3-D Reconstruction of the Human Ribcage Based on Chest X-Ray Images

Christopher Koehler, Thomas Wischgoll, Forouzan Golshani

Abstract

In order to avoid later stages of lung diseases, preventative methods such as regular preemptive screenings are an absolute necessity. For these screenings, costly methods, such as CT scans, are often times cost prohibitive. X-ray imagery is significantly more cost effective. However, it does not provide a three-dimensional view of the patient's torso. Hence, this paper describes a technique that, based on a series of geometric template models and two X-ray images of a subject's chest, creates a patient specific 3-D reconstruction of the ribcage. The 3-D reconstruction helps improve visualization of the chest without needing to use expensive 3-D scanners and can be combined with existing computer-aided disease detection techniques to help with diagnosis.

Keywords: 3-D modeling, chest radiographs, computer-aided diagnosis (CAD), geometric reconstruction, rib segmentation.

Introduction

Two-dimensional chest radiographs are the most common method for diagnosing pulmonary diseases. They are widely used due to the fact that they are inexpensive and expose the patient to less radiation than 3-D CT scans. However, X-rays are limited because they only generate two-dimensional data. Since an X-ray image is just a projection of the patient's body, it is difficult to visualize the three-dimensional configuration of the organs and bones. From a visualization standpoint, it is useful to have a 3-D reconstruction of the ribs to compliment the X-ray images because they help provide a frame of reference for locating other organs in the chest and for getting a better idea exactly where in the chest a disease such as lung cancer is present and how far it has spread. Figure 1 shows an example of an initial pair of PA and lateral X-ray images that the 3-D reconstruction is based on.

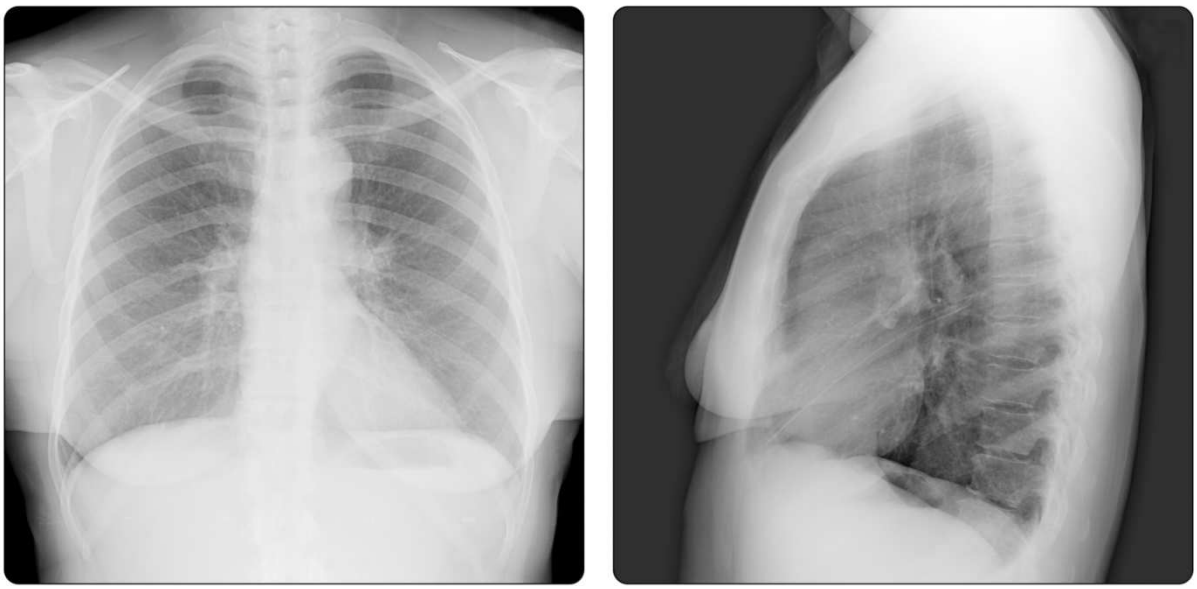

Fig. 1. Initial PA and lateral X-ray images.

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

Generating a 3-D reconstruction of a ribcage based on two 2-D X-ray images is a difficult problem for several reasons. First, most traditional 3-D from 2-D algorithms fail on this kind of data because they were not designed to work with X-ray images and the detailed information therein. Typical 3-D from 2-D algorithms focus on the surfaces of objects in multiple images in order to get clues about the 3-D shapes of the objects. For example, photometric stereo approaches [1] use changes in lighting information to generate 3-D, and shape from texture techniques [2] look at how the size of textures on objects change to generate 3-D. However, in an X-ray image the bones, soft tissue, organs and vascular structure are all superimposed on top of each other, so the surface data necessary for these algorithms to work properly is not available. Volumetric reconstruction techniques like the inverse radon transform [3] also produce unsatisfactory results because they require more X-ray images taken at different angles around the patient in order to generate a reasonably accurate 3-D volume.

Another reason why it is difficult to generate 3-D ribcage reconstructions from X-ray images is the huge diversity found in the shape of human ribs. Simply starting with a complete 3-D ribcage atlas and then linearly transforming it as a whole could never generate every possible alignment of the ribs found in nature.

It is very important that the border of the projection of the 3-D reconstruction matches the edges of the ribs in the PA X-ray image closely or it is likely to cause more confusion than benefit for applications like computer aided disease detection. The overlapped cross hatching pattern of the ribs in an X-ray image can cause a lot of false positives and negatives for automated disease detection software, especially when a diseased area is partially overlapped by one or more ribs. However, when the borders of the ribs are known along with the approximate density of the bone at each pixel, the ribs can be partially removed from the images and the negative affects of the superimposed ribs can be lessened. An additional benefit of the proposed methodology is that the ribcage reconstruction provides a frame in which estimates of disease size and spread can be visualized more effectively than in 2-D X-ray images**.**

Algorithm Overview

To address the limitations 2-D X-ray images have for visualization and disease diagnosis and the limitations of other 3-D from 2-D algorithms when used on X-ray images, this paper introduces a knowledge-based algorithm to generate an approximate 3-D reconstruction of a human ribcage by combining geometric knowledge about true 3-D ribcages with interactive rib segmentation techniques. Measurements taken from the X-ray images are combined with knowledge about the general shape of individual ribs in order to create a good starting 3-D geometric template for each rib so that it will require a minimal amount of data gathered from the X-ray images to transform them to the shape of the subject's true ribs. The vertices in these starting rib surface meshes are grouped into clusters, which are transformed in a way that minimizes 3-D distortion when aligning them with 2-D data.

The modality of the original data used by our algorithm is two 2-D chest X-ray images. One of the X-ray images is from the side (lateral) and the other is from the rear (PA). Without loss of generality, only the steps for the left side of the ribcage are described because it is straightforward to see how the algorithm applies to the right side. Figure 2 illustrates the main steps of the algorithm as well as its applications and related tasks. The process of generating a 3-D rib reconstruction is divided into three main tasks.

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

- Ribs in the PA X-ray images are segmented and classified, and the outer border of the ribcage in the lateral X-ray image is interactively identified.
- Geometric rib template models are generated and their vertices are grouped into several clusters and sub-clusters corresponding to the different parts of a single rib.
- Rib template vertices are transformed in a two-phase cluster-by-cluster process to line up with the segmented edges.

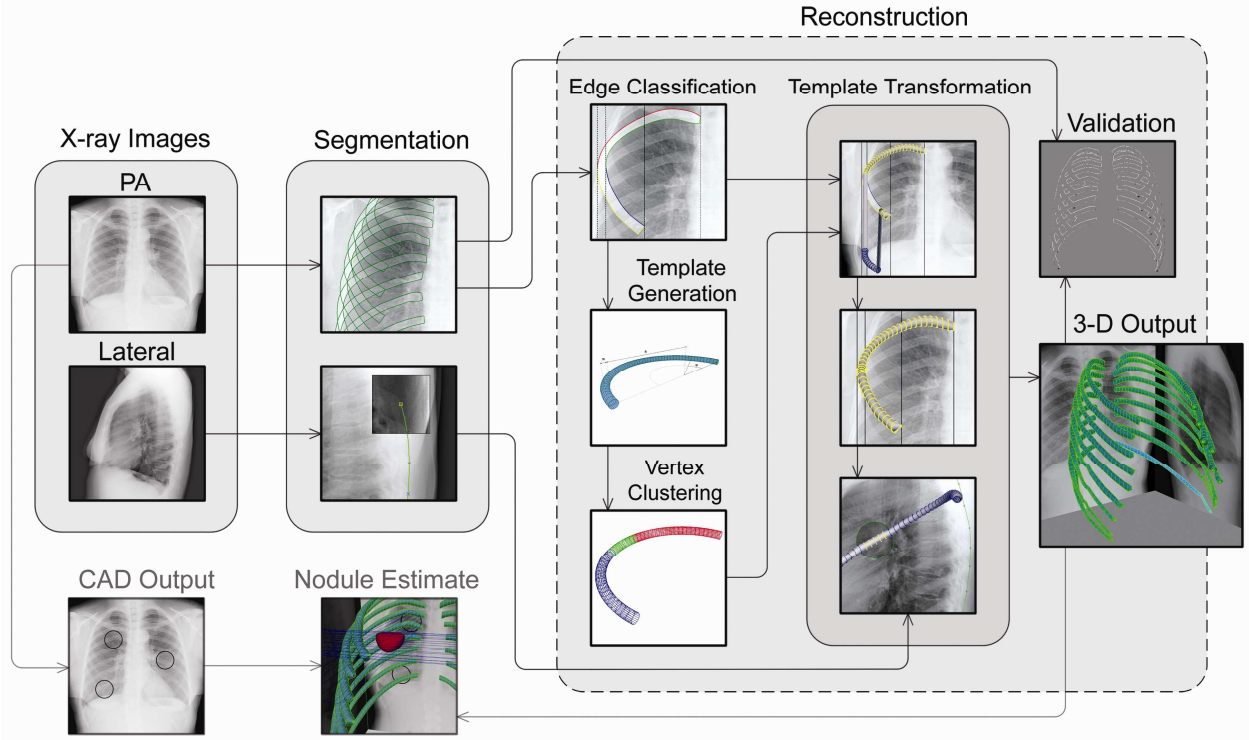

Fig. 2. Overview of the reconstruction algorithm and related processes.

Rib Segmentation

The 3-D ribcage reconstruction algorithm starts with a segmentation of the ribs in the PA X-ray image and a segmentation of the outer ribcage border in the lateral X-ray image. Early attempts to segment the ribs in X-ray images have involved fitting parabolas to the ribs [4]. Vertical derivative sections through the lung fields have been used to extract points on the edges of ribs [5]. Sinusoidal functions have also been fitted to vertical sections of X-ray images to detect ribs [6][7]. R-Snakes have been used to get a more precise segmentation of rib borders after finding initial approximations of the borders with a modified Hough transform [8]. Another approach has been to fit a statistical model of a ribcage to patient specific data in order to segment the rib borders [9]. Pixel-based iterative classification has been used [10] to segment the posterior ribs. Segmentation of scoliotic rib borders has also been achieved with a semi-automated technique using directional filters and edge following [11]. Due to the fact that it is

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

the only published algorithm capable of segmenting the anterior ribs in a PA X-ray image, the technique described in [11] was used while testing the 3-D reconstruction algorithm.

There is currently no reliable way to automatically segment the individual ribs or even the outer border of the ribcage in lateral X-ray images due to the fact that there is a lot of noise in the lateral X-ray images and the ribs from both sides of the body are superimposed on top of each other. For this reason, an interactive approach was used to identify the outer border of the ribcage in the lateral X-ray images. A technician identifies the points at the front and back of each rib. A b-spline curve is fit to the selected points in order to form an approximate boundary around the outside of the ribcage. Histogram equalization is used in a local window surrounding the technician's mouse cursor in order to clarify image features in areas where intensity change is subtle. Figure 3 gives examples of the X-ray images during and after the segmentation process.

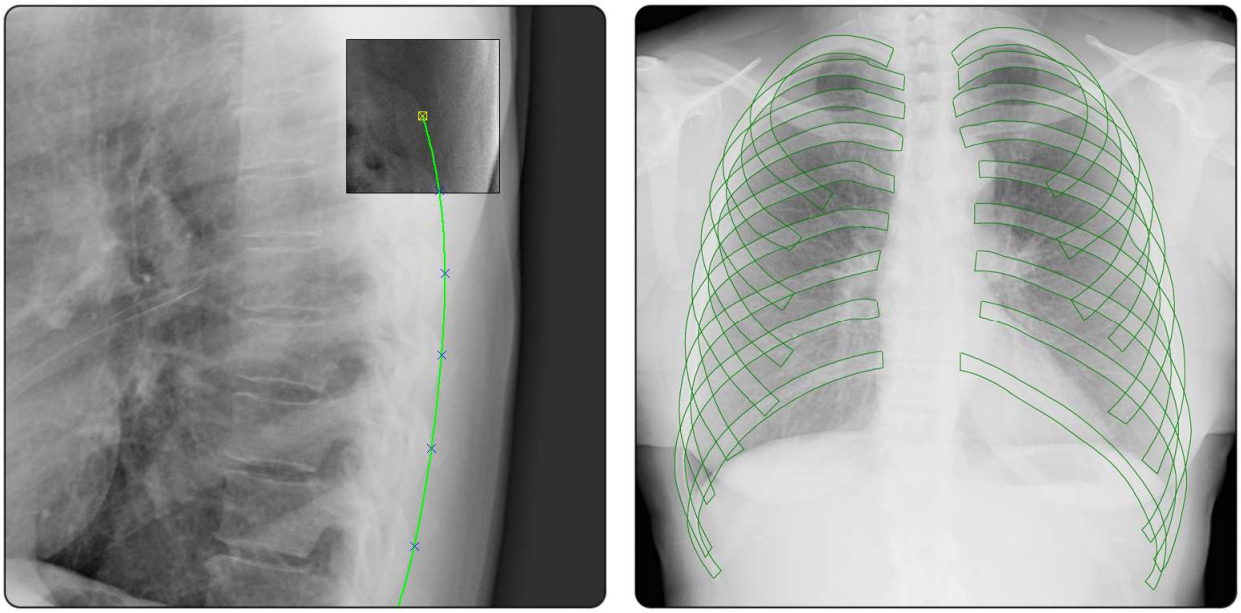

a. b.

Fig. 3. (a.) Interactive segmentation of the ribcage outer border. (b.) Segmented ribs in a PA X-ray image.

Locating Anchor Points

The segmentation step identifies the outline of each rib, however further classification of these edges is needed before inputting them into the rib template deformation algorithm. In particular the vertical lines corresponding to the maximum and minimum x-coordinates of the lateral rib are marked as well as the minimum x-coordinates of the posterior and anterior ribs. Also, the top and bottom edges of the posterior and anterior portions of the ribs are marked with different intensities.

Several quantities, which will be used to generate each rib template, are also defined. The sectional radius Sr is defined as half the difference between the lateral maximum and lateral minimum x-coordinates, the rib width W is defined as the difference between the lateral maximum and posterior minimum x-coordinates, and the rib radius R is defined as the rib width minus the sectional radius $W - Sr$. Figure 4 shows the 5^{th} rib on the left side of the ribcage after

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

the segmented edges have been further classified to contain all the necessary information to generate and transform a 3-D rib template.

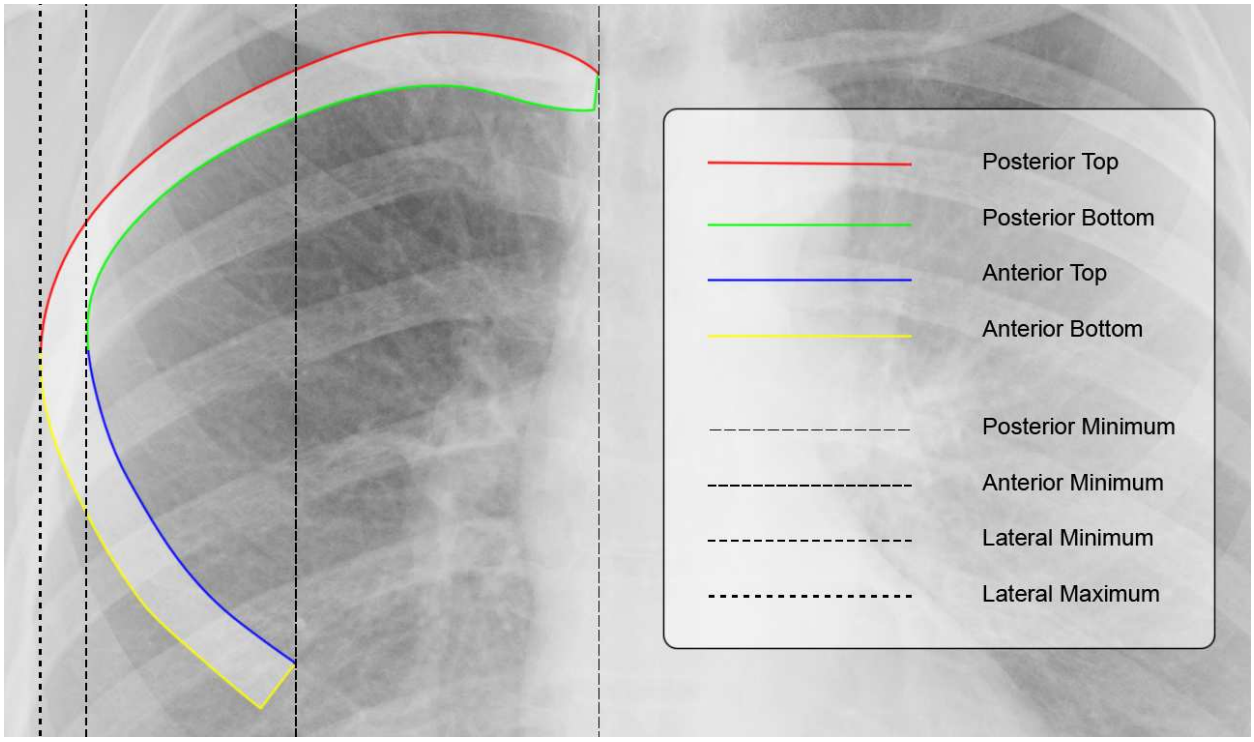

Fig. 4. Classified edges and anchor points.

3-D Template Generation

Since the approximate shape of a single rib is known, a 3-D half-torus mesh can be generated for each rib as a starting point. The radius and sectional radius of each newly generated rib template are taken to be the R and Sr values measured from the corresponding segmented rib. In order to ensure enough resolution in the rib template, there must be one or more axial subdivisions for every pixel in the x-direction of the corresponding segmented rib. In order to more clearly illustrate the steps of the algorithm, a rib template with only 50 axial subdivisions is used for the remainder of the images. Figure 5 shows an example of this initial rib template. However, to ensure the resolution is high enough in practice, the number of axial subdivisions Q in each rib template is defined as follows:

$$Q = \frac{\pi}{(\sin^{-1}(1/Sr))} - 1$$

Starting with a shape that is similar to a true 3-D rib helps deal with the fact that there is not enough data in the X-ray images alone to know what their original 3-D shapes would have been. Linearly transforming the template as a whole will not give enough flexibility to align it to the segmented rib edges, so it still needs to be broken into separate vertex clusters which can be individually aligned to the segmented edges.

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

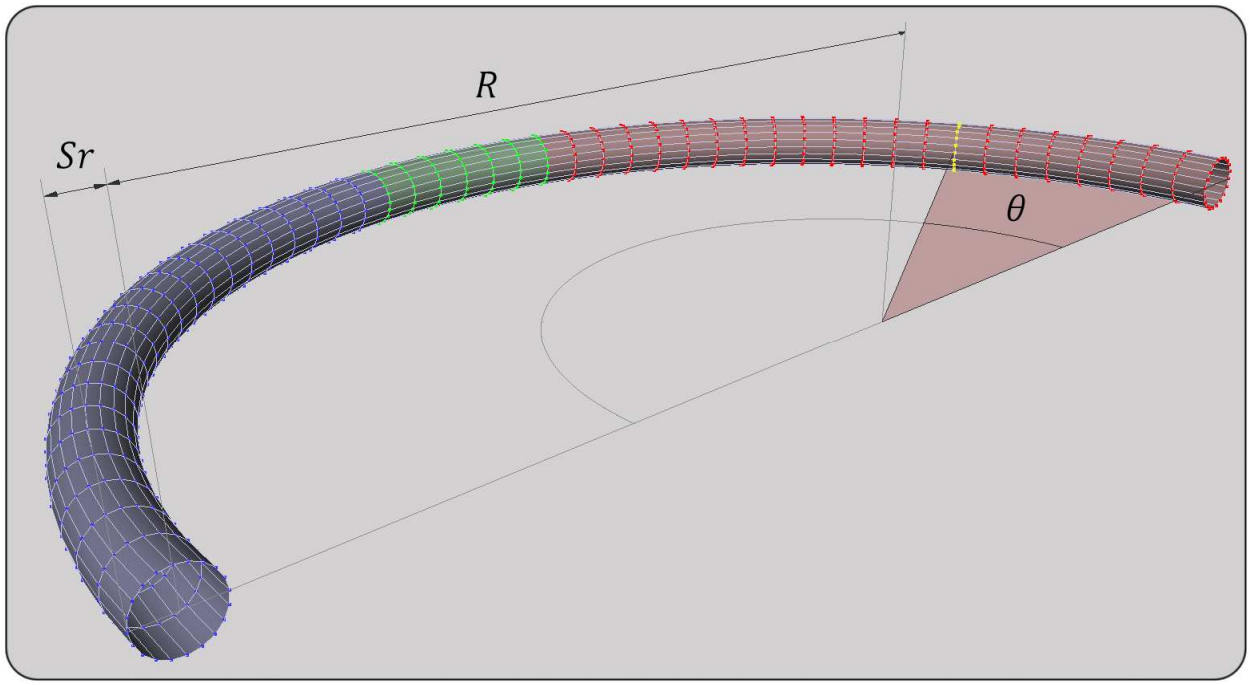

Fig. 5. Initial geometric rib template with classified vertex clusters.

Vertex Clustering

Before performing any transformations, the vertices in the rib template meshes have to be classified in a similar way that the edges in the segmentation output were classified. The vertices are divided into clusters and sub-clusters based on the portions of the rib template that need to be deformed simultaneously. The algorithm starts by defining several groups of vertex sub-clusters, which will later be classified as belonging to the anterior, lateral or posterior portion of the rib template. The vertex sub-clusters are created by defining the point between the two ends of the half-torus as the pole in a polar coordinate system, and grouping all vertices that have the same angle θ into a sub-cluster. Figure 5 illustrates the process of grouping the vertices into sub-clusters.

Once the sub-clusters have been formed they are classified. To do this the entire rib template is rotated to approximately the same angle that a real rib would be rotated. Since it is unreliable to automatically segment rib edges in lateral X-ray images to determine the true rotation amount an experiment was preformed to determine a default rotation angle. A manual segmentation was done on 80 of the most distinguishable rib edges from 5 randomly selected lateral X-ray images. A least squares regression line was fit to each of the segmented edges and its angle was measured. The average angle of these least squares regression lines was found to be 24.8°. For this reason, a rotation of 24.8° is applied to each rib template. It is important that the template be rotated to a fairly accurate starting angle in order to properly classify the sub-clusters into the lateral, posterior and anterior vertex clusters and to reduce distortions in the final output. Properly classifying the sub-clusters assures that they will be aligned to the correct segmented edge.

To classify the sub-clusters, the algorithm starts by finding the vertices in each sub-cluster with the maximum and minimum y-coordinate values. Sub-clusters whose maximum and minimum y-coordinates are not within an experimentally determined threshold value of the

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

border of the 3-D model when it is projected onto the xy-plane are assigned to the lateral vertex cluster. The remaining sub-clusters are assigned to the posterior and anterior vertex clusters based on whether their maximum y-coordinate values are above or below the maximum y-coordinate value of the entire lateral vertex cluster. Figure 5 shows the properly classified anterior (red), lateral (green), and posterior (blue) groups of sub-clusters. Again, this is a simple example template, but it works the same way with a more detailed rib template.

Vertex Cluster Transformation

Once the vertex clustering is complete, it is possible to transform the vertex clusters so that they will line up with the segmentation data that was taken from the original X-ray images. Since our templates are shaped approximately like real ribs and their thickness was measured from the rib segmentation, it is sufficient to only align them to the segmented ribs in the PA X-ray image and then to scale them to fit the ribcage borders in the lateral X-ray image. In our example the segmented edge data from the PA X-ray image is placed on the xy-plane with the origin at the bottom right corner and the segmented lateral X-ray image is placed on the yz-plane with the origin at the bottom left corner.

Due to the fact that it simplifies the methodology if the vertex clusters are always scaled up, the algorithm starts by translating the entire template down so that the projection of its vertex with the maximum y value is below the pixel on the corresponding segmented rib with the minimum y value. Also, before starting to transform the template, vertex sub-clusters should be removed from the end of the anterior vertex cluster until the ratio of the distances between the lateral minimum and the posterior minimum anchor points and the lateral minimum and anterior minimum anchor points is as close as possible to the ratio of distances between the maximum and minimum x-coordinates of the posterior and anterior vertex clusters.

After the template has been aligned along the x-axis, the sub-clusters of vertices are aligned one by one along the y-axis. This step gives the final 3-D model a near perfect alignment with the segmented ribs in the PA X-ray image. Each sub-cluster in the posterior cluster is scaled and transformed so that its vertices with the maximum and minimum y-coordinate values are aligned with the y-coordinate values of the segmented edges corresponding to the top and bottom of the posterior rib. The anterior sub-clusters are transformed in a similar way. Due to the fact that the maximum and minimum vertices of the sub-clusters in the lateral cluster do not necessarily lie on the border of the 3-D model when it is projected onto the xy-plane, the entire lateral group is translated along the y-axis all at once so that it lines up with the lateral maximum anchor point. This process is illustrated in Figure 6.

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

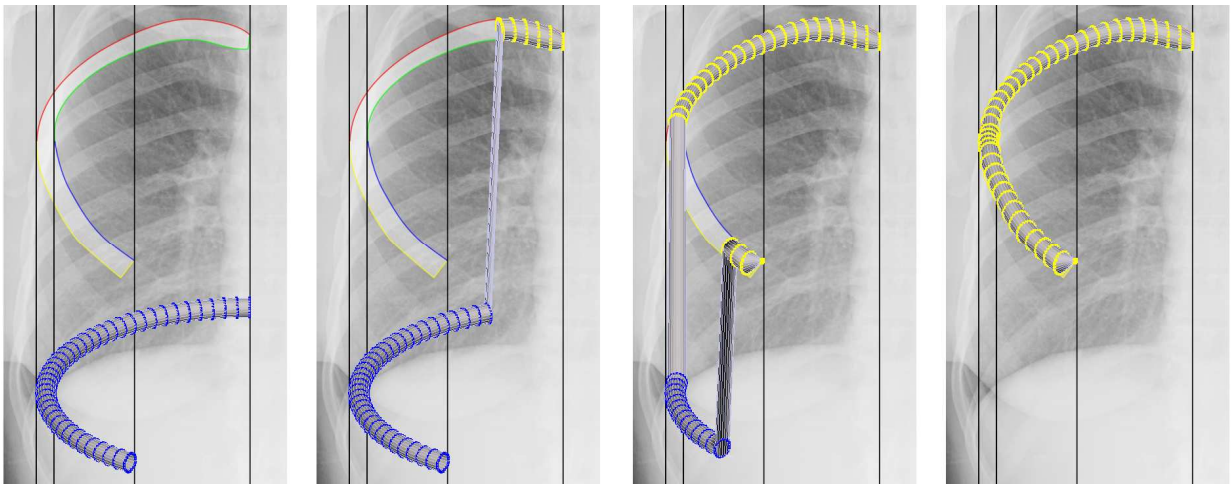

Fig. 6. Transforming a rib template along the y-axis.

At this point the 3-D template is completely aligned with the segmentation data from the PA X-ray image. We have not altered the template along the z-dimension, so to finish the transformation process the entire rib template is scaled and translated along the z-axis so that its vertices with the minimum and maximum z-coordinate values are aligned with the anterior and posterior borders of the ribcage in the lateral X-ray image. The lateral cluster is then rotated around the z-axis so that its angle matches the angle between the two vertices on the template with the maximum and minimum z-coordinate values. This last set of transformations is shown in Figure 7.

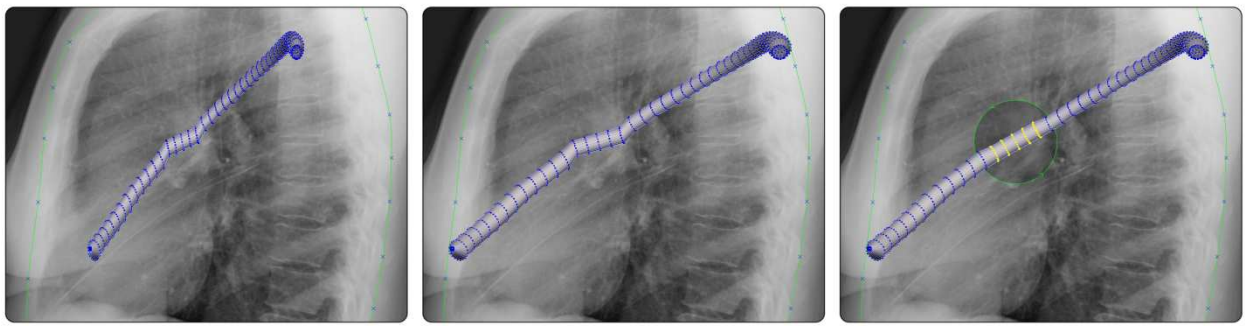

Fig. 7. Fitting the deformed template to the lateral X-ray image.

In order to generate the final 3-D reconstruction, this vertex cluster transformation process is applied to a new 3-D template mesh for each segmented rib. Figure 8 shows perspective and orthographic views of the final output of the reconstruction process. Ribs 1, 11 and 12 have been omitted due to the fact that they can not be segmented reliably. The reconstructions of ribs 2 through 10 are very good approximations considering that so much of the true 3-D information is lost in the X-ray imaging process.

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

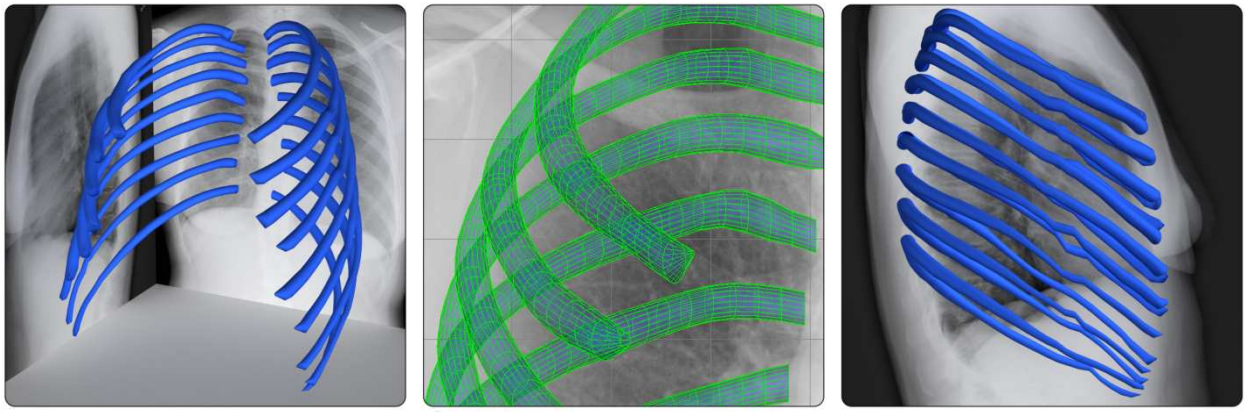

Fig. 8. Final 3-D ribcage reconstruction (a.) Perspective projection (b.) PA projection (c.) Lateral projection

Results and Applications

The goal of the presented methodology is to get a good 3-D approximation, whose projection aligns with the segmented ribs in the PA X-ray image, to use when 3-D scanner output is not available. In order to quantitatively measure the accuracy of the reconstruction compared to the original segmentation, the reconstructed mesh is converted to a volume and then rendered with a maximum intensity projection. The resulting projection of the reconstruction is then compared to the original segmented ribs on a pixel-by-pixel basis. Figure 9 shows the differences between the original segmentation of an X-ray image and the projected silhouette of its 3-D reconstruction. Pixels that were in the projected silhouette but not in the segmentation are marked in white and account for 6.973% of the total silhouette, and pixels that were in the segmentation but not in the silhouette are marked in black and account for 0.537% of the total segmentation in this case. As can be seen in the image, mostly single pixels on the boundary of the rib are detected as different, which is due to the partial pixel-effect within the original X-ray images and rounding during the voxelization process. Overall, this comparison testifies to the good match between the reconstructed ribcage and the X-ray image.

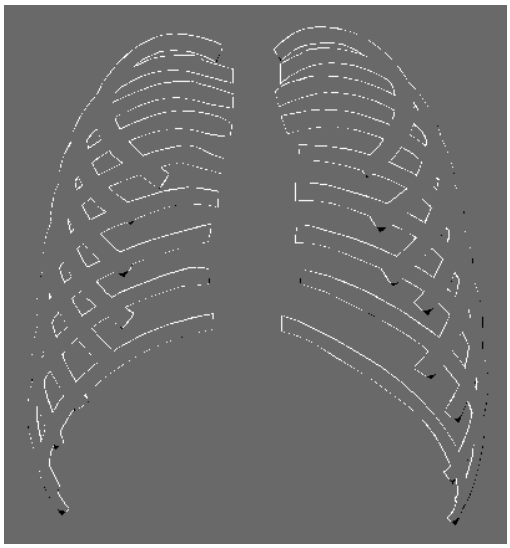

Fig. 9.

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

One of the major uses of the 3-D reconstruction software is as a visualization tool. It is much easier for the human brain to comprehend the real shape of a person's ribcage by seeing it rotating in 3-D instead of looking at two 2-D X-ray images. The 3-D reconstruction software can also be combined with Computer-Aided Diagnosis (CAD) software in applications like early detection of lung cancer, to generate 3-D size estimates and improve visualization of affected regions. For example, if the CAD software is capable of segmenting several potentially affected areas in the PA and lateral X-ray images, the algorithm could determine where the segmented regions from both X-ray images intersect and show a 3-D estimate of the size of the affected region framed by the ribs.

Future Work

The Multimedia based approach of integrating 3-D geometric modeling with medical scanner output to perform reconstructions in situations where other 3-D from 2-D algorithms fail is applicable to other tasks as well. In the future we hope to build on the rib reconstruction algorithm by using its output to reconstruct the lungs. The interior of the 3-D ribcage reconstruction would provide a good starting template to reconstruct the lungs and there are several algorithms available already to segment the lung fields in PA and lateral chest X-rays. The reconstruction algorithm will also be further tested in the future by reconstructing X-rays that were simulated by volume rendering CT scans and then comparing the reconstruction output with a segmentation of the ribs in the original CT scan.

Acknowledgements

The authors would like to thank Riverain Medical for providing the data sets used in this paper as well as the Ohio Department of Development for providing funding for the Early Lung Disease Detection Alliance to pursue the presented research agenda.

References

[1] R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Optical Engineering,* vol. 19, pp. 139-144, 1980.

[2] D. A. Forsyth, "Shape from texture without boundaries," in 2002, pp. 225-239.

[3] P. Toft, "The Radon Transform --- Theory and Implementation," June. 1996.

[4] H. Wechsler, "Image processing algorithms applied to rib boundary detection in chest radiographs," *Computer Vision, Graphics, and Image Processing,* vol. 7, pp. 375-390, 1978.

[5] P. de Souza, "Automatic rib detection in chest radiographs," *Computer Vision, Graphics, and Image Processing,* vol. 23, pp. 129-161, 1983.

[6] G. F. Powell, K. Doi and S. Katsuragawa, "Localization of inter-rib spaces for lung texture analysis and computer-aided diagnosis in digital chest images," *Med. Phys.,* vol. 15, pp. 581-587, 1988.

[7] S. Sanada, K. Doi and H. MacMahon, "Image feature analysis and computer-aided diagnosis in digital radiography: Automated delineation of posterior ribs in chest images," *Med. Phys.,* vol. 18, pp. 964-971, September 1991. 1991.

[8] Zhanjun Yue, A. Goshtasby and L. V. Ackerman, "Automatic detection of rib borders in chest radiographs," *Medical Imaging, IEEE Transactions on,* vol. 14, pp. 525-536, 1995.

[9] B. van Ginneken and B. M. ter Haar Romeny, "Automatic delineation of ribs in frontal chest radiographs," *Proc SPIE Int Soc Opt Eng,* vol. 3979, pp. 825-836, 2000.

[10] M. Loog and B. Ginneken, "Segmentation of the posterior ribs in chest radiographs using iterated contextual pixel classification," *Medical Imaging, IEEE Transactions on,* vol. 25, pp. 602-611, 2006.

[11] F. C. Plourde F. and J. Dansereau, "Semi-Automatic Detection of Scoliotic Rib Borders Using Chest Radiographs," *Studies in Health Technology & Informatics,* vol. 123, pp. 533-537, 2006.

Related Work

The main goal of the reconstruction process described in this paper is to use only the minimum amount of data that is typically collected in early lung cancer detection screenings (PA and lateral X-ray images) and to get a good 3-D approximation, where the 2-D projection aligns perfectly to the segmented rib borders in the PA X-ray. There exists a lot of previous work describing how to create ribcage reconstructions using 3-D CT scans [1] [2] [3], multiple X-ray images at different angles [4], and other attempts at creating approximate reconstructions of anatomical features using only two X-ray images [5] [6].

Delorme et al. [4] evaluate a technique to produce personalized 3-D models of the bone structure of patients with scoliosis. A set of 3-D reference bones from a CT scan of a dry specimen is deformed by free-form deformation based on several manually determined anchor points. They start with one lateral X-ray image and two PA X-ray images taken at different angles. A technician then identifies 10 points on each vertebra, 21 points on the pelvis and 11 points on the midlines of each rib, which are then used to deform the reference bones. The geometry of the ribs specifically is reconstructed with a line reconstruction technique.

Benameur et al. [5] take a model-based approach to generating a 3-D reconstruction of a scoliotic ribcage from a set of PA and lateral radiographs. They align shape samples to a database of 3-D ribcages and then apply probabilistic principle component analysis to the resulting shapes to generate a set of probabilistic prior models that correspond to different deformations seen in scoliotic ribcages. They then use an energy minimization function to select the 3-D reconstruction whose projection is closest to the segmented ribcage. This yields good results when finding the midlines of the 3-D ribs but the reconstructions do not align perfectly with the edges of the ribs in the X-ray images.

Staal et al. [1] presented a method to segment and label the entire ribcage in a volumetric CT scan. They start by detecting 1D ridges in the 3-D data and then create tube shaped primitives from the ridge points. A classifier is trained to determine which portion of the ribs the

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE

primitives correspond to, and each rib is assigned a centerline. Finally, a region-growing algorithm starts at the primitive centerlines to produce the final segmentation of the ribs.

Lamecker et al. [6] developed an atlas-based method to reconstruct a 3-D pelvis from two X-ray images in order to aid surgery planning. They generate a 3D-statistical shape model of a pelvis from a set of training shapes, and then optimize a similarity measure between the projected silhouettes of the shape model and an X-ray image.

Klinder et al. [2] used a model-based method for automatic segmentation and labeling of the ribs and vertebral column in CT scans. They generated an average ribcage from 29 CT data sets. Rib centerlines are then extracted from test data sets and registered with the centerlines from the average model using an iterative closest point algorithm. Finally, vertebras are positioned with an iterative object relation model.

Shen et al. [3] presented an automatic tracing-based algorithm for extracting the ribs from CT volume data sets. Their algorithm uses edge detection and intensity information. They assume a smooth rib surface with perpendicular normals and combine this information with edges to determine the tracing direction. The advantage of this approach is that it finds both centerlines and contours.

References

[1] J. Staal, B. van Ginneken and M. A. Viergever, "Automatic rib segmentation and labeling in computed tomography scans using a general framework for detection, recognition and segmentation of objects in volumetric data," *Medical Image Analysis,* vol. 11, pp. 35-46, 2. 2007.

[2] T. Klinder, C. Lorenz, J. von Berg, S. P. M. Dries, T. Bulow and J. Ostermann, "Automated model-based rib cage segmentation and labeling in CT images," in 2007, pp. 195-202.

[3] H. Shen, L. Liang, M. Shao and S. Qing, "Tracing based segmentation for the labeling of individual rib structures in chest CT volume data," in 2004, pp. 967-974.

[4] S. Delorme, Y. Petit, J. A. de Guise, H. Labelle, C. -. Aubin and J. Dansereau, "Assessment of the 3-D reconstruction and high-resolution geometrical modeling of the human skeletal trunk from 2-D radiographic images," *Biomedical Engineering, IEEE Transactions on,* vol. 50, pp. 989-998, 2003.

[5] S. Benameur, M. Mignotte, F. Destrempes and J. A. De Guise, "Three-dimensional biplanar reconstruction of scoliotic rib cage using the estimation of a mixture of probabilistic prior models," *Biomedical Engineering, IEEE Transactions on,* vol. 52, pp. 1713-1728, 2005.

[6] H. Lamecker, T. H. Wenckebach and H. Hege, "Atlas-based 3D-Shape Reconstruction from X-Ray Images," *Pattern Recognition, International Conference on,* vol. 1, pp. 371-374, 2006.

Digital Object Identifier 10.1109/MMUL.2009.57 1070-986X/$26.00 © 2009 IEEE