Manifold Learning for 4D CT Reconstruction of the Lung

Manfred Georg *, Richard Souvenir †, Andrew Hope ‡, Robert Pless *

Washington University in St. Louis
* Saint Louis, MO 63130 USA
{mgeorg,pless}@cse.wustl.edu

† UNC Charlotte
Charlotte, NC 28223 USA
souvenir@uncc.edu

‡ University of Toronto
Toronto, ON, M5G 2M9, Canada
Andrew.Hope@rmp.uhn.on.ca

Abstract

Computed Tomography is used to create models of lung dynamics because it provides high contrast images of lung tissue. Creating 4D CT models which capture dynamics is complicated because clinical CT scanners capture data in slabs that comprise only a small part of the tissue. Commonly, creating 4D reconstruction requires stitching together different lung segments based on an external measure of lung volume. This paper presents a novel method for assembling 4D CT datasets using only the CT data. We use a manifold learning algorithm to parameterize each slab data with respect to the breathing cycle, and an alignment method to coordinate these parameterizations for different sections of the lung. Comparing this data driven parameterization with physiological measurements captured by a belt around the abdomen, we are able to generate slightly smoother reconstructions.

1. Introduction

CT imaging is often used for imaging tissues, such as lungs, that do not provide high contrast in MR-images. CT imaging is often used in lung cancer diagnosis and treatment, which has the additional problem that lung tumors move with the lung. Since radiation treatments affect both the tumor and the healthy tissues near the tumor, effective radiation treatment planning requires modeling the motion of the tumor and the surrounding healthy tissues.

Simple solutions to the motion issue include using fluoroscopy to monitor tumor movement and using fixed, breath hold images for treatment planning. In some cases, patient breathing has been controlled in an attempt to predict tumor position and create motion maps [11, 15]. Recently, four-dimensional CT (4D CT) has been developed, which uses an external measure of breath phase to create multiple CT datasets at various breath phases [1, 4, 5, 14]. This is more effective than fluoroscopy at determining tumor movement [13]. This technology continues to advance in complexity, but the need for an external measure of breathing has remained a key component of all methods thus far.

A standard way of creating a 4D model uses the following 3 step procedure [5, 6]. First, a common multi-slice clinical CT scanner is configured to capture a volume roughly one inch in height. As shown in figure 1, a cross sectional volume (slab) of the lung usually about one inch high is collected repeatedly from the same location in the lung while the patient breaths. The patient is then moved one inch through the CT scanner, to the next couch position, to begin capturing the next slabs. Each data acquisition takes approximately one second with several seconds between acquisitions. This is slow enough and patient breathing is irregular enough that breathing cannot be estimated from the time of image acquisition; instead, an external breath measurement such as a belt measuring abdomen circumference is used to estimate the volume of the patient's lung. Next, a complete 3D volume is generated for each part of the breathing cycle by picking a target lung volume, and choosing the data at each couch position acquired closest to the target lung volume. Figure 2 shows, on the left side, the lung volume measurements for each slab plotted against the CT scanner position with respect to the lung. For a target lung volume shown as a vertical line, the right side shows the reconstruction of the lung corresponding to the slabs closest to the target lung volume. Techniques such as optic flow [4] and template matching [2] can be used to compute the motion of tumor tissue.

The newest methods for creating 4D CT reconstructions of lung motion directly infer a deformation map from the slab data. These methods use different techniques to determine the reference volume, but all methods require that each slab is annotated with the lung volume measured at the time of acquisition [3, 8, 16].

In this paper, we present a substitute for the external breath measurement which will only require the image data itself. It is our intention that any algorithm that would use an external physiological measurement can instead be modified to use our breath measure. The breath measure stems

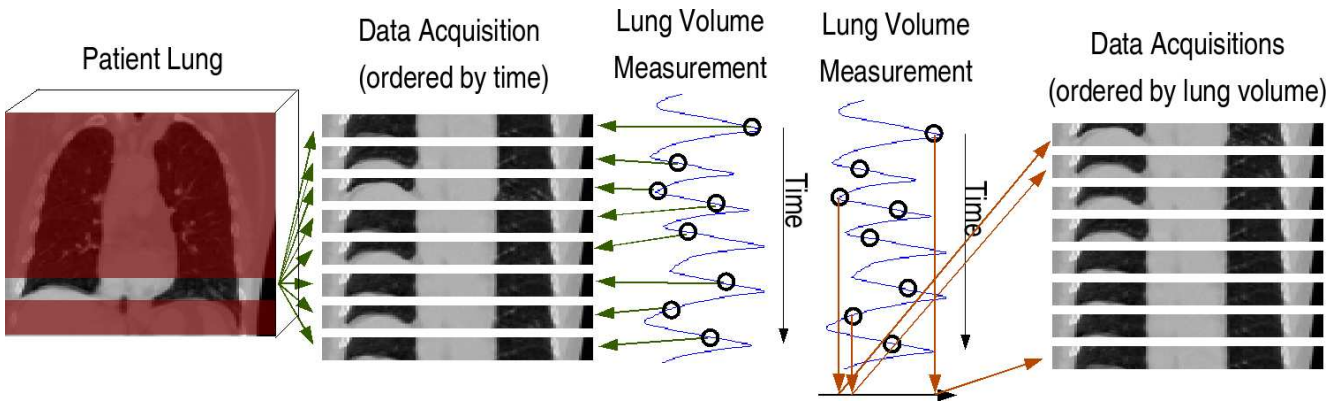

Figure 1. A CT scanner is used to repeatedly image slabs of the lung approximately 1 inch in height at the same location while a patient breaths. At the same time a belt around the abdomen of the patient measures lung volume. Although, the data acquisitions ordered by time does not have a coherent pattern with respect to breathing, the slabs can be reordered by the belt measurement.

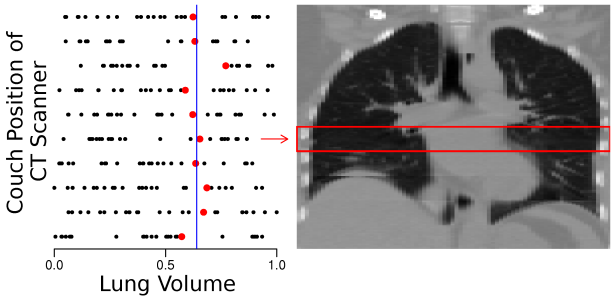

Figure 2. On the left, the external breath measurements for each 1 inch high image acquisition is shown organized by scanner couch position. A lung volume is created by picking the slab at each couch position of the CT scanner with a lung volume close to a target volume which is shown as a vertical line. On the right, a coronal cross section of the 3D lung reconstruction corresponding to the target lung volume is shown. The red box shows the size of image acquisitions at a single couch position. Several artifacts are visible near the diaphragm.

from an analysis of the intrinsic redundancy in the lung dataset from the same tissue being viewed at different parts of the breathing cycle. In the current work, we demonstrate a novel post-processing technique based on manifold learning that can generate breathing measurement estimates to allow 4D CT algorithms to be effective with just imaging data.

2. Prior Work

The use of 4D aware imaging and radiation treatment is a growing area with many applications (see [7] for a review). Due to a diversity of scanner technologies 4D CT of the lung can be executed in a number of different ways. Some examples include spiral and cine mode CT with an external breath measurement [5, 14]. On a lower level, the X-ray beam can be gated based on a physiological signal and specialized CT algorithms used to reconstruct the medical image [9].

McClelland *et al.* briefly investigated the creation of a breath measure from the slab data by measuring the anterior-posterior movement of the skin surface in the slab [8]. A sine wave is fit to this signal to create a breath measure. Ultimately this method was abandoned in favor of an external breath marker.

Our contribution is to use manifold learning to automatically determine breath phase. Manifold learning is a field of research that seeks to discover and parametrize the intrinsic degrees of freedom inside large data sets. Recently, computationally efficient algorithms such as Isomap [12] have made it possible to apply manifold learning to large data sets. The goal of Isomap is to assign each image a point in a low dimensional intrinsic parameter space such that similar images are close to each other without using any prior knowledge. This has previously been applied to cine-MR images to estimate the parameters of the breathing and heartbeat [17]. In that work all data items captured the entire region of interest and therefore could be parameterized simultaneously. In general, the parameterization given by Isomap is not consistently oriented between similar datasets. In the case of our lung data, this means that the breath measure learned by Isomap for each couch position is of an arbitrary scale and orientation requiring an additional alignment step.

3. Data

4D CT images were generated for 25 lung cancer patients using the following protocol. Transverse volumes of the lung were acquired using a 16 slice CT scanner in cine mode while the patient was free breathing. At each couch position, 25 slabs (512 by 512 by 16) were collected consecutively, before moving the patient to the next couch position and collecting 25 slabs of the next section of lung. Manual

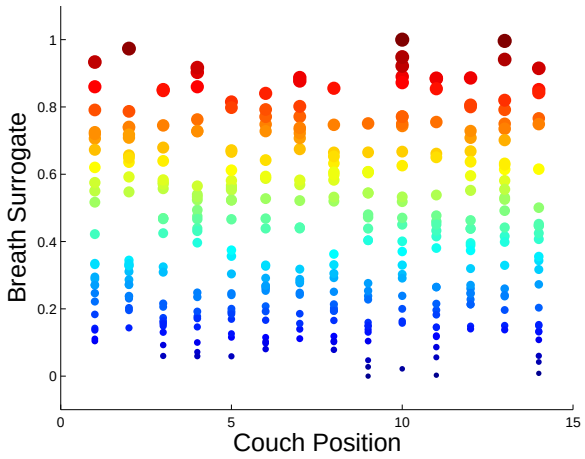

Figure 3. External breath measurement using a belt

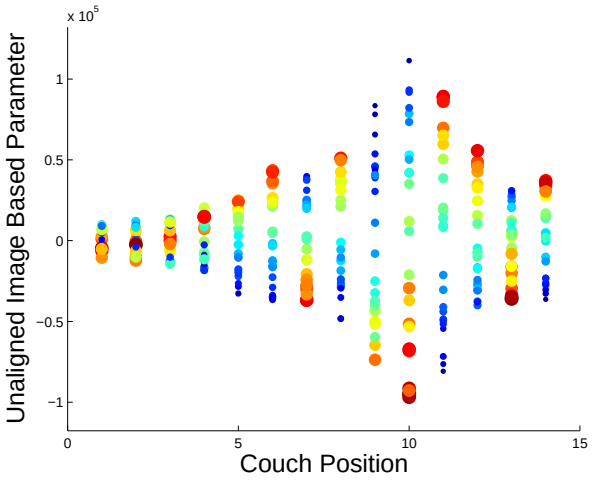

Figure 4. Isomap based breath measure

inspection was used to crop slabs to the smallest rectangular volume which contained lung tissue at any point in the breath cycle. The transverse plane was down-sampled to a factor of 0.3 of the original. On average the data volume used for processing was 70 by 130 by 176 voxels.

During image acquisition, an external breath surrogate was collected using a belt measuring abdomen circumference. Figure 3 shows the belt measurement for each image: the x-axis shows the position of the slab in the lung, while the y-axis gives the value of the belt measurement for a particular slab. The size and hue of the dot corresponds to the value of the belt measurement and therefore, in this case, is exactly equivalent to the y-axis. All later figures of breath measures will use the same format as this figure.

4. Algorithm

The estimation of the breath phase of a slab requires two parts. First, each couch position is analyzed independently to determine relative positions in the breath cycle for each slab. We call this a local breath measure. Second, the local breath measures from each couch position are combined into a globally consistent breath measure which is valid over all slabs from any couch position.

4.1. Measuring Breath Within a Couch Position

Although breathing is a cyclic process, we find that images of the lung at the same lung volume are similar whether the patient is inhaling or exhaling. This allows us to characterize the breath with a single parameter: lung volume. We apply the Isomap algorithm exactly as described in [12] to the slabs captured at a single couch position, using the square root of sum of squared distance (L_2 norm, and the k-nearest neighbors as a neighborhood criteria, with a value of 8 for k). We use the 1D parameterization produced by Isomap as the local breath measure for slabs computed at a given couch position. This breath measure is termed f_i, and defined as the mapping (computed through Isomap):

$$f_i : S_i \rightarrow \mathbb{R} \tag{1}$$

where S_i is the set of slabs at couch position i.

Figure 4 shows the local breath measure coordinates generated by Isomap, shown for different couch positions (labeled on the x-axis). Ground truth estimates (measured by the extension of a belt around the abdomen) are depicted by the size and hue of the points in the figure. This figure highlights that the Isomap parameter has a high correlation with the ground truth, but the Isomap parameters are sometimes oriented differently, and the scale also varies. This is because the Isomap algorithm parameterizes the slabs based on their similarity to each other: there is no global reference to orient the parameterization, and different parts of the lung have different overall contrasts, changing the relative scale.

4.2. Global Breath Measure

Since the local breath measure is only meaningful within a single couch position, extra work is needed to align the local breath measures to create a globally consistent breath measure valid over the entire lung. Figure 4 shows that the local breath measure and the belt measurements are related by a roughly linear relationship. This means that by scaling and translating the local breath measure with an affine transformation we can align all the local breath measures to create a global breath measure. We disregard until section 4.3 the problem of orienting each local breath measure, and focus on the scaling problem.

As an example of one way of scaling the local breath measures to create a global breath measure we naively rescale the range of each local breath measure to be 0 to

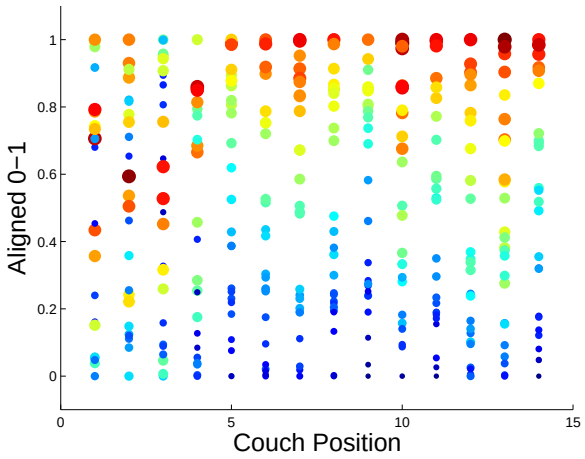

Figure 5. Isomap based breath measures scaled to range [0,1]

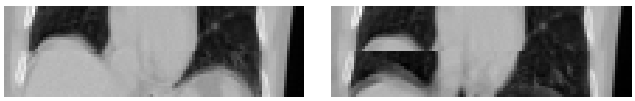

(a) ρ = 18 (b) ρ = 42

Figure 6. Coronal cross sections of two slabs from adjacent couch positions that come from (a) the same breath phase (b) different breath phases as determined by the discontinuity measure.

1 as shown in figure 5. If each couch position has a slab from both complete inhalation and complete exhalation, this would create a consistent global breath measure. Unfortunately, patients do not breath in a perfectly regular pattern, and for some couch positions the patients do not reach close to maximum inhale or exhale. Thus, this alignment process must be based on comparison of the slabs captured at neighboring couch positions.

The basic idea of the algorithm is to identify slabs from adjacent couch positions that were captured from similar parts of the breathing cycle. These slabs can be put together and should show no large discontinuities at their boundary. Thus, we define a discontinuity measure that finds for each slab, the two slabs in the adjacent couch position that are from similar parts of the breathing cycle.

We then search for the affine transformation (the scaling and translation) of the parameterization at one couch position so that each slab is best aligned with the slabs at the next couch position taken at the same part of the breathing cycle.

4.2.1 Measuring Boundary Discontinuities

slabs with similar breath phases should have little discontinuity between their borders when they are placed next to each other. Figure 6 shows two pairs of slabs one pair matches well and comes from the same breath phase while

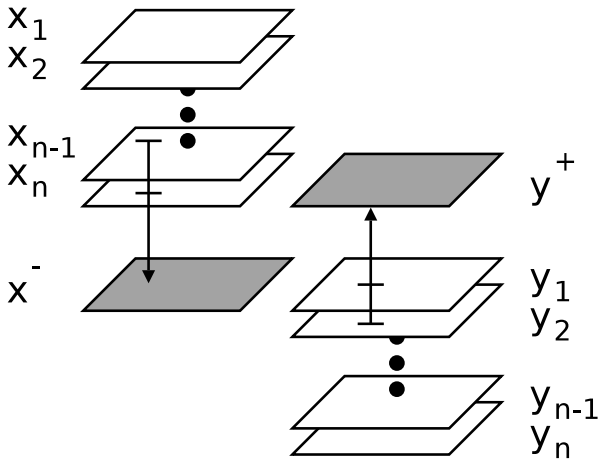

Figure 7. The discontinuity measure between two slabs is created by using the lower slices of the upper slab to predict what the slice directly below it will look like (x^-) and comparing that slice to the top slice of the bottom slab (y_1). The reverse computation of comparing y^+ and x_n is also performed.

the other pair has a lot of border discontinuity and does not come from the same breath phase.

We denote the set of slabs within a couch position as S_i. For simplicity we consider only the case where S_1 is above S_2, but a similar formulation is possible for the opposite case. Given two slabs $x \in S_1$ and $y \in S_2$ we compute a discontinuity measure over the boundary between the slabs. We number the transverse slices which comprise the slab x from top to bottom as $x_1, x_2, \cdots, x_n$. The discontinuity measure is computed by predicting the slice directly below the upper slab and comparing it to the top slice of the bottom slab and vice versa as shown in figure 7. The bottom two slices of x are used to linearly predict the slice below them which we denote as x^- ($x^- = 2x_n - x_{n-1}$). This slice is then compared to the top most slice in y. The top two slices in y are also used to predict the slice above them ($y^+ = 2y_1 - y_2$). The discontinuity measure ρ is the sum of the L_2 distance between the predictions and the actual slices from the adjacent slab.

$$\rho(x, y) = \|x^- - y_1\| + \|y^+ - x_n\| \quad (2)$$

When the discontinuity measure ρ is small as in figure 6(a) the two slabs are assumed to come from the same breath phase.

The discontinuity measure can be used to create a paired list of matching slabs which are likely to come from the same breath phase. For each slab $x \in S_1$ the two slabs in S_2, $y_{x,1}$ and $y_{x,2}$, which minimize the discontinuity measure are found and the pairs $(x, y_{x,1})$ and $(x, y_{x,2})$ are added to the match list. Conversely, for each slab $y \in S_2$ the two slabs in S_1 which are least discontinuous with y are found

and the pairs $(x_{y,1}, y)$ and $(x_{y,2}, y)$ added to the match list. This paired list of matching slabs is used to align the local breath measures. We find that using the two best matches rather than just the first yields a more robust solution and that computing the matches in both directions ensures that all slabs in both couch positions are considered.

4.2.2 Aligning Local Breath Measures

A distance measure is defined between slabs as the squared distance between the local breath measures $f_1(x)$ and $f_2(y)$ of the two slabs after alignment by the transformation ϕ.

$$D(x, y) = (f_1(x) - \phi(f_2(y)))^2 \tag{3}$$

The affine transformation which we use to align the coordinate spaces is a function with two parameters: the scaling value a and the translation value b.

$$\phi : v \mapsto av + b \tag{4}$$

We require a cost function which measures the quality of an aligning transformation ϕ between the local breath measures of two couch positions. The cost function for a transformation is the distance between the local breath measures after the transformation for each pair of slabs in the match list. The distances are divided by the discontinuity measure in order to minimize the impact of slabs which do not have good corresponding matches on the alignment.

$$\begin{aligned} C(a,b) \;=\; & \sqrt{\sum_{x \in S_1} \frac{D(x, y_{x,1})}{\rho(x, y_{x,1})} + \frac{D(x, y_{x,2})}{\rho(x, y_{x,2})}} \\ & + \sqrt{\sum_{y \in S_2} \frac{D(x_{y,1}, y)}{\rho(x_{y,1}, y)} + \frac{D(x_{y,2}, y)}{\rho(x_{y,2}, y)}} \end{aligned} \tag{5}$$

To find the parameters a and b of the affine transformation which minimize the cost function we use sequential quadratic programming, a standard constrained optimization algorithm (implemented in Matlab as `fmincon`).

$$\arg\min_{a,b} C(a, b), \quad a > 0 \tag{6}$$

The constraint $a > 0$ is used so that the transformation cannot invert exhalation and inhalation for any local breath measure from that which was set explicitly using the technique outlined in section 4.3.

4.2.3 Aligning all Couch Positions

To align the local breath measures of all couch positions we start with the central couch position (position 7 in the figures) and rescale it to the range 0 to 1. Then, moving outward from this couch position, we align the local breath measures of each couch position to that of its neighbor closest to the center. This yields a global breath measure which consistently describes the breath phase for every slab in all couch positions as shown in figure 8.

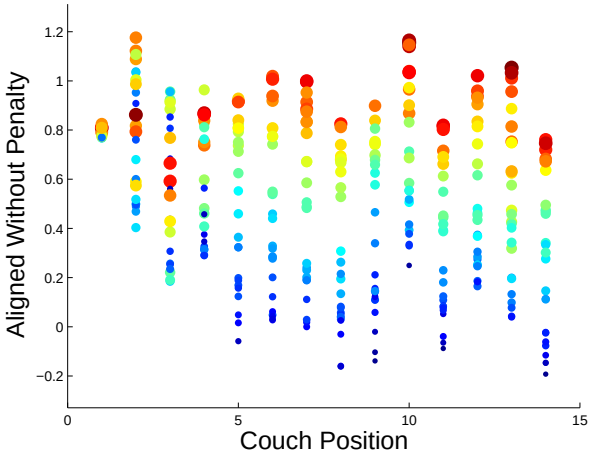

Figure 8. Breath measure without penalty function

4.3. Orienting the Local Breath Measures

The local breath measure of each couch position has an uncertain orientation with respect to inhalation and exhalation. Therefore, we explicitly orient the local breath measures before optimization.

First, the couch position with the largest range of local breath measure is found. Since the scale on the local breath measures is based on L_2 distances on the slabs, the couch position with the greatest range will have the most dramatic voxel value changes, which is caused by the most diaphragm movement. The slabs with the largest and smallest breath measures will display the most difference in amount of diaphragm tissue they contain. In the inhaled position, less diaphragm tissue is visible causing the average voxel value to be low. In the exhaled position, more diaphragm tissue is visible and the average voxel value will be high. The local breath measure is inverted if necessary so that a small breath measure value corresponds to an exhaled state.

The following method is used to explicitly orient all local breath measures to be consistent with that of the couch position with the diaphragm. For two adjacent couch positions, determine the correlation between the local breath measures of all pairs of slabs in the match list. If this correlation is negative, then the breath measure is inverted. Moving outward from the couch position with the most diaphragm movement, this procedure is repeatedly applied until all local breath measures are correctly oriented.

(a) Unaligned (b) Scaled 0-1 (c) Aligned, No Penalty (d) Aligned (e) Breath Surrogate

Figure 9. Coronal cross-section of a reconstructed image halfway between inhalation and exhalation using several methods.

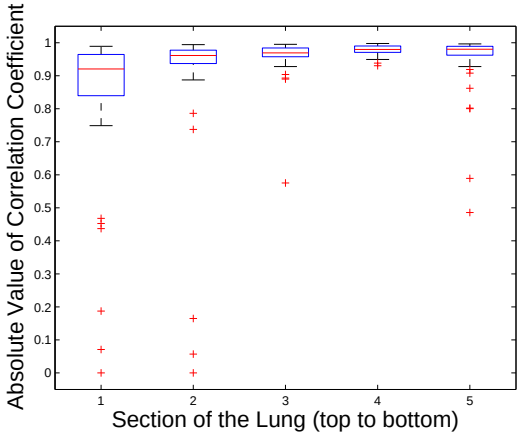

Figure 10. Correlation coefficient of the Isomap based breath measure and the belt measurements for each couch position for 25 patients. The lung is grouped into 5 sections from the top to the bottom. A boxplot of the absolute value of the correlation coefficient is shown for each section.

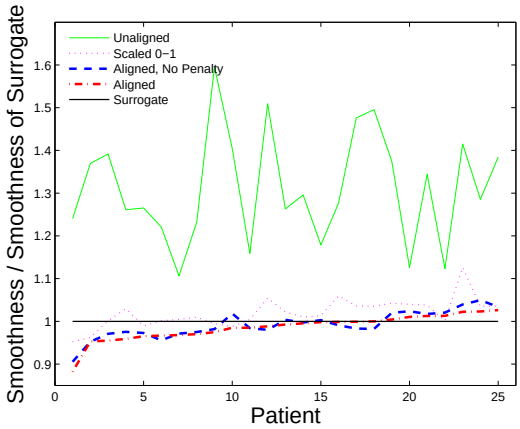

Figure 11. Discontinuity metric for reconstructions based on different breath measures.

5. Discussion

The ability to parameterize the slabs captured during lung CT is a key pre-processing step to many subsequent analysis steps. In this section we discuss applications of the automatic estimation of lung volume, explicit comparison to physiological models, and heuristics that improve the overall alignment.

5.1. Conventional Lung Models and 4D CT

The global breath measure can be used to create a lung image at a particular lung volume in the same way as the belt measurements can be used. The slab from each couch position with global breath measure closest to a target value is chosen. Figure 9 shows lung images obtained halfway between exhalation and inhalation for the different options of computing the breath measure discussed in the last section.

5.2. Physiological Verification of Breath Measure

We next compare the global breath measure against estimates of the lung motion captured by measuring the extension of a belt around the abdomen of the patient. Figure 10 shows the correlation of the breath measure with the belt measurements for all couch positions of all 25 patients grouped by location in the lung. The correlation is worse at the top of and above the lung where there is little breath related motion.

We compare images using the Isomap based breath measure to images based on the belt measurements. The metric used for comparison is the artifacts across each couch position boundary in the reconstruction computed by the discontinuity measure of equation 2. Figure 11 shows the ratio of boundary discontinuity for a reconstruction based on a particular breath measure to the discontinuity of the reconstruction based on the physiological breath surrogate. Patients are sorted by the ratio for the aligned with penalty term reconstruction. These results show that using the Isomap parameterization of breathing produces reconstructions with comparable (or slightly less) discontinuities at the boundaries between couch positions.

5.3. Local Breath Measures

The reconstruction algorithm hinges on the ability of Isomap to parameterize the breath phase of slabs. Several methods other than Isomap were investigated. Figure 12

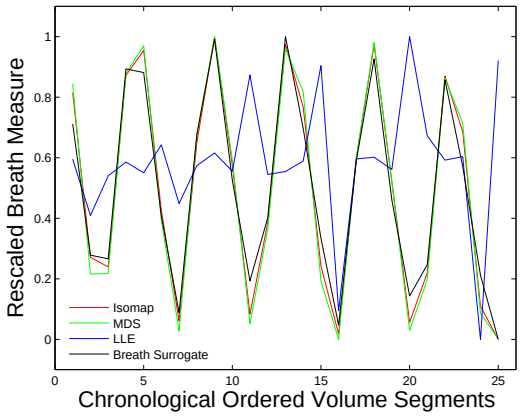

Figure 12. Several different breath measures collected at a single scanner position near the diaphragm. All breath measures are rescaled to the range [0,1].

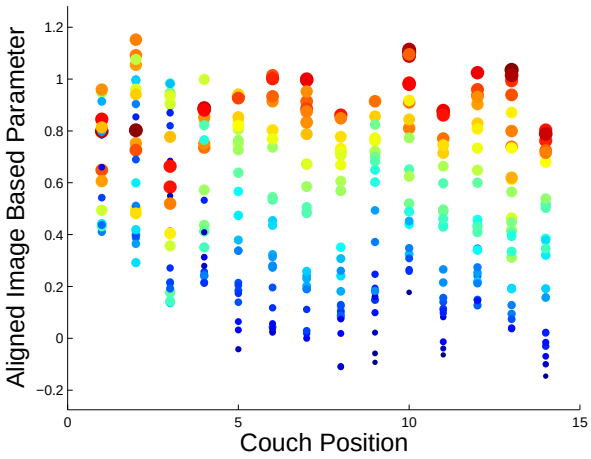

Figure 13. Breath measure with standard penalty function

shows the breath measures obtained for a single couch position near the diaphragm by different methods. Multidimensional scaling which is identical to Principal Component Analysis (PCA) in this setting and upon which Isomap is based has a very similar output to Isomap. The small number of slab samples requires a relatively large neighborhood in the graph used by Isomap, causing the graph distances to degenerate into Euclidean distances and Isomap to be equivalent to MDS. Locally Linear Embedding (LLE) produces a breath measure significantly worse than the other methods [10]. This is probably due to the small number of samples at each couch position.

5.4. Improved Global Alignment

Since alignments are only performed between neighbors, errors can compound from one end of the lung to the other. This problem is mitigated by starting the alignment process at the center of the lung, where less steps are chained together and the correlation with breathing is stronger. In figure 8, the first couch position has a much smaller range in its breath space than the other couch positions. This is because the closest two slabs in the neighboring couch position are the same for all slabs. When this occurs, the cost function favors shrinking the transformed space to have a small range which is close to the mean.

This type of problem can be explicitly corrected by adding a penalty term to the cost function to maintain the range of the breath measure near 0 to 1.

$$C' = C + \frac{(\max(f_i(S_i)) - 1)^2 + (\min(f_i(S_i)) - 0)^2}{\alpha} \tag{7}$$

The scaling constant α is set to 20 in our experiments. The optimization without penalty term is taken as the initial condition for the optimization with penalty so as to avoid convergence to a local minimum caused by the penalty term. Figure 13 shows the breath measure with this penalty term.

5.5. Orientation by Breath Marker

We can verify the accuracy of the local breath measure orientation algorithm by correlating the breath measures with the belt measurements. For all but one couch position of one patient we are able to correctly orient the breath measure, using the method presented earlier. The patient on which our image based method fails is breathing erratically and the method fails on a couch position with a correlation of 0.05.

6. Conclusion

Previous work in 4D CT of a lung has required a physiological measurement of the lung volume in order to create reconstructions of the lung. We presented a method for modeling a lung as it breaths using only the data captured by the CT scanner. Our experiments highlight that estimated lung volume correlates well with the physiological measurement and gives 3D reconstructions that are similar (or marginally better) in quality than current methods. This allows a collection of retrospective data analyses to be done in cases where physiological data was not captured, and offers the potential for new data acquisition protocols.

References

[1] J. Ehrhardt, R. Werner, D. Säring, T. Frenzel, W. Lu, D. Low, and H. Handels. An optical flow based method for improved reconstruction of 4D CT data sets acquired during free breathing. *Med Phys*, 34(2):711–21, 2007.

[2] M. Georg, J. J. Cannon, A. J. Hope, W. Lu, D. A. Low, and R. B. Pless. Automating 4D CT Reconstruction Using Manifold Learning. In *American Radium Society Annual Meeting*, May 2007.

[3] M. Georg, R. Souvenir, A. Hope, and R. Pless. Simultaneous Data Volume Reconstruction and Pose Estimation from Slice Samples. In *Computer Vision and Pattern Recognition*, June 2008.

[4] T. Guerrero, G. Zhang, T.-C. Huang, and K.-P. Lin. Intrathoracic tumour motion estimation from CT imaging using the 3D optical flow method. *Physics in Medicine and Biology*, 49:4147–4161, August 2004.

[5] H. Handels, R. Werner, T. Frenzel, D. Säring, W. Lu, D. Low, and J. Ehrhardt. Generation of 4D CT Image Data and Analysis of Lung Tumour Mobility during the Breathing Cycle. *Stud Health Technol Inform*, 124, 2006.

[6] A. J. Hope, M. Georg, J. J. Cannon, J. Hubenschmidt, W. Lu, D. A. Low, and R. B. Pless. Applications of Manifold Learning Techniques in 4D-CT reconstruction. In *International Conference on the use of Computers in Radiation Therapy*, June 2007. Reviewer's Choice.

[7] G. Li, D. Citrin, K. Camphausen, B. Mueller, C. Burman, B. Mychalczak, R. W. Miller, and Y. Song. Advances in 4D Medical Imaging and 4D Radiation Therapy. *TCRT*, 7(1):67–82, February 2008.

[8] J. R. McClelland, J. M. Blackall, S. Tarte, A. C. Chandler, S. Hughes, S. Ahmad, D. B. Landau, and D. J. Hawkes. A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy. *Medical Physics*, 33(9):3348–3358, September 2006.

[9] T. Nielsen, R. Manzke, R. Proksa, and M. Grass. Cardiac cone-beam CT volume reconstruction using ART. *Medical Physics*, 32(4):851–860, April 2005.

[10] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, December 2000.

[11] D. Sarrut, V. Boldea, S. Miguet, and C. Ginestet. Simulation of four-dimensional CT images from deformable registration between inhale and exhale breath-hold CT scans. *Medical Physics*, 33(3):605–617, March 2006.

[12] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[13] Y. G. van der Geld, S. Senan, J. R. van Sörnsen de Koste, H. van Tinteren, B. J. Slotman, R. W. Underberg, and F. J. Lagerwaard. Evaluating mobility for radiotherapy planning of lung tumors: A comparison of virtual fluoroscopy and 4DCT. *Lung Cancer*, 53(1):31–37, July 2006.

[14] S. S. Vedam, P. J. Keall, V. R. Kini, H. Mostafavi, H. P. Shukla, and R. Mohan. Acquiring a four-dimensional computed tomography dataset using an external respiratory signal. *Physics in Medicine and Biology*, 48(1):45–62, January 2003.

[15] J. W. Wong, M. B. Sharpe, D. A. Jaffray, V. R. Kini, J. M. Robertson, J. S. Stromberg, and A. A. Martinez. The use of active breathing control (ABC) to reduce margin for breathing motion. *International Journal Radiation Oncolology Biology Physics*, 44(4):911–919, July 1999.

[16] S. Xu, R. Taylor, G. Fichtinger, and K. Cleary. Lung Deformation Estimation and Four-dimensional CT Lung Reconstruction. *Academic Radiology*, 13(9):1082–1092, September 2006.

[17] Q. Zhang, R. Souvenir, and R. Pless. On Manifold Structure of Cardiac MRI Data: Application to Segmentation. *IEEE Conference on Computer Vision and Pattern Recognition*, 1:1092–1098, 2006.