Locally Switching Between Cost Functions in Iterative Non-rigid Registration

William Mullally[1], Margrit Betke[1], Carissa Bellardine[2], and Kenneth Lutchen[2]

[1] Computer Science Department, Boston University
[2] Department of Biomedical Engineering, Boston University,
Boston, MA 02215, USA

Abstract. In non-rigid image registration problems, it can be difficult to construct a single cost function that adequately captures concepts of similarity for multiple structures, for example when one structure changes in density while another structure does not. We propose a method that locally switches between cost functions at each iteration of the registration process. This allows more specific similarity criteria to be embedded in the registration process and prevents costs from being applied to structures for which they are inappropriate. We tested our method by registering chest computed tomography (CT) scans containing a healthy lung to scans of the same lung afflicted with acute respiratory distress syndrome (ARDS). We evaluated our method both visually and with the use of landmarks and show improvement over existing methodology.

1 Introduction

Registration methods are increasingly in demand by medical practitioners to accurately model the transformations they observe in their image data [1]-[4]. The problem of registering lung images has recently gained notice [5]-[9] largely due to an interest in supporting lung cancer diagnosis. Accurate registration of lung anatomy, however, remains an open problem, especially in the presence of pervasive pathology where normal structures have changed drastically in appearance as when the lung is afflicted with emphysema or acute respiratory distress syndrome (ARDS). In this work, we present a method for registering a chest CT containing a healthy lung to a chest CT of the same lung afflicted with ARDS and present a general framework to solve non-rigid registration problems in which the intensity values of corresponding structures do not change in the same way for all structures in the images.

ARDS is a disease that involves severe flooding and collapse of the lung, thereby making it difficult to breathe and maintain adequate gas exchange. ARDS has a high mortality rate (32% – 45%) [10]. Mechanical ventilation is often a required therapy in order to maintain the necessary O_2 and CO_2 levels of the body. Mechanical ventilation can impose non-physiological forces on the lung that can exacerbate lung injury. Characterizing the heterogeneous nature of the disease and how it is impacted by mechanical ventilation is important for basic survival of these patients [11]. On CT images ARDS is characterized by

Y. Liu, T. Jiang, and C. Zhang (Eds.): CVBIA 2005, LNCS 3765, pp. 367–377, 2005.

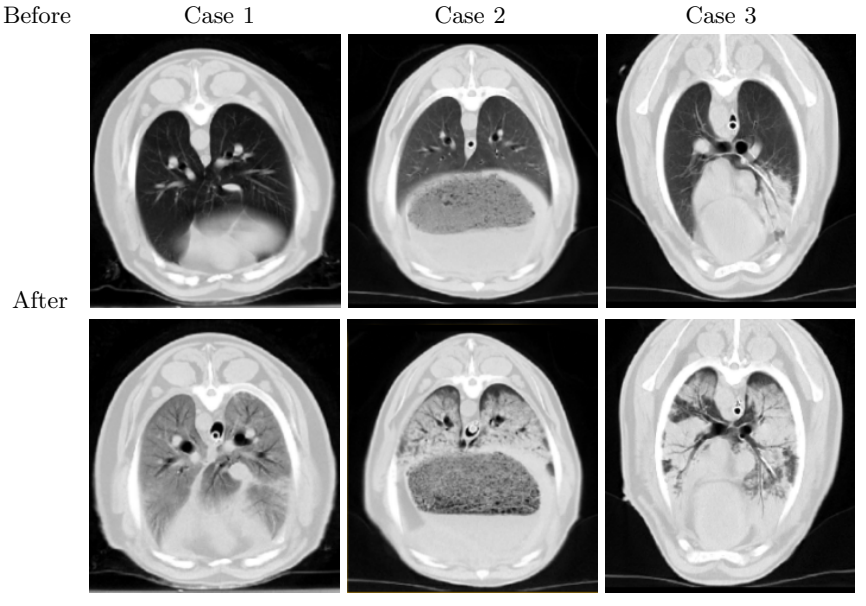

Fig. 1. Axial view of three sheep chests before and after a saline wash of their lungs induced ARDS. A saline wash affected the entire lung of case 1, but did not uniformly affect the lungs of cases 2 and 3 as can be seen from the dark air-filled patches at the top of case 2's lungs and throughout case 3's lungs.

regions within the lung that can be as dense as the tissue surrounding the lung (Fig. 1). Measures on the mechanics of ARDS have been limited to the order of quadrants within the lung [11] rather than the order of the finer anatomical structures visible even in low resolution CT. Non-rigid registration of healthy lung scans to scans of lungs with ARDS can help researchers understand the syndrome and may lead to better treatment options than currently exist. To the best of our knowledge, automatic non-rigid registration approaches have not been applied to this problem.

There are two principle branches of image registration techniques: those guided by image similarity cost functions and those guided by feature detectors. This paper presents a modification to non-rigid registration methods using image-based cost functions. Specifically, we modify the method proposed by Rueckert et al. [12] in which images are aligned by minimizing the "cost" of correspondence between the images by comparing intensity values. In registration problems like this, a similarity function that might be adequate for capturing the spatial deformations of temporally constant density tissue, for example the bony anatomy, might not be sufficient to capture both spatial deformations of structures and temporal changes in tissues density, as in an ARDS inflicted lung. We do not attempt to solve the correspondence problem of lung anatomy imaged before and after the onset of ARDS. Instead, we propose a method to non-rigidly register the tissues surrounding the

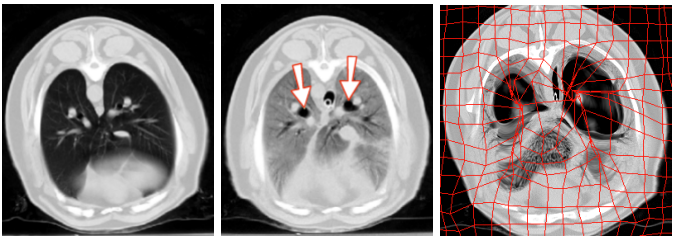

Fig. 2. Axial view of a healthy lung (left) and an ARDS inflicted lung before (middle) and after (right) non-rigid registration using correlation as the only cost to guide registration. The transformation grid is overlaid in the right image. The voxels of the bronchi (indicated by arrows in their original form in the middle image) are inappropriately warped to fill the lung (dark regions in right image). In the iterative registration process, the correlation gradient leads away from correct correspondence within the lung. Correspondence outside the lung is somewhat better.

lung without adversely affecting the topology of the lung. No segmentation is used in our approach. Our approach adaptively changes the character of the cost function at each iteration of the registration process. These changes allow cost functions to be applied only on appropriate anatomy.

2 Methods

We follow the framework proposed by Rueckert et al. [12] in which a multi-level non-rigid registration algorithm is built upon a hierarchical grid of b-splines. Multi-level approaches progress from coarse to fine image resolutions during the process of registration to avoid local minima problems and reduce computation time. A hierarchical b-spline grid controls the transformation T of the source image I_1 into the coordinate system of the target image I_2. The similarity between the two images is evaluated with a cost function of the form:

$$Cost(\theta) = C_{ImageSimilarity}(T(\theta, I_1), I_2) + C_{Regularization}(T(\theta)),$$

where θ designates the positions of the grid points controlling the b-spline. The gradient of this cost function is followed in steps of size μ to find the optimal registration between the two images. The registration process iteratively checks the gradient and modifies the transformation parameters.

The above approach will not work when the cost function does not appropriately embody all the changes that occur between image acquisitions. In Figure 2, for example, correlation was used as the only cost to guide the registration between healthy and ARDS inflicted lungs. The correlation gradient, in effect, pushed the transformation in an anatomically inappropriate direction.

We would like to construct a framework for using cost functions only on the anatomy for which a particular cost function is most appropriate. We can do this by associating each knot θ_i of the b-spline grid for $i = 0, ..., \mathrm{N}$, where N is the

number of grid control points, with a particular cost function. For $k = 1, ..., \mathrm{K}$ cost functions, this takes the form:

$$Cost(\theta_i) = \begin{cases} Cost_1(\theta) & \text{if } f_1(\theta_i, T(\theta, I_1), I_2) > \tau_1 \\ Cost_2(\theta) & \text{if } f_2(\theta_i, T(\theta, I_1), I_2) > \tau_2 \\ \vdots & \\ Cost_K(\theta) & \text{otherwise} \end{cases} \tag{1}$$

where f_k are $K - 1$ functions used to decide which cost function to apply by comparing each of them to a threshold τ_k. The cost function association of each control point is decided at each iteration. The algorithm, a modification of Rueckert et al.'s framework that includes our method of cost function selection, is detailed below.

Non-rigid Multi-scale Registration Algorithm with Local Cost Switching

Input: Two 3D scans

calculate an initial rigid registration
initialize the control points $\theta^{(l)}$ and down sample the images to the coarsest resolution, $l = 1$,
repeat
  determine the cost function $Cost(\theta_i) = Cost_k(\theta)$ to associate with each control point θ_i (Eq. 1)
  calculate $\nabla C = \frac{\partial Cost(\theta_i^{(l)})}{\partial \theta_i^{(l)}}$
  for $\mu =$ max. step size **until** $\mu =$ min. step size
    while $\|\nabla C\| >$ threshold ϵ **or** iteration count $<$ max. iterations **do**
      recalculate the control points $\theta = \theta + \mu \frac{\nabla C}{\|\nabla C\|}$
      determine the cost function $Cost_k(\theta)$ to use at each control point θ_i
      recalculate the gradient vector ∇C
    decrease step size μ
  increase the control point resolution and image resolution l
until finest level of resolution, $l = L$, is reached

Output: θ

We adapted the local cost switching algorithm to the specific application of registering healthy lungs to lungs inflicted with ARDS. We needed to construct cost functions that could appropriately express tissue transformations in this problem. First, we incorporated transformation costs that discourage dramatic bending of the transformation grid and extreme volume changes across the entire image. Secondly, the correspondence of anatomy outside of the lung can be captured by an image-intensity similarity term. We constructed a single cost function incorporating these terms and weighted their contributions to the overall cost using α, β, and γ as follows:

$$\begin{aligned} Cost() = \alpha\, C_{ImageSimilarity}(T(\theta, I_1), I_2) + \beta\, C_{BendingEnergy}(T(\theta)) \\ + \gamma\, C_{VolumePreservation}(T(\theta)) \end{aligned} \tag{2}$$

Specifically, we define $C_{ImageSimilarity}()$ to be the correlation coefficient $\frac{Covariance(I_1, I_2)}{Std.Deviation(I_1) * Std.Deviation(I_2)}$, $C_{BendingEnergy}(T(\theta))$ to be the bending energy

$\frac{1}{V}\int_0^X \int_0^Y \int_0^Z [(\frac{\partial^2 T}{\partial x^2})^2 + (\frac{\partial^2 T}{\partial y^2})^2 + (\frac{\partial^2 T}{\partial z^2})^2 + 2(\frac{\partial^2 T}{\partial xy})^2 + 2(\frac{\partial^2 T}{\partial xz})^2 + 2(\frac{\partial^2 T}{\partial yz})^2] dxdydz$ of a thin metal plate [12], where X, Y, and Z are the dimensions and V is the volume of the scan, and $C_{VolumePreservation}(T(\theta)) = \int_V |\log(det(Jacobian(T(\theta))))|\, d\theta$ [13].

For the single cost in Eq. 2 to be successful if applied by itself, it must not only relate the anatomy outside the lung which has undergone minimal intensity changes but also the anatomy within the lung which has drastically changed because of the infliction of ARDS. In general, $C_{ImageSimilarity}(T(\theta, I_1), I_2)$ is used to capture such a relationship, but as already shown in Fig. 2, correct anatomic alignment cannot be achieved using the correlation coefficient as the image similarity cost. Moreover, it may not be possible to construct such an image-intensity similarity function that can capture the entire range of anatomical changes seen in comparing scans of healthy lungs to scans of diseased lungs. In such cases where image similarity costs guide registration toward inappropriate solutions, our cost switching framework can be used to allow regions with strong image correlation to non-rigidly align themselves while still correctly deforming regions with weak image similarity. To formulate this, we specify Eq. 1 for two cost functions, $Cost_1$ with an image similarity term and $Cost_2$ without:

$$\begin{aligned} Cost_1() &= \alpha\, C_{ImageSimilarity}(T(\theta, I_1), I_2) + \beta\, C_{BendingEnergy}(T(\theta)) \\ &\quad + \gamma\, C_{VolumePreservation}(T(\theta)) \\ Cost_2() &= \beta\, C_{BendingEnergy}(T(\theta)) + \gamma\, C_{VolumePreservation}(T(\theta)), \end{aligned} \tag{3}$$

Because Eq. 1 is applied for $K = 2$, we had to specify one decision function f_1 for Eq. 3. We defined f_1 to be the normalized correlation coefficient with threshold $\tau_1 = 0.5$. This function controls how to switch between the two cost functions at each grid point and differentiates between correlated and uncorrelated image regions. A fixed size patch around the location of each knot point was used as the correlation template. An initial alignment is achieved by performing a rigid registration on high density bony anatomy. The two costs defined in Eq. 3 are assigned in lines 4 and 9 of our cost switching registration algorithm.

It should be noted that with the costs in Eq. 3, we do not expect the anatomy within the injured portions of the lung to achieve perfect alignment. There is no mechanism to draw these regions into correct correspondence except for the bending and volume preservation costs responding to the deformations of surrounding tissues. Tissue outside of these injured regions should achieve good alignment.

3 Experiments

To evaluate our cost switching approach, we used CT scans taken from six sheep (Fig. 1). CT scans were taken both before and after ARDS was induced by treating their lungs with a saline wash. The sheep were placed on ventilators to control the pressure level of air in their lungs. The CT images had resolution $0.71 \times 0.71 \times 10$ mm^3 and captured the entire area of the lung.

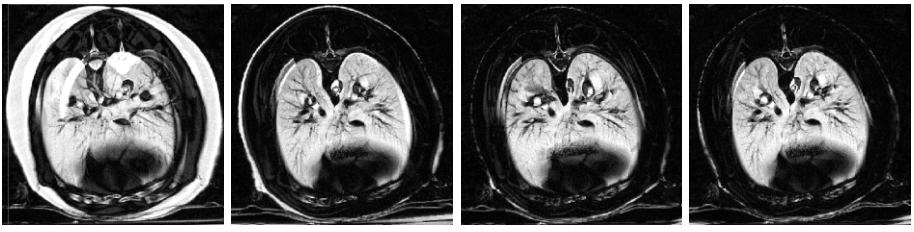

Fig. 3. Absolute difference in intensity values between unregistered scans (left), rigidly registered scans (middle left), single cost non-rigid registration (middle right), and cost switching non-rigid registration (right). The lungs are visible as a result of the density shift induced by the saline wash. Outside of the lungs, the non-rigid approaches improve upon the results of the rigid registration and in this case, our cost switching approach more accurately registered the ribs.

Table 1. Root Mean Squared Error Between Registered Landmarks in the Lung (mm). Rigid and cost switching non-rigid registration driven by non-lung structures result in similar error in lung landmark alignment.

Sheep	Unregistered	Rigid	Single Cost	Cost Switching
1	38.5	15.4	11.9	14.2
2	10.2	7.9	12.1	8.1
3	7.9	6.7	13.3	5.7
4	20.3	8.1	7.2	7.8
5	13.2	9.9	16.8	11.1
6	27.5	9.8	24.2	14.5
Average	19.6	9.6	14.3	10.2

We tested our method by registering images of the same animal imaged at the same pressure level before and after the infliction of ARDS. We validated our registration results both visually and by comparison to ground truth correspondence of landmarks. Figure 3 shows that non-rigid registration approaches reduced the intensity difference between scans for the anatomy outside the lung. As the anatomy outside the lung has not changed in density between imaging, a reduction in the intensity difference reflects a more accurate alignment of anatomy.

Within the lung, since significant density changes have occurred between data acquisitions, density differences are not a useful measure of misalignment. We can, however, evaluate changes in landmark location. The landmarks used consist of branching points of the bronchi and blood vessels within the lung. We used 18-22 landmarks per case. We show the root mean squared error in millimeters between corresponding landmarks in Table 1. We report an average improvement in landmark alignment of 28% by the cost switching approach in comparison to the single cost approach. We report an average degradation in landmark alignment of 6% for the cost switching approach in comparison to the

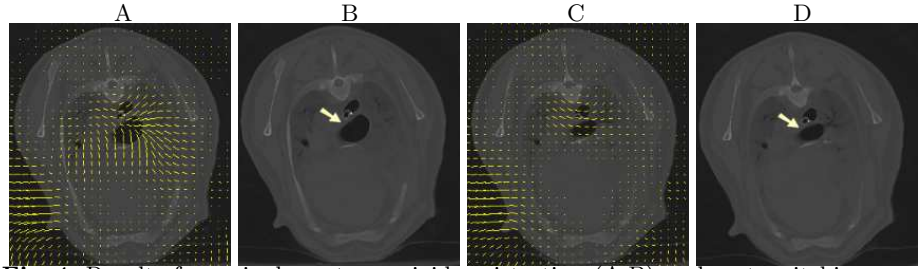

Fig. 4. Results from single cost non-rigid registration (A,B) and cost switching non-rigid registration (C,D) in sheep 1. We show the registered images with an overlay of the local deformation quantities (A,C). Notice the difference in the shape of the trachea (indicated with an arrow). The single cost approach greatly increased the size of the trachea. Our approach stayed closer to physically viable transformation.

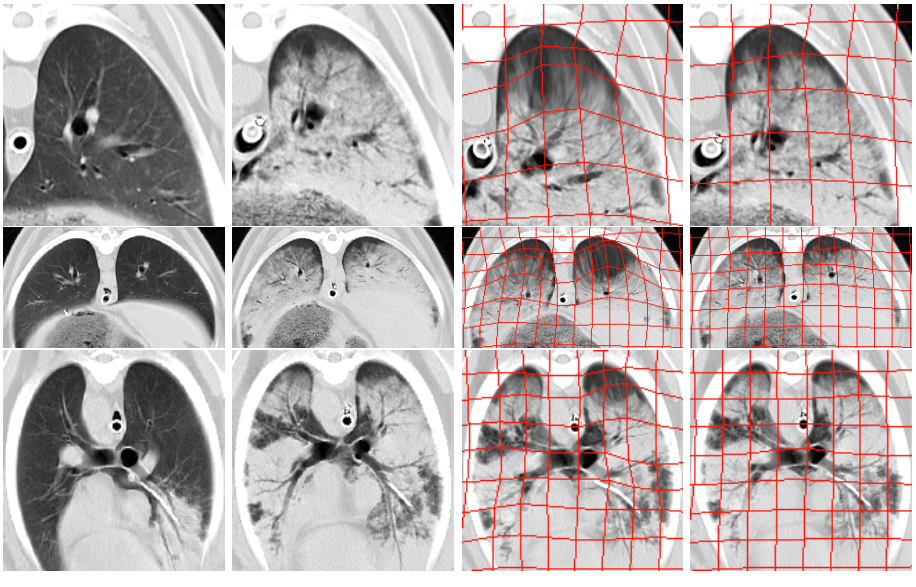

Fig. 5. Axial view of sheep 2 (top) and 3 (middle and bottom). Images are shown with (middle left) and without (left) ARDS before registration. Also shown are the single cost solution (middle right) and our cost switching solution (right) with an overlay of the deformation grid. Notice that the single cost solution expands the dark region at the top of the lung, in the process moving vessel structures away from their corresponding structures. Our approach maintains the character of the rigid transformation within the lung.

rigid approach, where as the single cost approach degraded 49% in comparison to the rigid approach. In all but one case, the cost switching approach maintained the error between landmarks obtained by rigidly registering the scans. In two

cases (the most homogeneous of the ARDS scans), the single cost non-rigid approach also maintained the same level of error but in four cases it significantly increased the error, separating the landmarks even more than if the scans had not been registered in three of those cases.

Visual evaluation shows that our cost switching method more accurately registers the ribs than the single cost non-rigid approach (Fig. 3). More significant, however, are the apparent anatomical changes of the trachea, esophagus, and lungs. In Fig. 4 the principal difference between the methods is in the apparent shape of the trachea. The single cost approach greatly increased the size of the trachea. As the only large air filled region in that portion of the body, the trachea was expanded to maximize the correlation with the air filled healthy lungs. Our cost switching approach thus stayed closer to a physically viable transformation. Note that some large deformations were needed to capture changes to the anatomy external to the lung. The single cost approach expanded air filled regions, most notably at the top of the lungs (Fig. 5). Our approach maintained an anatomically appropriate transformations.

Instead of cross-correlation as the image similarity function to guide our registration, other formulations are possible, most notably mutual information. We did test our approach using mutual information and while we do not present in depth results, we note that mutual information drew the anatomy within the lung away from correct correspondence. This performance may indicate that no general purpose similarity function can be defined to represent the density transformation that occurs between a healthy lung and a lung inflicted with ARDS.

4 Discussion and Conclusion

Assigning cost functions at each iteration allows the registration process itself to influence what cost function is used at a particular location. This is similar to several approaches [14,15] for adaptively deciding the impact a particular image point will have in an image registration process. Shen et al. [14] use a hierarchical ordering of "driving voxels" to register images. They propose a similarity measure in a high dimensional space incorporating tissue classification, image intensity, and geometric moment invariants. Using all the voxels in the image, they then find clusters in this space that represent each tissue type. They allow voxels close to the cluster centers to drive registration early in the registration process. They gradually relax the distance they use for determining the driving voxels until all voxel in the image are used. This approach, however, is not appropriate for the application of registering injured lung images in large part because the approach is built upon good segmentation, which has not been reliably demonstrated for injured lungs. Furthermore, Shen et al. eventually apply a single similarity function for all voxels in the images. In the application of injured lung registration, correlation draws some image regions away from correct correspondence, therefore we explicitly deny regions with low correlation from driving the registration process. Moreover, our framework allows for the use of multiple similarity functions, tailored to the requirements of a particular

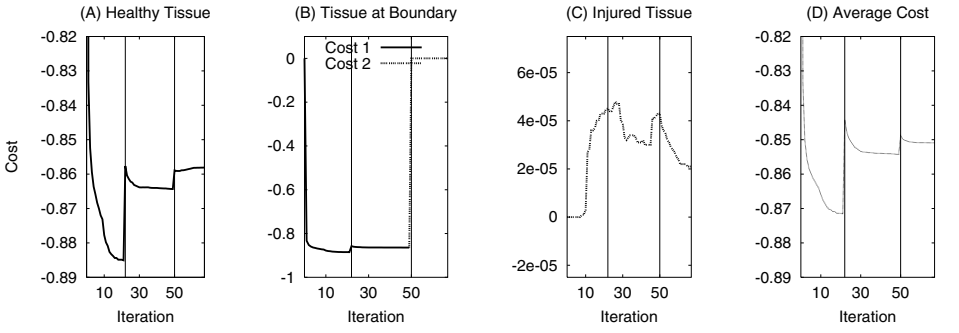

Fig. 6. Typical examples of gradient descent during the registration process at one grid point θ_i for healthy tissue (A), at boundary of healthy and ARDS inflicted tissue (B), within ARDS inflicted tissue (C), and averaged across all transformation grid locations θ (D). Changes in resolution level l occurred at iteration 22 and 50 and are indicated with vertical lines. (A) Outside the lung and within healthy regions of the lung, correlation is strong so $Cost_1$ is used. (B) At the boundary between healthy tissue and injured tissue, correlation can dominate the cost during the low resolution stages of the registration. At the highest resolution, however, strong correlation no longer exists so the registration process automatically switches to $Cost_2$ which does not include correlation costs. (C) Within injured regions of the lung, correlation is weak so $Cost_2$ is used throughout the registration. As the regularizing costs are affected by the movement of neighboring points on the transformation grid, the cost initially increases as other image patches move in response to the gradient of $Cost_1$. The cost begins to decrease once $Cost_1$ no longer drives large movements in other regions of the image. (D) On average, the cost always moves in the direction of a minimum except when the resolution level changes.

application, that can simultaneously drive the registration process on many different image regions of arbitrarily complex description.

Guest et al. [15] propose a reliability measure that can be used on any similarity measure and in any method for computing a registration transformation. The "reliability" of a measure at a particular image point is high if the point matches well to a single point or line in the corresponding image and is low if it matches well to a large region. Our framework can incorporate this measure or any other measure of reliability that is appropriate to a particular application.

We show typical examples of how the cost changed during the registration process in Figure 6. Parameters were set to $\alpha = -1$, $\beta = .001$, and $\gamma = .001$ in Eq. 3. Dues to the low weight given to them, the bending and volume preservation terms did not noticeably impact the solution until the later half of the registration process when the correlation gradient is small. Notice that our algorithm minimizes cost with respect to the transformation parameters and not with respect to the cost function used. The algorithm is not allowed to choose a cost function because it has the lowest value, otherwise, transitions from $Cost_1$ to $Cost_2$ would not occur as seen in Fig. 6 B at iteration 50.

We were limited to using six data sets for a proof of concept test of our approach for registering healthy lungs to diseased lungs. The images came from biomedical study into ARDS where animals were deliberately injured [11]. It is not feasible to obtain a large collection of such images. Our approach may have application to the study of other diseases and injuries to the lung as in pneumonia, asthma, or the effects of near drowning.

Ultimately, to solve the problem of registering CT scans containing healthy lungs to CT scans containing injured lungs, it may be necessary to fuse a feature based approach, perhaps using vessel branching points, to an image similarity approach. Segmentation of these structures, however, is still challenging, especially in the presence of an injury or disease like ARDS. Also, the cost functions used for in our current approach are not directly drawn from biomechanics. The elasticity properties of bones, ligaments, and other tissues should replace the bending energy of thin metal plates and volume preservation costs used here. In addition to formalizing such properties as costs, work remains to apply them to the appropriate anatomy within the registration framework.

In summary, we have presented a novel method for incorporating multiple cost functions into a single registration process. This allows for greater specificity and variation in the definitions of anatomical correspondence used in non-rigid registration problems. We have tested our approach on the difficult problem of registering chest CT scans before and after the infliction of ARDS. By providing a method to register healthy and ARDS inflicted lung scans, we hope to help researchers better understand the syndrome and find new treatment options. Our results demonstrate that being able to selectively apply cost functions on appropriate anatomy does increase the accuracy of the resulting registration.

Acknowledgments

We would like to thank Daniel Rueckert of Imperial College London for providing code which we modified to build our system and Jingbin Wang for discussions. Funding was provided by the National Science Foundation (IIS-0308213, IIS-039009, IIS-0093367, P200A01031, and EIA-0202067).

References

1. H. Lester and S. R. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32:129–149, 1999.
2. M. A. Audette, F. P. Ferrie, and T. M. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Med Image Anal*, 4(3):201–217, 2000.
3. Mäkelä Timo, P. Clarysse, O. Sipilä, N. Pauna, Q. C. Pham, T. Katila, and I. E. Magnin. A review of cardiac image registration methods. *IEEE Trans Med Imag*, 21(9):1011–1021, 2002.
4. J. P. W. Pluim, J.B. A. Maintz, and M. A. Viergever. Mutual-information-based registration of medical images: A survey. *IEEE Trans Med Imag*, 22(8):986–1003, August 2003.

5. M. Betke, H. Hong, D. Thomas, C. Prince, and J. P. Ko. Landmark detection in the chest and registration of lung surfaces with an application to nodule registration. *Med Image Anal*, 7(3):265–281, September 2003.
6. V. Boldea, D. Sarrut, and S. Clippe. Lung deformation estimation with non-rigid registration for radiotherapy treatment. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI'03*, pages 770–777. Springer-Verlag, Berlin, 2003.
7. I. Bricault, G. Ferretti, and P. Cinquin. Registration of real and CT-derived virtual bronchoscopic images to assist transbronchial biopsy. *IEEE Trans Med Imag*, 17(5):703–714, 1998.
8. Baojun Li. *The construction of a normative human lung atlas by inter-subject registration and warping of CT images*. PhD thesis, The University of Iowa, 2004.
9. J. N. Yu, F. H. Rahey, H. D. Gage, C. G. Eades B. A. Harkness, C. A. Pelizzari, and J. W. Keyes Jr. Intermodality, retrospective image registration in the thorax. *The Journal of Nuclear Medicine*, 36(12):2333–2338, December 1995.
10. K. F. Udobi, E. Childs, and K. Touijer. Acute respiratory distress syndrome. *American Family Physician*, 67(2):315–322, 2003.
11. C. L. Bellardine, E. P. Ingenito, A. Hoffman, F. Lopez, W. Sandborn, B. Suki, and K. R. Lutchen. Relating heterogeneous mechanics to gas exchange function during mechanical ventilation. *Annls. Biomed. Eng.*, 2005. In press.
12. D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Trans Med Imag*, 18(8):712–721, August 1999.
13. T. Rohlfing, C. R. Maurer, D. A. Bluemke, and M. A. Jacobs. Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Trans Med Imag*, 22(6):730–741, June 2003.
14. D. Shen and C. Davatzikos. Hammer: Hierarchical attribute matching mechanism for elastic registration. *IEEE Trans Med Imag*, 21(11):1421–1439, November 2002.
15. E. Guest, E. Berry, R. A. Baldock, M. Fidrich, and M. A. Smith. Robjust point correspondence applied to two- and three- dimensional image registration. *IEEE Trans Pattern Anal Mach Intell*, 23(2):165–179, February 2001.