introduction

DEFORMABLE registration of CT lung volumes has several clinical applications. It may be used for atlas-based segmentation of the lungs and its lobes [1], where expert segmentations from a database are propagated to a new subject. Longitudinal CT scans from the same patient can be used to monitor treatment or disease progression, e.g. for lung nodules [2]. 4DCT scans are now widely used for motion estimation in radiotherapy planning to increase the accuracy of dose delivery [3]. Deformable registration of dynamic CT scans can also enable direct estimation of lung ventilation [4] to assess patients with breathing disorders or spare well-functioning lung tissue from radiotherapy.

Non-rigid registration algorithms aim to solve an ill-posed, non-convex optimisation problem with several million degrees of freedom. There are three particular challenges in lung registration: large motions of small features, sliding motions between organs, and changing image contrast due to compression.

A. Large motion of small features

Motion within the lungs can often be larger than the scale of the features (vessels and airways). This may cause a registration algorithm to get trapped in a local minimum, and may lead to an erroneous registration. Local minima are frequently encountered in lung registration. In a recent comparison study on pulmonary CT registration [5], most algorithms have used continuous optimisation, which is particularly susceptible to local minima. In [6] a hybrid approach consisting of both a local intensity-based and sparse descriptor matching has been introduced within a variational framework. An alternative approach to avoid local minima is the use of discrete optimisation, which is usually formulated on a Markov random field (MRF). Discrete optimisation offers numerous advantages, in particular a greater control over the displacement space, to overcome these limitations. However, the space $\mathcal{L}$ of possible displacements has to be quantised, leading to many more degrees of freedoms. One state-of-the-art MRF-based medical image registration method, **drop** [7], reduces the dimensionality of the problem by sampling the displacement space only sparsely (along the three axes). This however may cause the optimal displacement to be missed. The authors attempt to address this problem by iteratively updating the transformation (thus warping the source image towards the target). Experimental results in [5] do not show an improvement of **drop** compared to continuous optimisation based methods, mainly because it lacks some of the attractive properties that a discrete framework potentially offers: most notably, avoidance of an iterative solution and the use of a dense sampling of the displacement space.

Instead, we use a dense stochastic sampling approach. The whole range of possible displacements is considered at each node in a dense manner (with a discretisation step of ≤ 2 voxels). We use a parametric B-spline transformation model with a regular control point grid. The voxels within non-overlapping cubes centred on the control points are assumed to move with the same translation vector. To reduce the complexity of the similarity cost calculations, it is calculated only for a random sample of all voxels within one cube. A similar approach has been introduced for continuous optimisation in [8]. The dense sampling approach does not require an iterative solution and enables accurate registration of small anatomical features undergoing large motion.

M.P. Heinrich and J.A. Schnabel are with the Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, OX3 7DQ, UK. (email: mattias.heinrich@eng.ox.ac.uk). M. Jenkinson is with the Oxford Centre for Functional MRI of the Brain, University of Oxford, OX3 9DU, UK. Sir M. Brady is with the Department of Oncology, University of Oxford, OX3 7DQ, UK.

Copyright (c) 2010 IEEE. Personal use is permitted. For any other purposes, permission must be obtained from the IEEE by emailing pubs-permissions@ieee.org.

B. Sliding motion at lung surfaces

Most registration algorithms include prior knowledge about the smoothness of deformations into the optimisation process to avoid physically implausible folding or gaps in the deformation field. The smoothness constraint typically assumes homogeneous motion and can be part of the transformation model (e.g. B-splines) or used as a penalty (regularisation) term. This assumption is violated in the case of a sliding motion between two objects (at their boundary), which commonly occurs during respiration. A homogeneous smoothness prior at sliding surfaces causes the registration to be inaccurate. Several authors address this problem by masking out the background objects that follow a different motion (e.g. a "motion mask" is used in [9]). Two separate registrations are then performed for the foreground and background objects. However, in order to analyse the motion close to the boundary this requires a non-trivial fusion of the two resulting motion fields. In [10], a direction-dependent regularisation is proposed that is based on an automatically detected mask.

In our previous work [11], we introduced a different graph structure, namely an intensity-derived minimum-spanning-tree, which effectively models sliding preserving motion. Here, we extend on this work, by using a more principled approach for the transformation model (including subpixel accuracy, a multi-level B-spline parameterisation, and incremental regularisation) and by introducing the concept of hyper-labels for simultaneous ventilation estimation.

C. Intensity variation due to lung compression

A local change in lung volume is expressed as a relative difference in the corresponding Hounsfield values within the breathing cycle. The change in density (and image intensity) can be problematic for deformable registration if a one-to-one intensity mapping is assumed. Similarity metrics that assume a globally linear relationship (e.g. cross-correlation) or statistical dependency (e.g. mutual information), cannot resolve for the locally varying contrast. Locally contrast invariant metrics, such as local cross-correlation [12], local phase [13] or modality independent neighbourhood descriptors [14] may be used instead. Recently, so called mass-preserving similarity terms have been introduced by [15] and [16]. Based on our discrete MRF-based registration setting, we propose an alternative solution, where we introduce an additional dimension into the label space $\mathcal{L}$. The fourth label dimension represents a local multiplicative intensity variation for each control point in the graph. The advantage of this approach is that we can simultaneously estimate a dense motion field and a regularised map of local density change. This directly provides a ventilation image of the lung functionality. Previous two-step approaches [4], [17], [18] first estimated a dense motion field and then computed a local intensity mapping based either on a Jacobian or on analysis of the Hounsfield units (HU) (based on the difference image after registration). Ventilation estimation based on a combination of Jacobian and HU, named the corrected Jacobian, has been presented in [19].

The remainder of this paper is structured as follows. In Sec. II the discrete optimisation framework is formulated as a labelling problem on an MRF. An intensity-driven minimum-spanning-tree representation is presented, which deals effectively with the complex sliding motion and offers attractive computational advantages over commonly used optimisers (such as graph cuts), because it can solve the labelling problem using only two passes of the min-sum message passing algorithm. A lower-envelope computation is introduced for a consecutive regularisation penalty, which enforces the smoothness of the full transformation (in a multi-level approach) and not only for the updates. A symmetric, diffeomorphic transformation model is employed, which both ensures inverse consistency and a singularity-free deformation field. Thereafter, the dense stochastic sampling is discussed and the concept of hyper-labels is introduced. An overview of the resulting algorithm is presented in Fig. 1. Section III describes the 4DCT lung dataset consisting of 10 cancer patients and a large set of manually defined landmarks. The landmarks are used in Sec. IV to evaluate the novel contributions made in Sec. II in terms of registration accuracy, complexity of deformations and computational efficiency. The results are also compared to the state-of-the-art continuous optimisation framework **gsyn** [12], which performed best in a recent lung registration challenge [5], and two methods from the literature, which have been evaluated using the same dataset. The capabilities of our approach to accurately align small features undergoing large deformations, the preservation of naturally occurring sliding motions, and simultaneous lung ventilation estimation are clearly shown.

II. Methods

Deformable registration using discrete optimisation is usually formulated as Markov random field (MRF) labelling. For the purposes of our parametric image registration framework, a graph is defined, in which the nodes $p \in \mathcal{P}$ (with spatial location $\mathbf{x}_p$) correspond to control points in a uniform B-spline grid and in which, for each node, there is a set of labels f_p, which correspond to discrete displacements. The energy function to be optimised consists of two terms: the data (also called unary) cost D (which is independent for each node); and the pair-wise regularisation cost $R(f_p, f_q)$ for any node q, which is directly connected ($\in \mathcal{N}$) with p:

$$E(f) = \underbrace{\sum_{p \in \mathcal{P}} D(f_p)}_{\text{data term}} + \alpha \underbrace{\sum_{(p,q) \in \mathcal{N}} R(f_p, f_q)}_{\text{regularisation term}} \tag{1}$$

The unary cost measures the similarity of the voxels around a control point p in one image and the set of voxels in the second image around the control point location, which is displaced by f_p. It is independent of the displacements of its neighbours. The pair-wise term enforces a globally smooth transformation by penalising deviations of the displacements of neighbouring nodes. The weighting parameter α sets the influence of the regularisation. For the case of lung CT registration, sum of absolute differences (SAD) has been widely used as similarity metric and the deformation field is regularised using

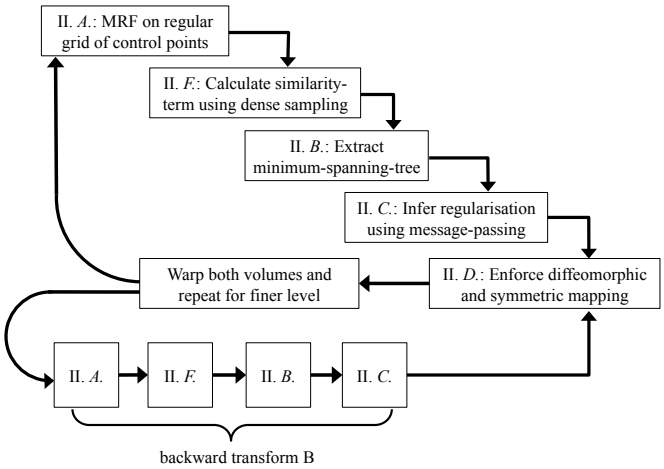

Fig. 1: Overview of presented approach for MRF-based lung registration. The flow-chart displays the algorithmic order of individual steps presented in Section II. A C++ implementation of the algorithm is available at http://users.ox.ac.uk/~shil3388.

the squared differences of the displacements of neighbouring control points. Let each label f_p describe a three dimensional displacement $f_p = \mathbf{u}_p = \{u_p, v_p, w_p\}$ between a control point p in the fixed image I and the moving image J. The energy term then becomes:

$$E(f) = \sum_{p \in \mathcal{P}} |I(\mathbf{x}_p) - J(\mathbf{x}_p + \mathbf{u}_p)| + \alpha \sum_{(p,q) \in \mathcal{N}} \frac{||\mathbf{u}_p - \mathbf{u}_q||^2}{||\mathbf{x}_p - \mathbf{x}_q||} \tag{2}$$

A. MRF-based optimisation

Methods to solve the MRF labelling problem can generally be categorised as one of two approaches: message passing and graph cuts. Message passing schemes include: loopy belief propagation (LBP) [20]; sequential tree-reweighted message passing (TRW-S) [21]; and dynamic programming on a tree [22]. Popular graph cut algorithms include: α-expansion moves (α-GC) [23]; and the fast primal-dual strategy (FastPD) [24]. Graph cuts can solve binary energy minimisation problems exactly by finding the minimum cut, which separates a graph, in which each node is connected to its neighbours and two additional nodes (source and sink). α-expansion moves are an extension of graph cuts to solve multi-labelling problems. Even though they are guaranteed to converge, they do not find the global optimum in most applications. Since α-GC relies on the pair-wise potential to be a metric, the most commonly used regularisation term, squared differences of displacements, cannot be used (only the square root of this, the L2 norm, is a metric). FastPD shows an improved performance compared to α-GC and relaxes the metric-requirement. However, this comes at the cost of substantially increased memory requirements (FastPD requires roughly 1000 bytes memory per degree of freedom, compared to 6 bytes for α-GC or message passing approaches).

Following the paradigm of pictorial structures [22], we propose that medical images can be more efficiently treated using a relaxed graph structure. Instead of connecting each node to its six immediate neighbours, only the most relevant edges are considered, leading to a minimum spanning tree (MST). The MST is a spanning tree with minimum total edge costs. The selection of the root node does not influence the result of the message passing. The MST sufficiently reflects the underlying anatomical connectivity in a medical image (see Fig. 2). Message passing on a tree finds the global minimum, without iterations, in only two passes.

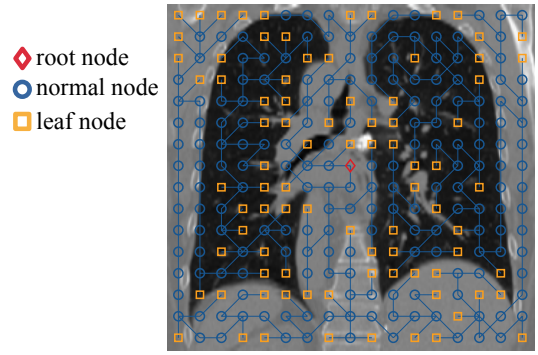

Fig. 2: Example of minimum spanning tree (MST) of a 2D coronal slice of a lung CT based on 8-connectedness (◇ root, ◯ normal, and □ leaf nodes).

B. Minimum spanning tree

Using Prim's algorithm [25], we can quickly find the unique MST given a set of nodes $p \in \mathcal{P}$ and edges e. The edge weight $w(p, q)$ is defined as the sum of absolute differences (SAD) between the intensities of all voxels within the influence region of a control point p and the respective voxels for a neighbouring control point q. The tree is well balanced, and, as a consequence, the maximum width is approximately $|\mathcal{P}|/\log|\mathcal{P}|$. The output of Prim's algorithm consists of a sorted list of all nodes (with increasing tree depth) and the index of each node's parent. A similar approach has been used for stereo correspondence [26], however other methods (LBP, TRW-S) perform better in that specific application.

Finding the best labelling for each node, i.e. the global optimum of Eq. 1, is possible using the min-sum message passing algorithm on the MST [22]. At each node p, the cost C_p of the best displacement can be found, given the displacement f_q of its parent q:

$$C_p(f_q) = \min_{f_p} \left(D(f_p) + \alpha R(f_p, f_q) + \sum_c C_c(f_p) \right) \tag{3}$$

where c are the children of p. The best displacement can be found by replacing $\min$ with argmin in Eq. 3. For any leaf node, Eq. 3 can be evaluated directly (since it has no children). Thereafter, the tree is traversed from its leaves down to the root node. It is worth noting that the costs C_p only have to be stored for the next tree level (only the argmin is needed to select the best displacement). Once the root node is reached, the best labelling for all nodes can be selected in another pass through the tree (from root to leaves). Another advantage of using message passing on a MST is that the exact marginal

distributions for each node can be directly obtained, which can be used to quantify the local uncertainty of the registration.

C. Message passing with incremental diffusion regularisation

Finding the minimum naïvely requires $|\mathcal{L}|^2$ calculations for the regularisation cost per pair of nodes. In [20] the min-convolution technique is introduced, which reduces the complexity to $|\mathcal{L}|$ by employing a lower envelope computation. For most commonly used (pair-wise) regularisation terms, such as diffusion (squared difference of displacements) and total variation (absolute difference) regularisation, this simplification is possible. Each label f_p can be represented by an upward facing parabola rooted at $(f_p, D^*(f_p))$, where $D^*(f_p) = D(f_p) + \sum_c C_c(f_p)$. The minimisation in Eq. 3 is defined by the lower envelope of these parabolas. In order to find this lower envelope in a single pass over all labels, the intersection between the parabola of the current label and the right-most parabola of the lower envelope needs to be calculated. This technique, which is described in detail in [20] requires the label spaces $\mathcal{L}$ of both nodes to be equivalent. We make an extension to this method, which enables the use of an incremental regularisation in a multi-level scheme. If a previous deformation field is known, we first warp the moving image towards the fixed image. Since, the regularisation cost depends only on the difference between two labelings, only the (subpixel) offset $\Delta = u_p - u_q, (p, q) \in \mathcal{N}$ for each dimension between displacements has to be considered. The lower envelope can be found in a similar way as in [20], however the coordinates of the intersections s of the parabolas now depend on the offset Δ:

$$s = \frac{(D^*(f_q) + (f_q + \Delta)^2)) - (D^*(f_p) + (f_p + \Delta)^2))}{2 \cdot (f_q - f_p)} \quad (4)$$

D. Symmetry and diffeomorphism

For many deformable registration algorithms, one image has to be chosen as the target, the other as the moving image. This biases the registration outcome and may additionally introduce an inverse consistency error (ICE). The ICE for a forward transform A and a backward transform B is defined as the difference between AB and the identity. In [12] a symmetric deformable registration is introduced, which calculates a transformation from both images to a common intermediate image and also ensures that $A(0.5) = B(0.5)^{-1}$, where 0.5 defines a transformation of half length. The full forward transformation is then $A(0.5) \circ B(0.5)^{-1}$.

We adopt a similar approach, in which we first estimate both full transformations A and B independently. In order to calculate a valid inverse transformation the transformations have to be diffeomorphic. In [27] it is shown that B-spline transformations can also be guaranteed to be diffeomorphic if the maximum displacement of each control point is limited to ≈ 0.4 times the grid spacing. However, this reduces the maximal possible deformations otherwise the grid spacing has to be increased by too much. In our approach, we first allow for displacements larger than this theoretical limit, and afterwards obtain a diffeomorphic mapping by applying the scaling and squaring method [28]. This approach avoids transformations for which physically implausible folding of volume occurs and yields a continuous-valued transformation. We then use a fast iterative inversion method, as presented in [29], to invert the half-length transformations and obtain $A(0.5)^{-1}$ and $B(0.5)^{-1}$. This ensures that the two new symmetric transformations $A^S = A(0.5) \circ B(0.5)^{-1}$ and $B^S = B(0.5) \circ A(0.5)^{-1}$ are inverse consistent. A different approach was presented by [30], which uses only one transformation grid (in the space of the intermediate mean) and ensures a symmetric transformation by enforcing the forward displacements to be the opposite (negative) of backward displacements. This removes the need to calculate two separate transformations, but has a lower accuracy, since the lowest step (without additional interpolation) is then two voxels.

E. B-spline transformation models

To avoid local minima, most continuous-optimisation-based registration algorithms use a multi-resolution scheme in which the images are downsampled after Gaussian smoothing. This may degrade the quality of the registration. We adopt a multi-level B-spline scheme [31], in which we only employ the highest resolution image. For a given level, the image is subdivided according to the B-spline grid into non-overlapping cubic groups of voxels. We choose non-overlapping cubes to preserve the independence of the unary term for all nodes. The similarity cost, which is incurred when translating a cube of voxels (or equivalently move a control point in the B-spline grid) is aggregated voxel-wise as explained in the next section. Subsequently, the regularisation term is calculated only for each control point in the B-spline grid. Using this approach, both high spatial accuracy and low computational complexity are achieved. For the next level, the previous deformation field is obtained using the B-spline interpolation for voxels between control points and used as the prior deformation. The order of the B-spline (linear, quadratic or cubic) can be chosen according to the application. While this parametric transformation model introduces a way of obtaining subpixel accuracy, a hybrid approach using an additional continuous refinement step (c.f. [32]) may be beneficial.

F. Dense stochastic sampling

In the previously presented **drop** method [7] the complexity (especially the memory requirements of the employed FastPD optimisation) of the registration had to be reduced by using a sparse sampling of the deformation space. Instead of densely sampling the deformations in all three dimensions, only displacements along the three axes are considered. This may lead to a non-optimal registration and severely reduce its accuracy. In [33] it has been shown that this leads to similar problems, which are common for gradient-based optimisation (local minima, bias of initialisation). Using the more efficient optimisation strategy presented above, the search space can be increased substantially. We therefore propose to use a dense displacement sampling with a discretisation step of only 2 voxels for the first level. While the complexity of the smoothness calculations is reduced using the parametric grid

model, the computational complexity of the similarity cost is only dependent on the number of voxels (and independent of the number of grid nodes). This can easily outweigh the performance gain, which is achieved with the more efficient MST-based regularisation. In the context of gradient-based medical image registration the concept of stochastic optimisation has been introduced by [8]. In the case of a parametric model there are many more voxels than control points in the transformation grid. The similarity cost of a certain control point displacement is then aggregated (or summed) for all voxels within the influence range of the control point. Similar to [34], we can make a stochastic approximation to this summation. Only a random subset K of all voxels is used for each control point. This makes the similarity term computation many times more efficient with very little sacrifice of registration accuracy. K is chosen experimentally in Sec. III-A.

G. Hyper-labels for dynamic imaging

In many applications of medical image registration, finding the transformation parameters is only an intermediate target. One may be interested in propagating segmentation labels, estimating contrast uptake in dynamic sequences or generating a patient-specific (predictive) motion model. Parameters other than geometric displacements are of interest. We introduce hyper-labels into our registration framework, for which a fourth (and possible higher) dimension is added to the label space. This additional parameter may correspond to a non-uniform multiplicative intensity variation (to estimate bias fields or pharmacokinetic parameters), a segmentation label (c.f. [35]) or a motion parameter (e.g. the time phase-shift for a sinusoidal motion in 4D data). In this paper, hyper-labels are used to directly estimate the density change of lung tissue and therefore the local lung ventilation.

Simultaneous image registration and intensity correction has been proposed to deal with non-uniform bias fields in combination with brain tissue segmentation in [36]. The coding complexity of the residual image is minimised in [37] to estimate a smooth intensity correction field. Similarly, in [38] a regularised correction function is employed in a variational optimisation framework to compensate for inhomogeneous intensity mappings. The disadvantage of such approaches is that they all rely on continuous optimisation schemes and are readily trapped in local minima. In contrast, our proposed hyper-label approach allows us to freely define the range of possible intensity correction values. It can also make use of the globally optimal regularisation, avoid local minima and is independent of the initialisation.

Two different metrics have primarily been used so far to derive ventilation images from dynamic CT [17]: Hounsfield unit (HU) change and Jacobian determinant of the deformations. Both methods were studied for ventilation estimation in [39], [4], and [40]. In [41] it was found that only the HU-based ventilation estimate had a significant correlation with lung functionality of emphysema patients. In most cases the ventilation estimation was performed in a post-processing step after deformable registration. In [15] and [16] the intensity change based on the Jacobian determinant is included in the similarity metric during the registration (so called mass-preserving registration). For MRF-based optimisation this would introduce an unwanted dependency of the unary and pair-wise potentials, so we adapt the HU-based approach. According to [41] the fractional local change of lung volume ΔV_{HU} is defined as:

$$\Delta V_{HU}(\mathbf{x}_p) = \frac{I(\mathbf{x}_p) - J(\mathbf{x}_p + \mathbf{u}_p)}{J(\mathbf{x}_p + \mathbf{u}_p) + 1000} \cdot V_I(\mathbf{x}_p)^{vox}\,, \quad (5)$$

where I represents the exhale and J the inhale volume, both measured in Hounsfield units, and $V_I(\mathbf{x}_p)^{vox} := 1$ the exhale voxel volume, which is constant across the image. To include this locally varying density correction function into our registration framework we introduce a fourth variable ν to the label space $\mathcal{L}$, so that $f_p = \{u_p, v_p, w_p, \nu_p\}$. For each node, ν_p takes quantised values of local volume change and the smoothness of ν is ensured in the regularisation term. We simplify Eq. 5 by adding 1000 HU to both images. This yields:

$$\begin{aligned} \nu_p = \Delta V_{HU}(\mathbf{x}_p) &= \frac{I(\mathbf{x}_p)}{J(\mathbf{x}_p + \mathbf{u}_p)} - 1 \\ \Rightarrow (1 + \nu_p) \cdot J(\mathbf{x}_p + \mathbf{u}_p) &= I(\mathbf{x}_p) \end{aligned} \quad (6)$$

The unified energy term to be minimised is then:

$$\begin{aligned} E(f) = \sum_{p \in \mathcal{P}} |I(\mathbf{x}_p) - (1 + \nu_p) \cdot J(\mathbf{x}_p + \mathbf{u}_p)| + \\ \alpha \sum_{(p,q) \in \mathcal{N}} \frac{||\mathbf{u}_p - \mathbf{u}_q||^2 + \beta(\nu_p - \nu_q)^2}{||\mathbf{x}_p - \mathbf{x}_q||} \end{aligned} \quad (7)$$

where β can be used to weight the influence of density change on the regularisation. The min-sum algorithm can be applied as before, as well as the lower-envelope computation, except this is now performed over a four-dimensional array of smoothness terms. The ventilation image V_{HU} is directly given by ν_p extracted from the given labelling f_p^* after optimising Eq. 7. The choice of α will be discussed in Sec. IV-A. Figure 3 shows an example of the estimated ventilation maps, using the Jacobian metric and our approach.

III. Experiments

We performed deformable registration on ten cases of the DIR-Lab 4DCT dataset acquired at inhale and exhale phase [42][1]. The patients were treated for esophageal or lung cancer, and a breathing cycle CT scan of thorax and upper abdomen was obtained, with slice thickness of 2.5 mm, and an in-plane resolution ranging from 0.97 to 1.16 mm. Since these 4DCT scans are already in rigid alignment, no linear registration step was required. Particular challenges for the deformable registration task are the changing contrast between tissue and air, because the gas density changes due to compression, discontinuous sliding motion between lung lobes and the lung rib cage interface, and large deformations of small features (lung vessels, airways). For each scan 300 anatomical landmarks have been carefully annotated by thoracic imaging experts. The maximum average landmark error before registration is 15 mm over all cases. The first five cases have been cropped

[1]This dataset is freely available at http://www.dir-lab.com

(to include the thorax) and resampled to form an in-plane dimension of 256x256 by [42]. We apply a similar cropping to the remaining five datasets, but apply no resampling. The number of voxels n in each scan ranges from $\approx 4-7\times 10^6$.

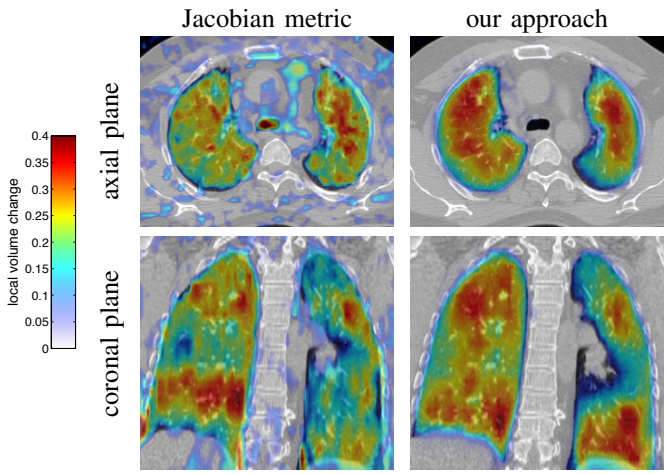

Fig. 3: Comparison of lung ventilation for Case 6. The local lung volume change is displayed as colour overlay (transparent for low values) over greyscale intensities of the inhale scan, without using segmentation masks. Our approach using simultaneous ventilation estimation with hyper-labels demonstrates a much smoother map of lung functionality, than the two-step approach using the Jacobian metric ($J-1$ as used e.g. in [17]). It can be seen that in this case the ventilation in the region around the tumour in the left lower lobe is substantially reduced.

A. Parameter choice

The challenges of the dataset require us to optimise over a large number of degrees of freedoms. Only the highest image resolution is used, but the number of nodes is reduced by using a uniform grid of control points. Three levels with decreasing grid-spacing of $g=\{8,6,4\}$ voxels are employed. This means that the number of nodes is increased for each subsequent level. The number of labels is correspondingly decreased, and the maximum search radius is set to $l_{\max}=\{8,6,4\}$ steps. The search space is defined as $\mathcal{L}=d\cdot\{0,\pm 1,\ldots,\pm l_{\max}\}^3$ voxels, where d is a discretisation step, which is defined as $\{2.0,1.0,0.5\}$ voxels for the three levels. Sub-pixel displacements are achieved by upsampling the moving image using trilinear interpolation.

For the similarity term $K=64$ random samples are used for each control point p within the cubic region $R_p=\{-g/2+1,-g/2+2,\ldots,g/2\}^3$. The sampling locations are uniformly distributed within the local support region R_p of a node. The same locations are used for all displacements $\mathbf{u}_p$ for one node, but the sampling is updated for every new node, to avoid a bias. Given a grid point spacing of 8 voxels, the standard deterministic similarity cost computation would require 512 calculations per node, the stochastic approach therefore yields an 8-fold improvement in computation time. By using the same numbers for g and $l_{\max}$ for each level the complexity is kept constant for each level, the number of similarity computations is $\sim K\cdot n/g^3\cdot(2l_{\max}+1)^3\approx 8Kn$, thus linear with the number of voxels. By using the lower envelope technique discussed in Sec. II-C the complexity of the regularisation term is also linear with n.

The B-spline order used to interpolate a deformation field to a higher level can be chosen according to the specific application. We found that a first order (linear) B-spline function provided best results, because it can better preserve the sliding motion. There is a small trade-off between accurate sliding motion preservation and enforcing of diffeomorphic, invertible transformations. We have, however, found in our experiments (see e.g. Fig. 6) that by using a small grid-spacing in combination with the dense displacement sampling a very sharp gradient in the motion field could be recovered, while at the same time preserving a transformation without singularities.

The hyper-labels, which are added as fourth dimension to the label space $\mathcal{L}$ are set to be $\nu=d\cdot\{0,1,\ldots,4\}$ with discretisation steps of $d=\{0.15,0.075,0.0375\}$. This means that the complexity is increased by a factor of 5. The maximum local volume change is therefore 60 %, which is well within the range of values measured in previous studies [41]. The weighting factor β in Eq. 7 is empirically set to 1.

IV. Results

To quantify the registration accuracy and the influence of each of our contributions, we use the 300 manually selected landmarks per scan, which are provided with the dataset [42]. The landmarks are well distributed throughout the lungs and have been selected with an intra-observer error of ~ 1 mm.

A. Influence of regularisation weighting α

In addition to the implicit regularisation of the parametric B-spline grid, the regularisation terms $R(f_p,f_q)$ in Eqs. 2 and 7 are of great importance. A higher value of α results in a smoother deformation field. Figure 4 shows the resulting registration accuracy and the mean squared error (MSE) between the deformed moving and fixed image with varying strength of the regularisation α for our proposed method. Interestingly, the lowest MSE coincides with the lowest target registration (TRE), which helps in selecting a suitable value for α. In order to compare different registration techniques, they should generally have a similar complexity of the resulting transformations. The standard deviation of the determinant of the Jacobian J of the deformation fields is used to measure the complexity, where $\mathrm{std}(J)=0$ describes a perfectly smooth transformation.

B. Evaluation of contributions

The deformable registrations for all ten cases between maximum inhalation and maximum exhalation are first performed using a symmetric transformation model, the dense stochastic sampling approach, without the use of hyper-labels and using the optimisation based on an image-derived MST. This forms the baseline for the subsequent experiments. An

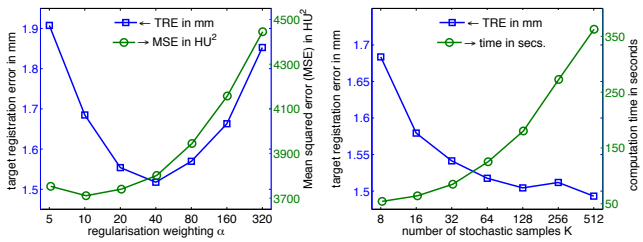

Fig. 4: Influence of α on registration accuracy (averaged for all cases) and mean squared error (MSE) is shown in the left plot. The right plot shows the accuracy with increasing number K of stochastic samples. The TRE does not improve substantially with more than 64 samples per control point.

TABLE I: Results for individual contributions assessed using average target registration error (TRE in mm) for 3000 expert selected landmarks. A significant ($p < 0.05$) improvement (lower median) compared to the baseline is indicated by **+**, a higher error by **-** (non-significant results are in brackets). The computation time (for optimisation and complete registration) using a two-threaded C++ implementation per scan pair is given in minutes. The optimal weighting for the regularisation α and the resulting complexity of the deformations, evaluated using the standard deviation of the Jacobians, are given.

	time (min) optim.	time (min) total	α	std(J)	TRE in mm	p-value
initial					8.46	
axial	0.12	2.37	40	0.103	2.25	$9{\cdot}10^{-19}$ **-**
non-adapt. tree	0.62	2.11	40	0.116	2.05	$7{\cdot}10^{-9}$ **-**
asymmetric	0.48	1.95	40	0.147	1.83	$9{\cdot}10^{-5}$ **-**
graph cut	17.8	19.3	5	0.118	1.73	0.152 (-)
non-stochastic	0.61	6.04	40	0.109	1.49	0.588 (+)
hyper-labels	4.57	7.97	40	0.099	1.43	0.017 **+**
baseline	0.62	2.09	40	0.109	1.52	

average registration accuracy of 1.52 mm, with a smoothness of std(J)=0.109 is obtained for the baseline of our proposed algorithm with a computation time of 2.04 minutes per case. Next, each individual contribution of this work is tested separately and compared to the baseline, thereby only one element is changed each time. The optimal regularisation weighting α is found for each method separately (as shown for the baseline in Fig. 4 left). All deformations are free from singularities. The results are summarised in Table I and Fig. 5.

1) Axial sampling: First, the dense displacement sampling is replaced by the sparse approach used in **drop** [7] (and similarly in [43] and [44]). Here only displacements along each of the three axes are considered. This substantially reduces the size of the label space from $|\mathcal{L}| = (2\,l_{\max}+1)^3$ to $|\mathcal{L}| = 3 \cdot 2l_{\max}+1$. Following [7] we introduce an iterative loop to partially compensate for the reduced degrees of freedom. The accuracy of the registration experiments significantly deteriorates (to 2.25 mm), confirming our previous assumption that an axial sampling of the displacements results in many false local minima of the registration problem.

2) Graph Cut: Second, the MST-based optimisation is replaced by α-expansion graph cuts [45], [46] using a six-

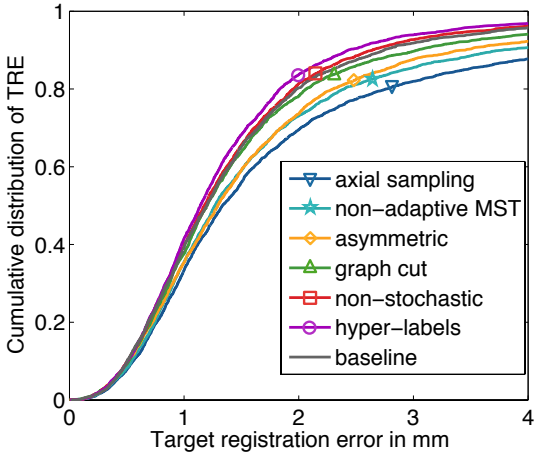

Fig. 5: Cumulative distribution of target registration error in mm of each method using 3000 manual expert landmarks. A significant improvement, using the Wilcoxon rank-sum test, can be seen for the baseline approach compared to asymmetric registration ($p = 9 \cdot 10^{-5}$), without using an image-adaptive MST ($p = 7{\cdot}10^{-9}$) and axial sampling ($p = 9{\cdot}10^{-19}$). Hyper-labels are significantly better than the presented baseline ($p = 0.017$).

connected graph. FastPD is not suitable because of its larger memory demand. Since our preferred diffusion regularisation term (squared differences) is not a metric and cannot not be optimised with α-GC, it is substituted with the L2 norm (Euclidean distance) (see also discussion in Sec. II-A). During the multi-level registration, the regularisation can only be applied to the update of the deformation field when using α-GC, and the computation time for the optimisation is increased by a factor of over 30, demonstrating the improvements achieved by the use of our MST-optimisation. The accuracy is also slightly decreased to a TRE of 1.73 mm.

3) Non-adaptive MST: As described in Sec. II-B, we employ an image-derived MST, which removes edges across locations with large intensity differences, which are likely to coincide with sliding motion. To evaluate the suitability of this approach, the alternative of choosing a non-adaptive MST, which is achieved by setting the edge-weights for the six neighbours of each node randomly, is included for comparison. Using a non-adaptive MST significantly reduces the accuracy to 2.05 mm. Figure 6 demonstrates that this is mainly due to the better preservation of sliding motion of our image-adaptive approach.

4) Asymmetric: In general, registration problems should be treated symmetrically to remove bias from the choice of fixed and moving image. The direct comparison of an asymmetric registration to the baseline also shows a significant improvement in terms of registration accuracy for the symmetric approach. The smoothness of the deformations of an asymmetric registration is also decreased, due to the fewer constraints.

5) Non-stochastic: As discussed in Sec. III-A the stochastic sampling approach for the similarity terms yields a greatly reduced computation time. Figure 4 shows that the registration

accuracy improves with an increasing number of samples, but converges at $K \approx 64$. The TRE of the deterministic approach is only slightly lower than the stochastic one and this change is not statistically significant.

6) Hyper-labels: Last, we study the influence of the proposed hyper-labels for simultaneous registration and intensity correction. Five quantised intensity correction labels are added to each geometric label, therefore this approach has a 5-fold increased complexity. The average TRE is further (statistical significantly) reduced to 1.43 mm. Additionally, the obtained intensity correction field can be directly used to quantify lung ventilation as described in Sec. II-G, an example for which is shown in Fig. 3. This has several advantages over calculating the ventilation as a post-processing step (c.f. [4],[40]), as it is not directly affected by misalignments, avoids local minima during registration and does not need a specific smoothing of the difference images (c.f. [41]). To further validate the ventilation estimation, an additional registration was performed for the POPI-model [47] (equivalent to case 17 of [5], which provides manual segmentations) and the total lung volume change was calculated by integrating all local ventilation values inside the lungs. A good agreement between the Jacobian method (246.6 ml), our hyper-label approach (246.1 ml), and manual segmentations (247.2 ml) was found.

C. Comparison to state-of-the-art

To demonstrate the impact of our MST-based registration framework, a comparison to previously proposed state-of-the-art algorithms is performed on the same dataset. Schmidt-Richberg et al. [10] proposed a variational lung registration method with direction-dependent regularisation to cope with the effects of discontinuous sliding motion. Gorbunova et al. [16] use a mass-preserving similarity term, which has been demonstrated to be less affected by the changes in contrast during respiration. Both methods were tested and evaluated by the authors on the same 4DCT dataset ([16] only uses a subset of 5 patients). We performed additional registration experiments using the **gsyn** method presented by Avants et al. [12], which is publicly available in the ANTS package[2]. **gsyn**, which is a symmetric, diffeomorphic, demons-like algorithm, has been chosen, because it performed best in a recent comparison study of pulmonary CT registration algorithms [5]. The following parameters were carefully chosen to obtain the best registration accuracy with similarly complex transformations as the proposed approach: 4 resolution levels; similarity metric: normalized cross correlation (NCC) (radius of 2 voxels); Gaussian smoothing of σ=3 and 1 voxels for gradient and deformation fields, respectively. In contrast to the results presented in [5] the full volumetric scans were considered and no lung masks were used to remove the outer lung information (and thereby invalidate the registration for locations outside of the mask). In a second experiment, we generate lung segmentations for all cases using thresholding and simple morphological operations and set the intensities of the outside voxels to the maximum of the ones inside the masks. We again optimised the smoothing parameters of **gsyn** for best registration accuracy, now yielding a more aggressive setting of σ=3.5 and 0.375 voxels respectively.

Table II shows the results for all compared approaches. Our proposed approach achieves the best results (TRE of 1.43 mm) with considerably lower computation times and without requiring lung/motion masks. Compared to **gsyn** applied to volumes with lung masking, the complexity of deformations is reduced by more than half. Figure 6 visualises the good registration quality and the well preserved sliding motion of our approach using an image-derived compared to a non-adaptive MST.

V. Discussion and Conclusion

A novel method based on discrete optimisation has been introduced that outperforms previous approaches for deformable CT lung registration, both in terms of accuracy and computational efficiency. A dense stochastic displacement sampling is performed for the similarity term on the highest available image resolution, which successfully avoids local minima, for which continuous optimisation methods are susceptible. The regularity of the deformation field is enforced using min-sum message passing on an image-derived minimum spanning tree (MST). This ensures a globally optimal solution without the need for an iterative scheme. The algorithm is efficiently implemented in a symmetric multi-level framework yielding lower computational complexity than other state-of-the-art algorithms, while employing many more degrees of freedom. The improvement in terms of computation times is ten-fold compared to the widely used continuous-based **gsyn** framework, as well as the discrete α-expansion graph cut optimisation. Since the data cost of all nodes and the pairwise potentials of nodes having the same depth in the MST can be computed independently, the approach is well suited for parallelisation on graphics processors.

An average TRE of 1.43 mm was found for a challenging dataset of inhale-exhale CT scans. This is a significant improvement over **gsyn** [12] (TRE=2.43 mm), when the full deformation field is required and no lung masks are employed. The accuracy of the proposed approach compares well with previously published results on the same dataset, which addressed the challenges of sliding motion [10] (TRE=2.13 mm, improved to TRE=1.55 mm in [48]) and locally varying contrast [16]. A globally smooth ventilation image is simultaneously estimated avoiding the previous two-step approaches to this task.

Despite only employing inhale and exhale scans, the accuracy of the proposed method is comparable to the best performing algorithms on this dataset [49] and [50], which use all time frames of the 4DCT cycle. Being able to estimate motion and ventilation with very high accuracy based on only inhale and exhale scans offers great clinical benefit, e.g. for diagnosis of breathing disorders, because the imaging radiation dose could be significantly reduced. However, for use in gated radiotherapy the motion of all phases throughout the breathing cycle has to be estimated.

In the presence of large deformations, our proposed dense stochastic displacement sampling yields a higher robustness

[2]http://www.picsl.upenn.edu/ANTS/

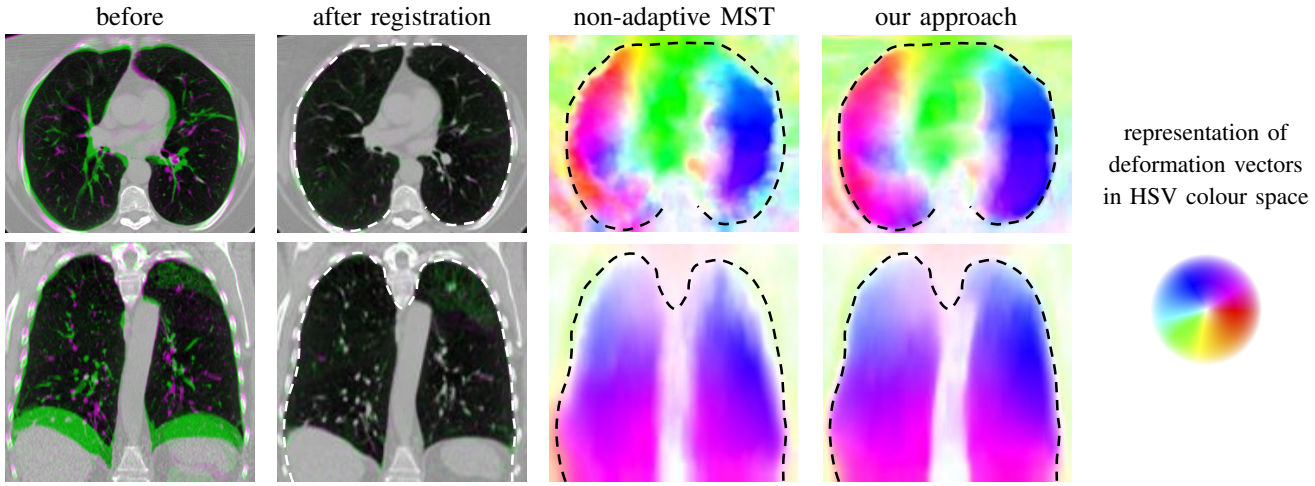

Fig. 6: Registration result for Case 8 of 4D CT dataset. The top row shows the axial plane, the bottom row the coronal plane. In columns from left to right is displayed: Overlay of inhale (green) and exhale (magenta) phase before and after registration, colour-coded deformation field of our method without using image information to derive MST and with the presented image-adaptive MST. The deformation fields demonstrate that the sliding of the lungs is better preserved using our MRF-based dense sampling approach with image-adaptive MST, which explains the improved registration accuracy. The outline of the thoracic cage is shown for visual guidance (dashed white or black line).

TABLE II: Results for deformable registration of inhale and exhale CT. Average target registration error (TRE) for 300 expert selected landmarks per scan pair, given in mm. Best performing algorithm per case is set in bold. Average computation time, (maximum) degrees of freedom (d.o.f.), and complexity of the deformation fields (std(J)) are given where available.

case #	initial	**gsyn** Avants[12] smooth deformations	symmetric (NCC) with lung masks	Gorbunova[16] mass-preserving	Schmidt-Richberg[10] sliding-preserving	proposed	proposed with hyper-labels
1	3.89±2.8	1.03±0.5	1.08±0.5	1.15±0.6	1.28±0.8	0.98±0.5	**0.97±0.5**
2	4.34±3.9	1.02±0.6	1.14±0.6	1.26±0.7	1.13±0.6	**0.95±0.5**	0.96±0.5
3	6.94±4.1	1.31±1.0	1.22±0.7	1.79±1.1	1.29±0.7	**1.18±0.7**	1.21±0.7
4	9.83±4.9	1.65±1.6	1.44±1.1	2.01±1.4	2.58±2.9	1.60±1.4	**1.39±1.0**
5	7.48±5.5	2.05±2.3	**1.45±1.5**	1.70±1.3	1.62±1.3	1.72±1.5	1.72±1.6
6	11.10±7.0	2.50±3.3	1.56±2.6	—	1.96±1.4	1.72±1.1	**1.49±1.0**
7	11.59±7.4	3.77±5.7	2.26±4.0	—	3.83±3.1	1.79±1.4	**1.58±1.2**
8	15.16±9.0	7.06±8.6	2.82±6.3	—	3.84±4.6	2.24±2.6	**2.11±2.4**
9	7.82±4.0	1.89±2.0	**1.17±1.3**	—	2.01±1.5	1.45±0.8	1.36±0.7
10	7.30±6.3	2.06±3.0	**1.50±2.3**	—	1.75±1.3	1.59±1.6	1.55±1.6
avg	8.46±6.58	2.43±4.1	1.57±2.1	—	2.13±1.8 1.55±1.1*[48]	1.52±1.4	**1.43±1.3**
std(J)		$12.9{\cdot}10^{-2}$	$20.9{\cdot}10^{-2}$	—	—	$10.9{\cdot}10^{-2}$	$\mathbf{9.9{\cdot}10^{-2}}$
avg. time		29 min	21 min	—	—	**2.09 min**	7.97 min
d.o.f.		$2.2{\times}10^{7}$	$2.2{\times}10^{7}$	—	—	$8.1{\times}10^{7}$	$40.4{\times}10^{7}$

*improved results from [48] using a fast explicit diffusion (FED) scheme with improved convergence

against misregistration and makes the registration outcome less dependent on the initialisation. We make use of an image-derived MST to represent the connectivity between control points and enable a very efficient optimisation. As demonstrated in the experiments, for a sufficiently small grid spacing the MST representation deals very well with the sliding motion at the pleural interface and accurately aligns the whole volume without the need of additional masking of the thoracic cage, because it reduces the regularisation across strong image gradients. While an MST representation might not be suitable for all applications of deformable image registration, the greatly reduced complexity may still make the approach better suited in terms of the trade-off between quality and efficiency. It has to be noted, that even a six-connected graph is only an approximation to the real connectivity in medical images. An interesting prospective development of this approach would be to incorporate anatomical knowledge (e.g. fissure segmentations [51]) into the construction of the MST.

We additionally introduced the concept of hyper-labels, which can carry complementary information in addition to the displacement labels and demonstrated its applicability to simultaneous ventilation estimation. Many more applications are possible using hyper-labels, such as estimation of MRI bias fields or contrast-uptake. In the future, we plan to employ this concept for other dynamic imaging modalities, such as dynamic contrast-enhanced MRI. A clinical validation of the ventilation estimation is intended in the future by correlating the obtained maps with hyper-polarised Xe-MR imaging. We also plan a submission of our new approach to the EMPIRE study [5] in order to enable a comparison to a wider range of state-of-the-art algorithms.

ACKNOWLEDGMENT

We thank EPSRC and Cancer Research UK for funding this work within the Oxford Cancer Imaging Centre.

REFERENCES

[1] L. Zhang, E. Hoffman, and J. Reinhardt, "Atlas-driven lung lobe segmentation in volumetric X-ray CT images," Medical Imaging, IEEE Transactions on, vol. 25, no. 1, pp. 1 –16, jan. 2006.

[2] Staring, M. and Pluim, J.P.W. and de Hoop, B.J. and Klein S. and van Ginneken, B. and Gietema, H. and Nossent, G. and Schaefer-Prokop, C.M. and van de Vorst, S. and Prokop, M., "Image Subtraction Facilitates Assessment of Volume and Density Change in Ground-Glass Opacities in Chest CT," Investigative Radiology, vol. 44, no. 2, pp. 61 – 66, February 2009.

[3] E. Weiss, K. Wijesooriya, S. V. Dill, and P. J. Keall, "Tumor and normal tissue motion in the thorax during respiration: Analysis of volumetric and positional variations using 4D CT," International Journal of Radiation Oncology*Biology*Physics, vol. 67, no. 1, pp. 296 – 307, 2007.

[4] R. Castillo, E. Castillo, J. Martinez, and T. Guerrero, "Ventilation from four dimensional computed tomography: Density versus jacobian methods," Physics in Medicine and Biology, vol. 55, pp. 4661–4685, 2010.

[5] K. Murphy, B. van Ginneken, J. Reinhardt, S. Kabus, K. Ding, and J. Pluim, "Evaluation of registration methods on thoracic CT: The EMPIRE10 challenge," Medical Imaging, IEEE Transactions on, vol. 30, no. 11, pp. 1901 –1920, 2011.

[6] T. Brox and J. Malik, "Large displacement optical flow: Descriptor matching in variational motion estimation," Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 33, no. 3, pp. 500 –513, march 2011.

[7] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios, "Dense image registration through MRFs and efficient linear programming," Medical Image Analysis, vol. 12, no. 6, pp. 731 – 741, 2008.

[8] S. Klein, M. Staring, and J. Pluim, "Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-Splines," Image Processing, IEEE Transactions on, vol. 16, no. 12, pp. 2879 –2890, dec. 2007.

[9] J. Vandemeulebroucke, O. Bernard, S. Rit, J. Kybic, P. Clarysse, and D. Sarrut, "Automated segmentation of a motion mask to preserve sliding motion in deformable registration of thoracic CT," Medical Physics, vol. 39, p. 1006, 2012.

[10] A. Schmidt-Richberg, R. Werner, H. Handels, and E. J., "Estimation of slipping organ motion by registration with direction-dependent regularization," Medical Image Analysis, vol. 16, no. 1, pp. 150–159, 2012.

[11] M. P. Heinrich, M. Jenkinson, S. M. Brady, and J. A. Schnabel, "Globally optimal registration on a minimum spanning tree using dense displacement sampling," in Medical Image Computing and Computer-Assisted Intervention MICCAI 2012, ser. Lecture Notes in Computer Science, H. Delingette, P. Golland, M. Kensaku, and N. Ayache, Eds. Springer Berlin / Heidelberg, 2012.

[12] B. Avants, C. Epstein, M. Grossman, and J. Gee, "Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain," Medical Image Analysis, vol. 12, no. 1, pp. 26 – 41, 2008.

[13] M. Mellor and M. Brady, "Phase mutual information as a similarity measure for registration," Medical Image Analysis, vol. 9, no. 4, pp. 330 – 343, 2005.

[14] M. P. Heinrich, M. Jenkinson, M. Bhushan, T. Matin, F. V. Gleeson, S. M. Brady, and J. A. Schnabel, "MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration," Medical Image Analysis, no. 0, pp. –, 2012.

[15] Y. Yin, E. Hoffman, and C. Lin, "Mass preserving nonrigid registration of CT lung images using cubic B-spline." Medical Physics, vol. 26, pp. 4213–22, 2009.

[16] V. Gorbunova, J. Sporring, P. Lo, M. Loeve, H. A. Tiddens, M. Nielsen, A. Dirksen, and M. de Bruijne, "Mass preserving image registration for lung ct," Medical Image Analysis, vol. 16, no. 4, pp. 786 – 795, 2012.

[17] T. Yamamoto, S. Kabus, T. Klinder, J. von Berg, C. Lorenz, B. Loo, and P. Keall, "Four-dimensional computed tomography pulmonary ventilation images vary with deformable image registration algorithms and metrics." Medical Physics, vol. 38, pp. 1348–58, 2011.

[18] J. M. Reinhardt, K. Ding, K. Cao, G. E. Christensen, E. A. Hoffman, and S. V. Bodas, "Registration-based estimates of local lung tissue expansion compared to xenon CT measures of specific ventilation," Medical Image Analysis, vol. 12, no. 6, pp. 752 – 763, 2008.

[19] K. Ding, K. Cao, M. Fuld, K. Du, G. Christensen, E. Hoffman, and J. Reinhardt, "Comparison of image registration based measures of regional lung ventilation from dynamic spiral ct with Xe-CT," Medical Physics, vol. 39, p. 5084, 2012.

[20] P. Felzenszwalb and D. Huttenlocher, "Efficient belief propagation for early vision," International Journal of Computer Vision, vol. 70, pp. 41–54, 2006.

[21] V. Kolmogorov, "Convergent tree-reweighted message passing for energy minimization," Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 28, no. 10, pp. 1568 –1583, 2006.

[22] P. F. Felzenszwalb and D. P. Huttenlocher, "Pictorial structures for object recognition," International Journal of Computer Vision, vol. 61, pp. 55–79, 2005.

[23] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max- flow algorithms for energy minimization in vision," Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 26, no. 9, pp. 1124 –1137, 2004.

[24] N. Komodakis, G. Tziritas, and N. Paragios, "Fast, approximately optimal solutions for single and dynamic MRFs," in Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on, june 2007, pp. 1 –8.

[25] R. C. Prim, "Shortest connection networks and some generalizations," Bell System Technical Journal, vol. 36, pp. 1389–1401–244, 1957.

[26] O. Veksler, "Stereo correspondence by dynamic programming on a tree," in Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on, vol. 2, 2005, pp. 384–390.

[27] D. Rueckert, P. Aljabar, R. A. Heckemann, J. V. Hajnal, and A. Hammers, "Diffeomorphic registration using B-Splines," in Medical Image Computing and Computer-Assisted Intervention MICCAI 2006, ser. Lecture Notes in Computer Science, R. Larsen, M. Nielsen, and J. Sporring, Eds. Springer Berlin / Heidelberg, 2006, vol. 4190, pp. 702–709.

[28] V. Arsigny, O. Commowick, X. Pennec, and N. Ayache, "A log-euclidean framework for statistics on diffeomorphisms," in Medical Image Computing and Computer-Assisted Intervention MICCAI 2006, ser. Lecture Notes in Computer Science, R. Larsen, M. Nielsen, and J. Sporring, Eds. Springer Berlin / Heidelberg, 2006, vol. 4190, pp. 924–931.

[29] M. Chen, W. Lu, Q. Chen, K. J. Ruchala, and G. H. Olivera, "A simple fixed-point approach to invert a deformation field," Medical Physics, vol. 35, no. 1, p. 81, 2007.

[30] A. Sotiras and N. Paragios, "Discrete symmetric image registration," in Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on, 2012, pp. 342–345.

[31] D. Rueckert, L. Sonoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes, "Nonrigid registration using free-form deformations: application to breast MR images," Medical Imaging, IEEE Transactions on, vol. 18, no. 8, pp. 712 –721, aug. 1999.

[32] D. Crandall, A. Owens, N. Snavely, and D. Huttenlocher, "Discrete-continuous optimization for large-scale structure from motion," in Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, june 2011, pp. 3001 –3008.

[33] R. W. So, T. W. Tang, and A. C. Chung, "Non-rigid image registration of brain magnetic resonance images using graph-cuts," Pattern Recognition, vol. 44, no. 10, pp. 2450 – 2467, 2011.

[34] H. Robbins and S. Monro, "A stochastic approximation method," Annals of Mathematical Statistics, vol. 22, no. 3, pp. 400–407, 1951.

[35] D. Mahapatra and Y. Sun, "Integrating segmentation information for improved MRF-based elastic image registration," IEEE Transactions on Image Processing, vol. 21, no. 1, pp. 170–183, 2012.

[36] J. Ashburner and K. J. Friston, "Unified segmentation," NeuroImage, vol. 26, no. 3, pp. 839 – 851, 2005.

[37] A. Myronenko and X. Song, "Intensity-based image registration by minimizing residual complexity," Medical Imaging, IEEE Transactions on, vol. 29, no. 11, pp. 1882 –1891, nov. 2010.

[38] J. Modersitzki and S. Wirtz, "Combining homogenization and registration," in Biomedical Image Registration, Third International Workshop, WBIR 2006, ser. Lecture Notes in Computer Science, J. P. W. Pluim, B. Likar, and F. A. Gerritsen, Eds., vol. 4057. Springer, 2006, pp. 257–263.

[39] T. Guerrero, K. Sanders, J. Noyola-Martinez, E. Castillo, Y. Zhang, R. Tapia, R. Guerra, Y. Borghero, and R. Komaki, "Quantification of regional ventilation from treatment planning CT," International Journal of Radiation Oncology, Biology, Physics, vol. 62, pp. 630–634, 2005.

[40] S. Kabus, J. von Berg, T. Yamamoto, R. Opfer, and P. Keall, "Lung ventilation estimation based on 4D-CT imaging." in Proc. of the first international workshop on pulmonary image analysis, MICCAI, 2008, pp. 73–82.

[41] T. Yamamoto, S. Kabus, T. Klinder, C. Lorenz, J. von Berg, T. Blaffert, B. W. L. Jr, and P. J. Keall, "Investigation of four-dimensional computed tomography-based pulmonary ventilation imaging in patients with emphysematous lung regions," Physics in Medicine and Biology, vol. 56, no. 7, p. 2279, 2011.

[42] R. Castillo, E. Castillo, R. Guerra, V. Johnson, T. McPhail, A. Garg, and T. Guerrero, "A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets," Physics in Medicine and Biology, vol. 54, no. 7, p. 1849, 2009.

[43] K. J. Lee, D. Kwon, I. D. Yun, and S. U. Lee, "Deformable 3d volume registration using efficient MRFs model with decomposed nodes," in British Machine Vision Conference 2008, M. Everingham, C. J. Needham, and R. Fraile, Eds., 2008.

[44] A. Shekhovtsov, I. Kovtun, and V. Hlavác, "Efficient MRF deformation model for non-rigid image matching," Computer Vision and Image Understanding, vol. 112, no. 1, pp. 91–99, 2008.

[45] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 23, no. 11, pp. 1222 –1239, 2001.

[46] V. Kolmogorov and R. Zabin, "What energy functions can be minimized via graph cuts?" Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 26, no. 2, pp. 147 –159, 2004.

[47] J. Vandemeulebroucke, D. Sarrut, and P. Clarysse, "The POPI-model, a pointvalidated pixel-based breathing thorax model," in Conference on the Use of Computers in Radiation Therapy, 2007.

[48] A. Schmidt-Richberg, J. Ehrhardt, R. Werner, and H. Handels, "Fast explicit diffusion for registration with direction-dependent regularization," Biomedical Image Registration, pp. 220–228, 2012.

[49] C. Metz, S. Klein, M. Schaap, T. van Walsum, and W. Niessen, "Nonrigid registration of dynamic medical imaging data using nD+t B-splines and a groupwise optimization approach," Medical Image Analysis, vol. 15, no. 2, pp. 238 – 249, 2011.

[50] E. Castillo, R. Castillo, J. Martinez, M. Shenoy, and T. Guerrero, "Four-dimensional deformable image registration using trajectory modeling," Physics in Medicine and Biology, vol. 55, no. 1, p. 305, 2010.

[51] A. Schmidt-Richberg, J. Ehrhardt, R. Werner, and H. Handels, "Lung registration with improved fissure alignment by integration of pulmonary lobe segmentation," Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012, pp. 74–81, 2012.