Knowledge-Assisted Reconstruction of the Human Rib Cage and Lungs

Christopher Koehler and Thomas Wischgoll ■ *Wright State University*

These knowledge-assisted 3D rib cage and lung reconstruction algorithms facilitate early detection of diseases at a fraction of the cost. Combining shared domain knowledge of human anatomy and solid-modeling techniques with knowledge extracted from x-ray images, this interactive approach transforms a series of primitive template meshes into reconstructed ribs and lungs.

Computer-aided disease detection software effectively helps medical professionals use PA (posterior-anterior) and lateral (side) x-ray images to detect diseases such as lung cancer at early stages, which can be quite difficult. Typically, such software focuses only on PA x-ray images and uses image-processing and soft-computation techniques to identify potentially diseased areas. There's been little focus on extracting knowledge from the x-rays and using knowledge of human anatomy to generate 3D reconstructions. Such reconstructions could help the detection process by providing a different way of visualizing the x-ray data to better investigate hard-to-diagnose regions.

The medical field has already compiled a great deal of tacit knowledge about early lung cancer detection. For example, experienced oncologists can often look at an x-ray image and know something is wrong before actually spotting the cancer nodule. This suggests that a novel approach to computer-aided disease detection that uses 3D reconstruction to give experts a new way to visualize medical data could be just as viable as attempting to automatically detect diseases with a certain error rate.

Using x-ray images that are already collected anyway and generating an approximate 3D reconstruction of the ribs and lungs could provide some of the benefits of a CT (computed tomography) scan at a fraction of the cost. Unfortunately, many traditional 3D-from-2D algorithms fail to work with x-ray data because of the angles at which the images are taken, the lack of lighting or surface texture data, and the human ribs' extremely complex shape. On the other hand, solid-model-based reconstructions are tedious, even for experts, in both geometric modeling and human anatomy.

So, we present a unique approach to generating the desired reconstruction. This approach combines

- an expert user's knowledge of modeling techniques,
- knowledge of the general shape of human ribs and lungs, and
- knowledge that segmentation can extract from the x-ray images

to generate and transform a series of primitive 3D geometric templates into approximate patient-specific rib cage and lung reconstructions. Users can then employ these models as they are or can slightly adjust them. (The "3D Reconstructions Based on X-ray Image Data" sidebar discusses other work in this area.)

Image Space Knowledge Extraction

Using shared knowledge to drive visualization applications can be beneficial when working with large and complex data sets.[1] To demonstrate such an application, we create approximate volumetric reconstructions by using the inverse radon transform on x-ray image pairs. We then use the geometric lung reconstructions as a mask to clip

 Published by the IEEE Computer Society

out the volumes' portions corresponding to the lungs, so that users can instantly see those portions without having to identify them manually in the volume. Figure 1 shows an overview of the reconstruction process.

To reconstruct the ribs and lungs, our algorithms align 3D models to 2D segmentation knowledge from the x-ray images and generate 3D meshes by extruding certain 2D features from the images. So, it's important to start with accurate segmentations of the x-rays. In particular, we must segment the lung fields and the rib borders in the PA x-ray image, and segment the rib cage's outer border in the lateral x-ray image. Segmenting anatomical features in 2D x-ray images is difficult, owing to the large genetic variations between patients and the subtle intensity variations from superimposing bones, organs, and soft tissue during the x-ray process. So, we combine automated and interactive segmentation techniques and integrate some of the best features of previously published segmentation work with our own unique algorithms to detect all the necessary features from a pair of x-ray images.

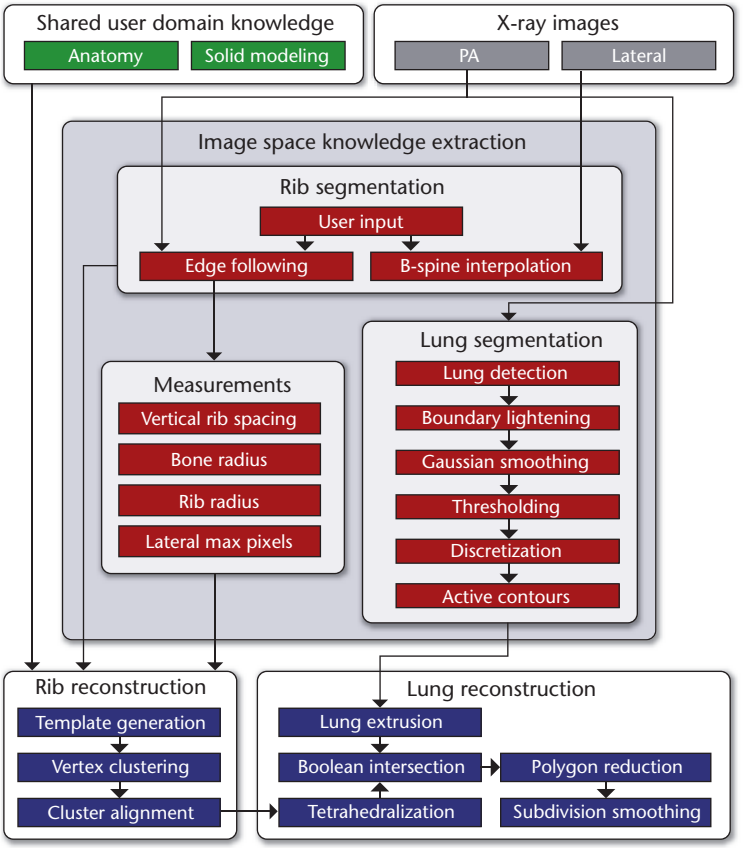

Figure. 1. The pipeline for knowledge-assisted rib and lung reconstruction. The image shows the individual steps for each component of the algorithms to derive the reconstructed rib cage and lung from PA (posterior-anterior) and lateral (side) x-ray images.

Rib Segmentation and Measurements

To generate segmentations to test our rib and lung reconstruction algorithms, we used the technique described by Frédéric Plourde and his colleagues because theirs is the only published algorithm that can segment entire ribs in PA x-ray images.[2] (We summarize related segmentation techniques in the "Rib and Lung Segmentation" sidebar.) After segmenting the rib borders, we identify several anchor points and measure the following distances. First, we identify the pixels corresponding to each lateral rib's outermost portion and use them to determine the distances between consecutive lateral ribs on the rib cage's left ($l_{i,i+1}$) and right ($r_{i,i+1}$) sides. Assuming the same number of segmentable ribs on each side, the following equation defines the average distance between the lateral portions of each rib ρ:

$$\rho = \frac{1}{2n}\left(\sum_{i=v}^{w-1} l_{i,i+1} + \sum_{i=v}^{w-1} r_{i,i+1}\right),$$

where n is the number of segmentable pairs of ribs, v is the number of the first segmentable rib, and w is the number of the last segmentable rib. (To make our algorithms work with x-ray images of variable resolutions where the actual rib and lung data accounts for variable amounts of the entire image, we use ρ as the basis for many other quantities, as we explain in more detail later.)

Also, the radius of a cross-section of each rib's bone on the left side Sr_i is defined as half the distance between the lateral rib's outermost portion and the last pixel directly to its right that's still in the same segmented rib. We define the approximate radius of each rib R_i as the distance between the innermost pixel of the rib's posterior portion and the outermost pixel of the rib's lateral portion minus Sr_i. Figure 2 explains these measurements further.

The 3D reconstruction also requires some knowledge about the rib structure from the lateral x-ray image. Lateral x-ray images can be extremely difficult to interpret because the patient's ribs on both sides and the vertebrae are superimposed atop one another. Thus, owing to the high amount of intuition necessary for this task, a completely automated segmentation isn't feasible. However, because technicians know the rib's general shape beforehand, they need to interactively identify only one point from each rib along the rib cage's rear border and one point on the lateral x-ray image's front border. We use a b-spline curve to interpolate between the marked points, forming an approximate segmentation of the rib cage boundary. A local window of histogram equalization and contrast adjustment follows the user's

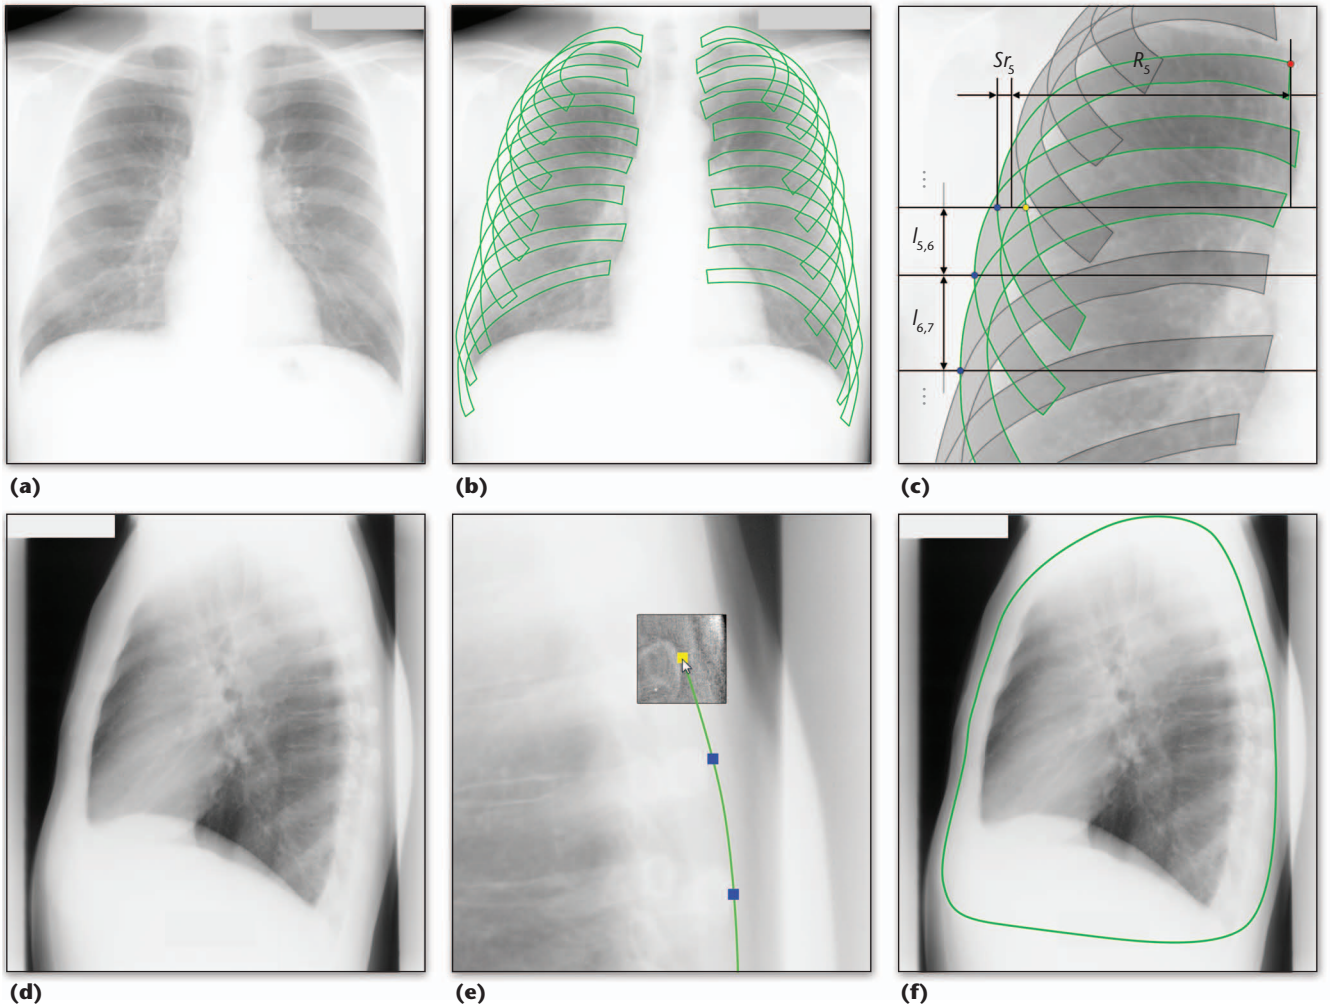

Figure. 2. The rib segmentation and results: (a) the initial PA x-ray; (b) a segmentation based on Frédéric Plourde's edge-following method; (c) inter-rib spacing between ribs 5 and 6 ($l_{5,6}$) and ribs 6 and 7 ($l_{6,7}$), bone radius (Sr_5), and rib radius measurements (R_5); (d) the initial lateral (side) x-ray; (e) interactive segmentation of the rib cage boundary; and (f) the final rib cage boundary segmentation. Overall, these images illustrate the segmentation of the ribs from the x-ray images, along with the derived measurements based on the segmentation.

mouse cursor throughout the process to simplify identifying ribs in the image's noisiest portions (see Figure 2). Having a segmentation of where the ribs connect to the vertebrae would improve reconstruction accuracy but would be too tedious for users.

Lung Field Segmentation

Our lung reconstruction algorithm is based on interactively refining the Boolean intersection of a reconstructed rib cage's convex hull and a 3D extrusion of the lung fields detected in the PA x-ray image. The goal is to produce the most accurate 3D lung reconstruction possible. Thus, there are criteria for our lung segmentations beyond just correctly classifying the highest percentage of pixels as being part of (or not part of) a lung. In particular, the segmented lung's edge must be smooth so that when we discretize the lung segmentation into a mesh, the distance between vertices will be uniform.

The lung fields appear as two large dark patches in the center of the images, surrounded by a lighter area due to the lateral ribs, shoulder bones, organs, and soft tissue. The corners of PA x-ray images are often dark, owing to the patients' bodies not extending to the image corners. Our lung segmentation algorithm starts by identifying pixels near the center of both lungs, separates approximate lung field boundaries from any other dark patches in the image, and then attempts to refine the segmentation border. Figure 3 illustrates this process.

We take the average pixel intensity in the PA x-ray image as a starting threshold value; we mark all pixels below this threshold as black, and all pixels above it as white. We consider a horizontal line through the image's middle row of pixels (see Figure 3a) and examine image intensities on that line. The white pixels closest to the image's center (colored green in Figure 3a) are located on the spi-

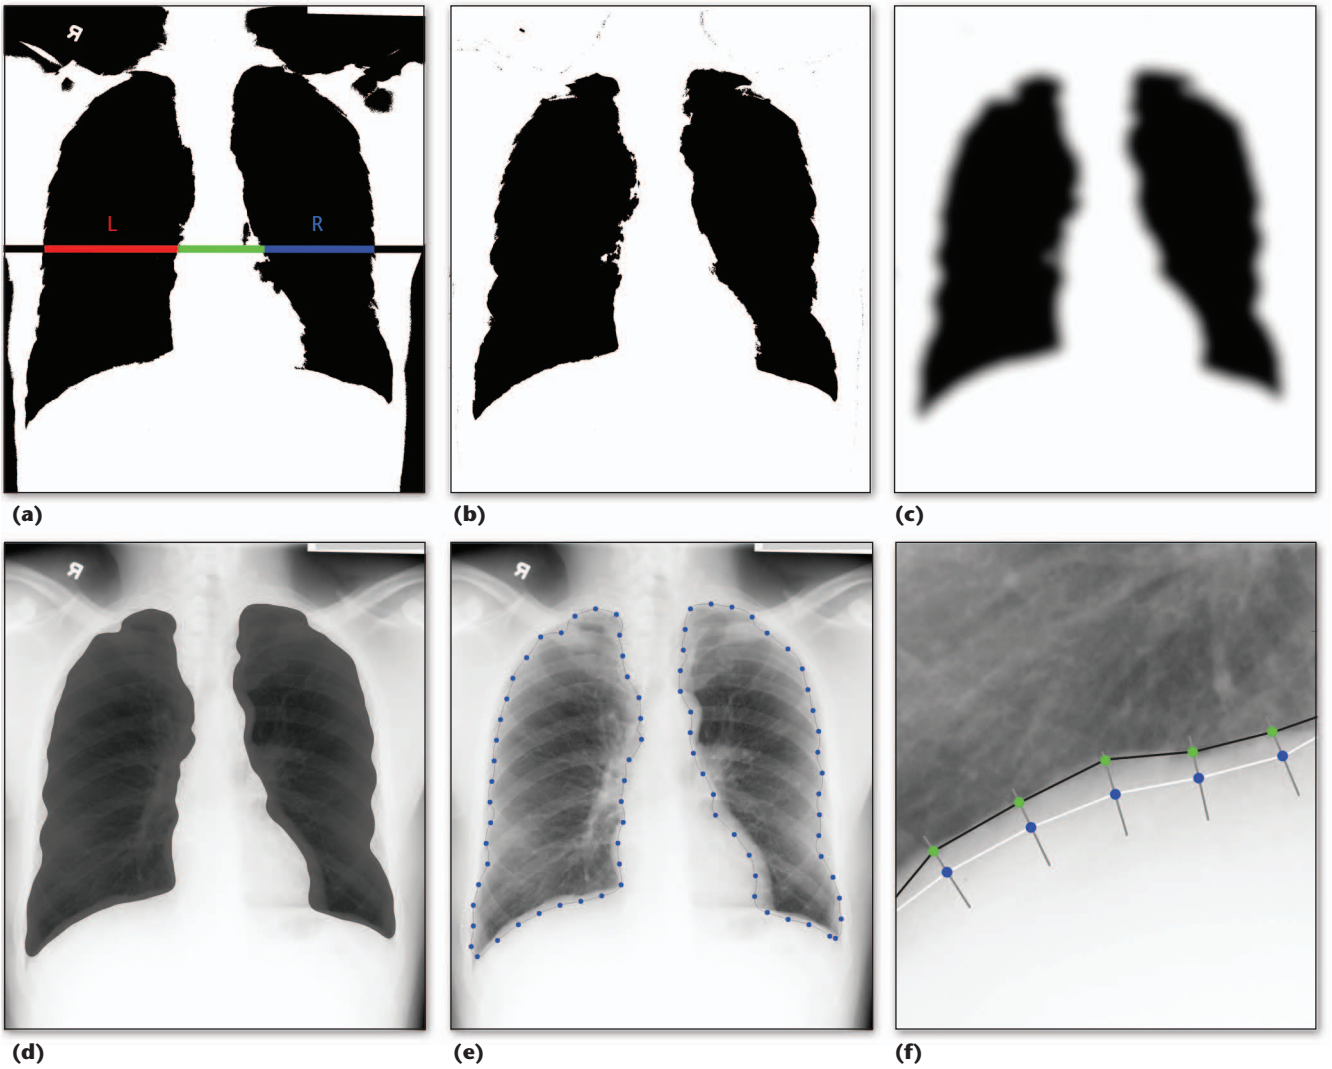

Figure. 3. Lung segmentation: (a) lung patches found after initial thresholding, (b) removal of black around the image border, (c) Gaussian smoothing to remove noise, (d) the initial lung segmentation, (e) the segmentation discretized into a series of vertices, and (f) the lower vertices refined on the basis of the maximum intensity change. The green points represent the new segmentation border. The horizontal line in (a) depicts one scan line, with the segmented areas marked in different colors. All the steps depicted are essential to achieving an accurate segmentation of the lung field in the PA x-ray image.

nal cord, whereas black pixels to the left and right of this white area belong to the lung fields (red and blue, respectively, in Figure 3a). We gradually reduce the threshold value until there are no more paths of black pixels from the lung field centers to the image borders. We then set all black pixels connected to the image borders as white, resulting in an image such as the one in Figure 3b. Gaussian blurring and thresholding smooth the lung borders and remove any remaining noisy spots in the image, resulting in another binary image (see Figure 3c). We then remove any remaining black patches except the two largest (see Figure 3d).

Next, we discretize the lungs into a series of connected vertices, which we refine to more closely match the true lung boundaries. Canny edge detection converts the initial segmented lung fields to contours; the pixel on each contour with the minimum y-coordinate serves as each lung's first vertex. We then recursively trace the contours and place a new vertex every $\rho/2$ pixels, yielding an image such as the one in Figure 3e.

Using a discretization interval of $\rho/2$ that's significantly larger than a pixel's size slightly reduces the extruded mesh's accuracy. However, the initial lung reconstruction resolution is already limited because no available data points exist between the ribs. Producing an initial reconstructed mesh containing many skinny triangular faces is undesirable, so this step intentionally limits resolution.

Next, we refine the discretized lung segmentations' bottom inner portions with an active contour-based technique. The slope of each discretized pixel's tangent serves as the slope of the

3D Reconstructions Based on X-ray Image Data

Researchers have presented several techniques for creating 3D reconstructions using x-ray image data. Hans Lamecker and his colleagues used statistical shape models and atlases to reconstruct a 3D pelvis on the basis of two x-ray images.[1] Researchers have also developed techniques for creating rib cage reconstructions to help treat scoliosis patients. Sébastien Delorme and his colleagues generated patient-specific scoliotic 3D bone models based on free-form deformation, one lateral (side) x-ray image, and two PA (posterior-anterior) x-ray images taken at different angles.[2] Also, Said Benameur and his colleagues performed x-ray image-based reconstructions for scoliosis patients by using an energy minimization function to select from a database of probabilistic prior models.[3]

Finally, although the application is entirely different, the techniques that Xuetao Yin and his colleagues developed to reconstruct buildings in 3D on the basis of architectural drawings also have much in common with reconstructing anatomical structures using x-ray image data.[4]

References

1. H. Lamecker, T.H. Wenckebach, and H.-C. Hege, "Atlas-Based 3D-Shape Reconstruction from X-ray Images," *Proc. 18th Int'l Conf. Pattern Recognition*, vol. 1, IEEE CS Press, 2006, pp. 371–374.
2. S. Delorme et al., "Assessment of the 3-D Reconstruction and High-Resolution Geometrical Modeling of the Human Skeletal Trunk from 2-D Radiographic Images," *IEEE Trans. Biomedical Eng.*, vol. 50, no. 8, 2003, pp. 989–998.
3. S. Benameur et al., "Three-Dimensional Biplanar Reconstruction of Scoliotic Rib Cage Using the Estimation of a Mixture of Probabilistic Prior Models," *IEEE Trans. Biomedical Eng.*, vol. 52, no. 10, 2005, pp. 1713–1728.
4. X. Yin, P. Wonka, and A. Razdan, "Generating 3D Building Models from Architectural Drawings: A Survey," *IEEE Computer Graphics and Applications*, vol. 29, no. 1, 2009, pp. 20–30.

line between the next and previous pixels. We use the slope's negative inverse to find endpoints of a line segment orthogonal to the current pixel; we connect the endpoints using the Bresenham algorithm. Determining the intensities of the original PA x-ray image corresponding to each point on this new line segment gives us an intensity histogram. We perform Gaussian smoothing on this histogram and find the intensities' first derivatives. We assume that the pixel with the maximum rate of intensity change is part of the lung field's new discretized boundary if the intensity change at that pixel is above a certain threshold. Figure 3e shows a close-up of several refined vertices on a discretized lung segmentation. To more closely match expert segmentations, we experimentally determined parameters such as this threshold, the orthogonal line segments' length, and the smoothing filter kernel size.

We tested the lung segmentation algorithm on 10 randomly selected PA x-ray images from 10 different patients for which expert segmentations were also available. We then compared the segmentation results with manual expert segmentations by two reviewers; in each case, we saved the number of true-positive (T_p), true-negative (T_n), false-positive (F_p), and false-negative (F_n) pixels. We measured each segmentation's accuracy using this formula:

$$Accuracy = \frac{T_n + T_p}{T_n + T_p + F_n + F_p}.$$

After the initial segmentation, the average accuracy was 0.9376 ± 0.0176. After we discretized the segmentation, it dropped slightly to 0.9363 ± 0.0175. Then, after we refined the discretization, it increased to 0.9428 ± 0.0172. The expert-segmented lungs used to calculate these accuracy measures occupied an average of 31.09 percent of the test images. Further validation of this technique is necessary, but, for our purposes, the lung reconstruction results are ultimately what determine the segmentation's merit. Figure 4 compares automatic lung segmentations and manually performed expert segmentations from one reviewer.

Knowledge-Assisted 3D Reconstruction

To perform the 3D reconstructions, we use the segmentation knowledge extracted from the x-ray images and combine it with our knowledge about the general shape of human ribs and lungs, and we automate various modeling pipelines to produce the desired shapes. This lets us compensate for the lack of 3D information in the original images. We reconstruct the rib cage first because it's one of the inputs necessary for the lung reconstruction algorithm. We've improved our template-based rib cage reconstruction method by using shared domain knowledge of anatomy and polygon modeling.[3]

Rib Cage Reconstruction

For each segmented rib, we generate a primitive geometric template similar to a real rib, using the knowledge extracted during segmentation. We then split the template's vertices into groups that can be individually aligned to the rib segmentation in the PA x-ray image, such that the entire template eventually lines up with the segmentation without adding distortions along the unmodified dimension. After modifying the templates to align them with the segmented ribs in the PA x-

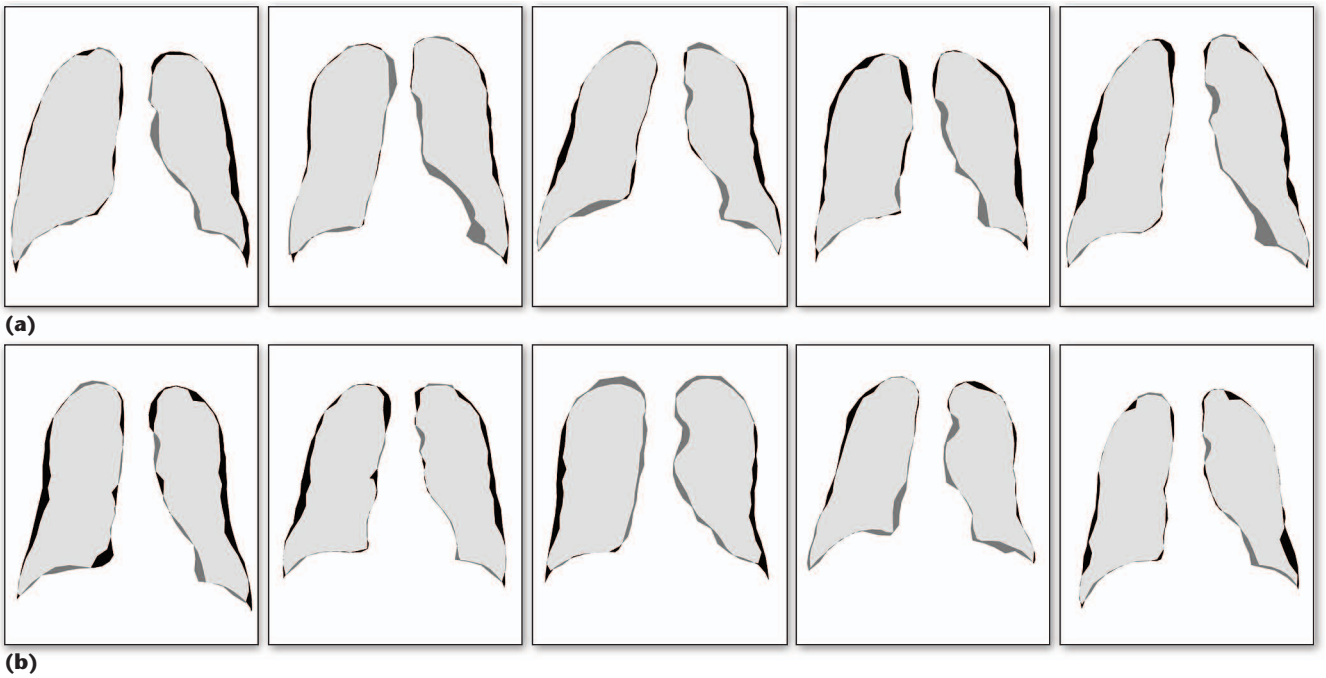

Figure 4. Comparison between automatic lung segmentations and manually performed expert segmentations for 10 different patient data sets. True-positive pixels are light gray, true negatives are white, false positives (areas considered as belonging to the lung by the algorithm but not by the experts) are dark gray, and false negatives (areas considered as not belonging to the lung by the algorithm but as belonging by the experts) are black.

ray, we translate and scale them horizontally so that their maximum and minimum vertices align with the rib cage boundary segmentation from the lateral x-ray image.

Knowledge-assisted primitive selection. A half-torus mesh serves as a geometric template for each rib. We use the measurements of the bone radius and the rib radius from the segmented ribs to parameterize the ith rib template:

$$x_i = [R_i + Sr_i(\cos v)](\cos u)$$

$$y_i = [R_i + Sr_i(\cos v)](\sin u)$$

$$z_i = Sr_i(\sin v)$$

where $v \in [0, 2\pi]$, and $u \in [0, \pi]$ for ribs on the left and $[\pi, 2\pi]$ for ribs on the right. We then sample parameters u and v a discrete number of times to yield a tube-shaped mesh for each template, in which the distance from the tube's center to the template's center is R_i, and the tube's radius is Sr_i.

Although higher resolution obviously yields a more accurate rib reconstruction, that isn't the case when using reconstructed ribs to generate a lung reconstruction. Regardless of how high the rib mesh's resolution is, no vertices will be available between the ribs. Thus, a higher resolution in the ribs will yield a lung mesh with many triangular faces that have a poor aspect ratio. So, when the algorithm has generated the mesh, it samples u with increments of π/ρ, resulting in ρ sections around the template's center.

The increment at which the algorithm samples v is less crucial in terms of generating a good lung reconstruction. For visualization purposes, the sampling resolution should be high enough that the rib surface appears smooth to the eye. This is an approximation, because it assumes cross-sections orthogonal to a rib are circular, which isn't the case. In the future, we hope to use knowledge extracted from the x-rays to modify the algorithm so that it can capture the true rib cross sections' rotating oval shape.

Vertex clustering. After generating the initial geometric templates but before transforming them, we group their vertices into clusters, which, according to the knowledge extracted through segmentation, produce a good rib reconstruction when transformed together. Linearly transforming all the templates couldn't possibly align them to the segmented ribs, and finding a nonlinear transformation to accomplish this would be difficult. But, by grouping the template's vertices, we can align them to the ribs through a series of linear transformations.

We refer to the two kinds of vertex clusters that are formed as *interior vertex clusters* and *radial vertex clusters*. Each rib template has only one

Rib and Lung Segmentation

Our reconstruction algorithms use knowledge extracted from x-ray images through segmentation. Researchers have proposed several techniques to segment rib borders in PA (posterior-anterior) x-ray images. Zhanjun Yue and others found approximate rib borders using the Hough transform and then refined them using R-Snakes.[1] Bram van Ginneken and Bart ter Haar Romeny aligned statistical rib cage models with patient-specific data.[2] Marco Loog and Bram van Ginneken used iterative pixel classification to segment rib borders.[3] Frédéric Plourde and his colleagues performed semiautomatic segmentation of both the anterior and posterior ribs using directional filtering and parallel-edge following.[4]

In the past decade, researchers have also presented several methods to segment lung fields. Osamu Tsujii and his colleagues used adaptive-sized hybrid neural networks to classify whether pixels were part of a lung, on the basis of measurements taken at those pixels.[5] Finally, Bram van Ginneken and Bart ter Haar Romeny effectively used a combination of pixel classification and rule-based segmentation to segment lung fields in x-ray images.[6]

References

1. Z. Yue, A. Goshtasby, and L.V. Ackerman, "Automatic Detection of Rib Borders in Chest Radiographs," *IEEE Trans. Medical Imaging*, vol. 14, no. 3, 1995, pp. 525–536.
2. B. van Ginneken and B.M. ter Haar Romeny, "Automatic Delineation of Ribs in Frontal Chest Radiographs," *Proc. Medical Imaging: Image Processing*, vol. 3979, SPIE, 2000, pp. 825–836.
3. M. Loog and B. van Ginneken, "Segmentation of the Posterior Ribs in Chest Radiographs Using Iterated Contextual Pixel Classification," *IEEE Trans. Medical Imaging*, vol. 25, no. 5, 2006, pp. 602–611.
4. F. Plourde, F. Cheriet, and J. Dansereau, "Semi-automatic Detection of Scoliotic Rib Borders Using Chest Radiographs," *Research into Spinal Deformities 5*, IOS Press, 2006, pp. 533–537.
5. O. Tsujii, M.T. Freedman, and S.K. Mun, "Automated Segmentation of Anatomic Regions in Chest Radiographs Using an Adaptive-Sized Hybrid Neural Network," *Medical Physics*, vol. 25, no. 6, 1998, pp. 998–1007.
6. B. van Ginneken and B.M. ter Haar Romeny, "Automatic Segmentation of Lung Fields in Chest Radiographs," *Medical Physics*, vol. 27, no. 10, 2000, pp. 2445–2455.

interior vertex cluster but multiple radial clusters. The radial vertex clusters are essentially circular groups of coplanar vertices that lie on a plane orthogonal to the rib template. Using the parametric equations for the templates, we can generate radial vertex clusters by holding u constant and incrementing v as usual.

Each template's interior vertex cluster is essentially the set of vertices from each radial cluster that are closest to the template's center of mass. Parametrically, we can generate the interior vertex cluster by setting $v = \pi$ and incrementing u in the normal way. The interior vertex cluster isn't part of the template transformation. It keeps track of the vertices in the reconstructed rib cage throughout reconstruction so that we can export them later to generate the lung reconstruction.

Vertex cluster alignment using image space knowledge. The next step individually transforms the radial vertex clusters to match the thickness of the corresponding portion of the rib segmentation. The resulting transformed template's bone radius at any point will closely match the true rib bone's radius. But first we remove extra radial clusters from the end of the template's anterior portion corresponding to the extent of the segmented anterior ribs. We translate the entire template so that the outermost vertex in its lateral section aligns with the outermost lateral pixel in the rib segmentation. We then translate radial vertex clusters on the template's posterior portion vertically upward until their uppermost vertices align with the top of the corresponding segmented rib edge. Similarly, we translate clusters from the rib's anterior portion downward. Figure 5 illustrates this process.

Next, we further refine the radial vertex clusters. We take the vertex most closely aligned with a segmented edge as a pivot point. We then scale the cluster until one other vertex lies on the opposite segmented edge and its remaining vertices lie between the segmented edges. At this point, the reconstructed ribs' PA projection should align with the segmentation from the PA x-ray image but not the lateral x-ray image. To correct this, we translate each template and scale it horizontally so that its front and rear border vertices lie on the outer rib cage boundary's segmentation from the lateral x-ray image. The resulting structure that forms after we've deformed all the templates in this manner is a good approximation of the true 3D structure of the patient's ribs from the initial x-ray images. Figure 6 illustrates refinement of the rib cage reconstruction.

Lung Reconstruction

The geometric lung reconstruction uses both the rib reconstruction knowledge that we've already generated and the lung segmentation knowledge. Essentially, we create a mesh that encompasses the inside of the reconstructed rib cage. Next, we carve out the portion of this mesh that corresponds to the

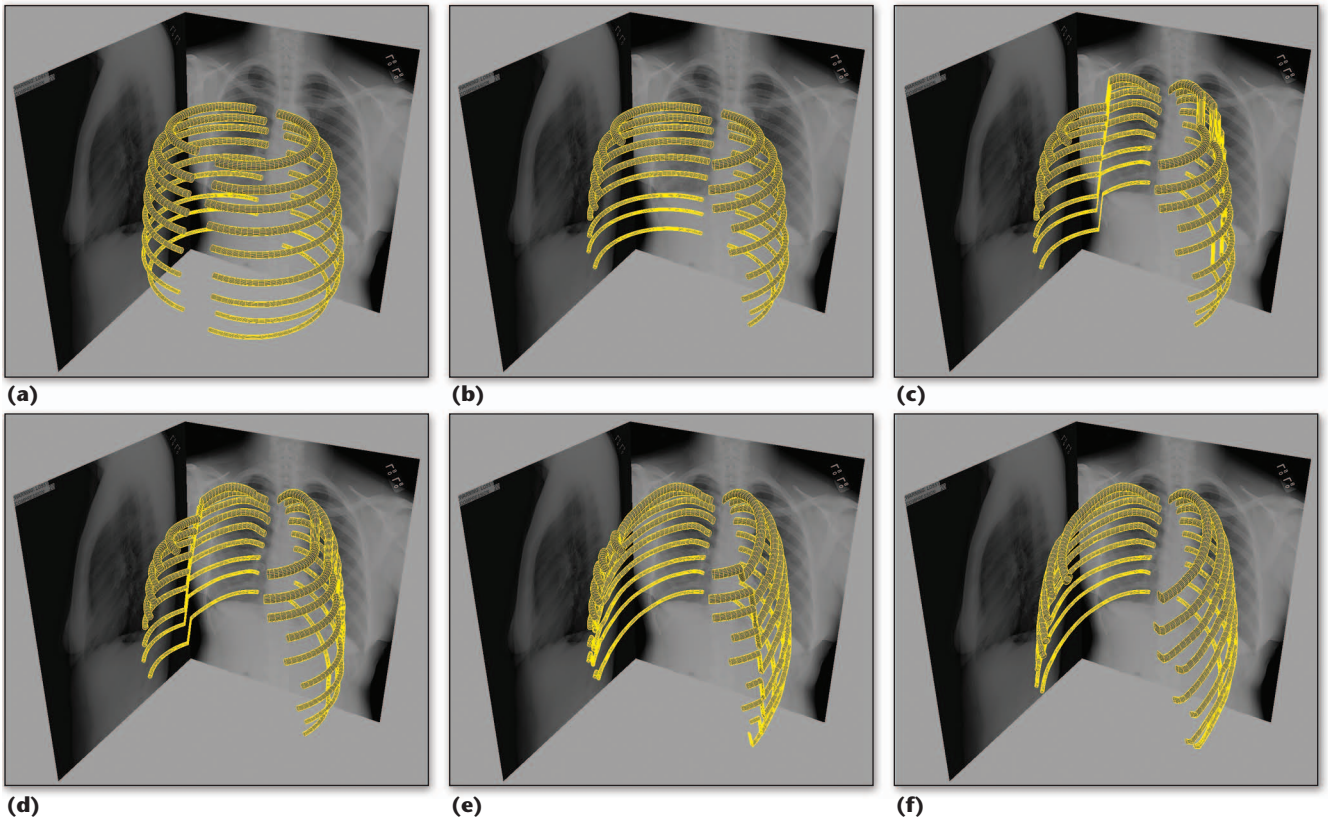

Figure 5. The initial pass of the radial-vertex-cluster transformation: (a) the initial rib templates, (b) rib templates with extra vertices removed from the end of the anterior section, (c–f) radial clusters translated vertically to align with the segmentation at different stages of this process. This shows the process of transforming the rib template to match the segmented ribs from the PA x-ray image.

lungs by extruding the segmented lungs up through it. We then refine the resulting mesh using polygon reduction and subdivision smoothing so that its surface is smooth like that of a real lung.

Segmented lung extrusion. We discretize the lung segmentation boundary into a series of vertices (as discussed earlier) so that we can extrude those vertices into 3D to generate a shape to serve as the basis for reconstructing the lungs. The result is two tube-like structures that resemble lungs from the front but are completely flat along the sides. We choose a discretization interval that's coarser than the segmentation's available accuracy. This prevents the resulting mesh from having many extremely skinny triangular faces when we merge it with another mesh whose accuracy is already limited by a lack of data between ribs.

Rib interior-vertex-cluster tetrahedralization. Real lungs don't have flat edges but curve along with the ribs' interior. The lung fields' extrusion alone doesn't capture this fact, and no information is available directly from the x-rays to generate this structure. So, we use the reconstructed rib cage we already generated. Recall that we grouped several vertices from each rib template into the interior vertex clusters before transforming the rib templates. There's nothing in the rib cage reconstruction process that removes these vertices from each rib's inner border after we've transformed the templates. So, we generate a new mesh that fills the rib cage's interior by computing the Delaunay tetrahedralization of the vertices from each deformed rib template's interior vertex cluster.[4]

Rib interior and lung intersection. The new mesh captures approximately how the lungs' outer border follows the inside of the rib cage, but it's solid in the center and near the bottom where the trachea, diaphragm, vertebrae, and heart normally would be. The silhouette of these features is white in the PA x-ray image, so the lung segmentation doesn't include them. Thus, the next step of the lung reconstruction is to carve the lung fields out by taking the Boolean intersection of the two meshes already obtained. Figure 7 illustrates the lung reconstruction up to this point.

Knowledge-assisted 3D lung boundary refinement. The initial reconstructed lung mesh has several problems that we must correct to improve the

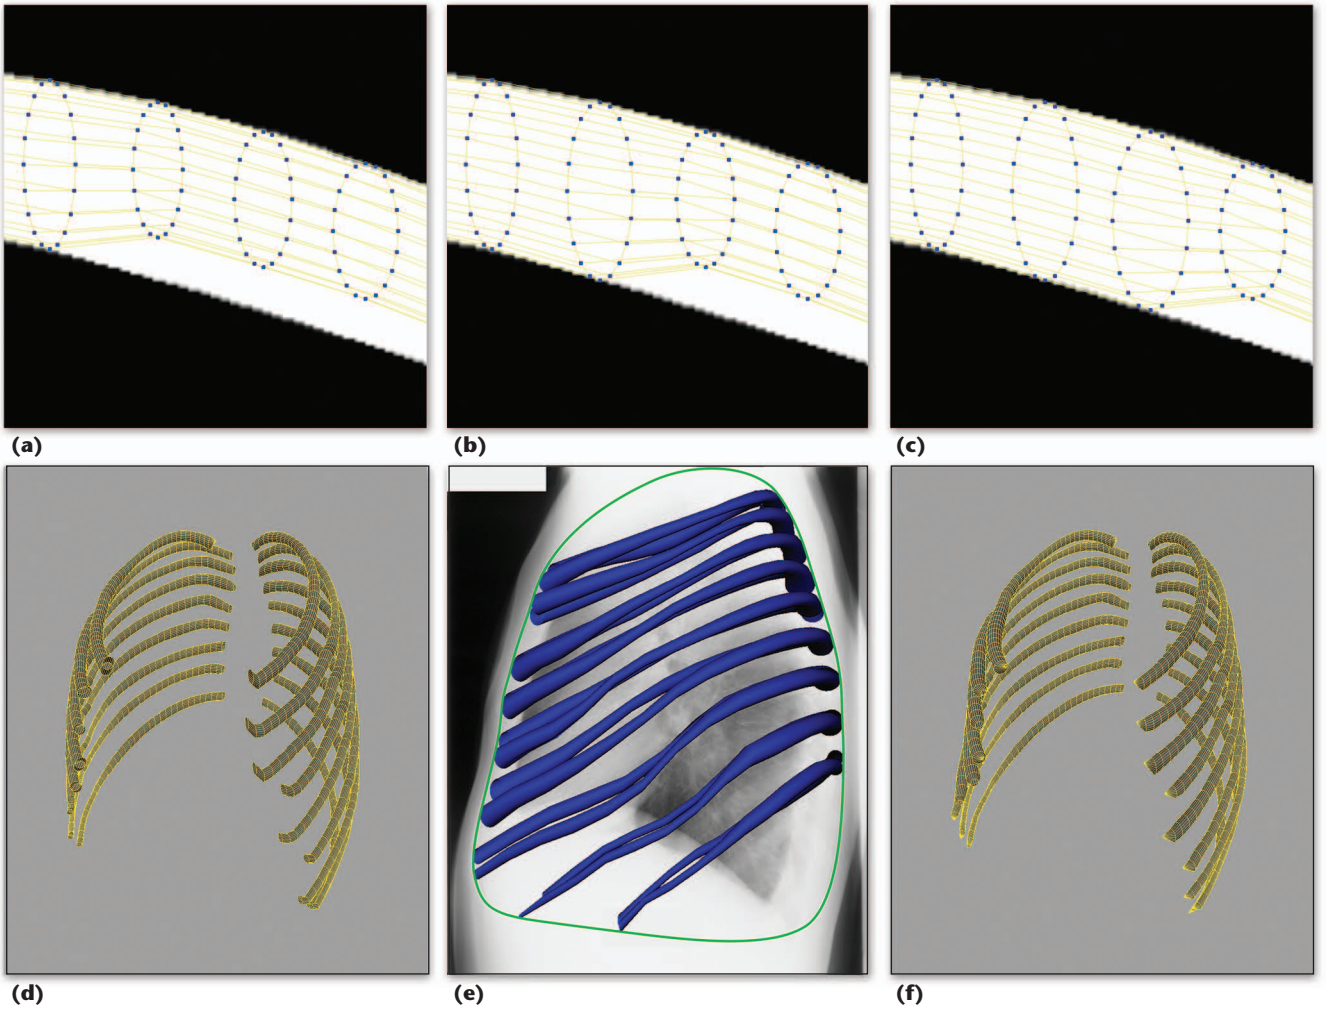

Figure 6. Refinement of a rib reconstruction: (a–c) radial vertex clusters scaled so that their projections align with a segmented rib's two opposing edges at different stages as the algorithm processes every radial cluster sequentially, (d) the reconstructed rib cage after the first pass of vertex cluster transformation, (e) rib templates translated and scaled to align with the segmentation of the lateral rib cage projection, and (f) the final reconstructed rib cage. These are the required steps for aligning the ribs to the lateral x-ray image.v

lung reconstruction's accuracy. In particular, it contains hard edges at the boundaries of where the rib cage interior and lung extrusion meshes intersect. We could correct this problem through a combination of polygon reduction and subdivision smoothing. However, this process isn't entirely straightforward. The initial mesh of the reconstructed lungs will inevitably contain a variable number of strangely shaped polygons, despite our best efforts to limit this outcome at several steps in the reconstruction. Subdivision smoothing alone produces undesirable results because of the poor aspect ratio of some of the mesh's faces. Even if we perform polygon reduction first to make the faces more uniform, artifacts of the reconstruction process will manifest themselves in the final smoothed model for variable percentages of polygon reduction. This makes determining a fixed reduction percentage difficult.

The solution is to use knowledge gathered from previous tests with similar input data to generate a good starting estimate of the reduction percentage, and then let users refine this estimate on the basis of their knowledge of real lungs' actual shape. The users employ a slider to control the polygon reduction percentage while the software system displays the result of polygon reduction and smoothing in real time. Ideally, we'd like to keep as much of the initial structure as possible, because as we remove more polygons, the reconstructed lungs slowly move further from the inside of the reconstructed ribs. On the other hand, geometric oddities will more likely occur at low polygon reduction percentages, so we must find an ideal intermediate percentage for each case.

Figure 8 illustrates this interactive refinement process based on lung reconstructions for two patients with significantly different bone structures.

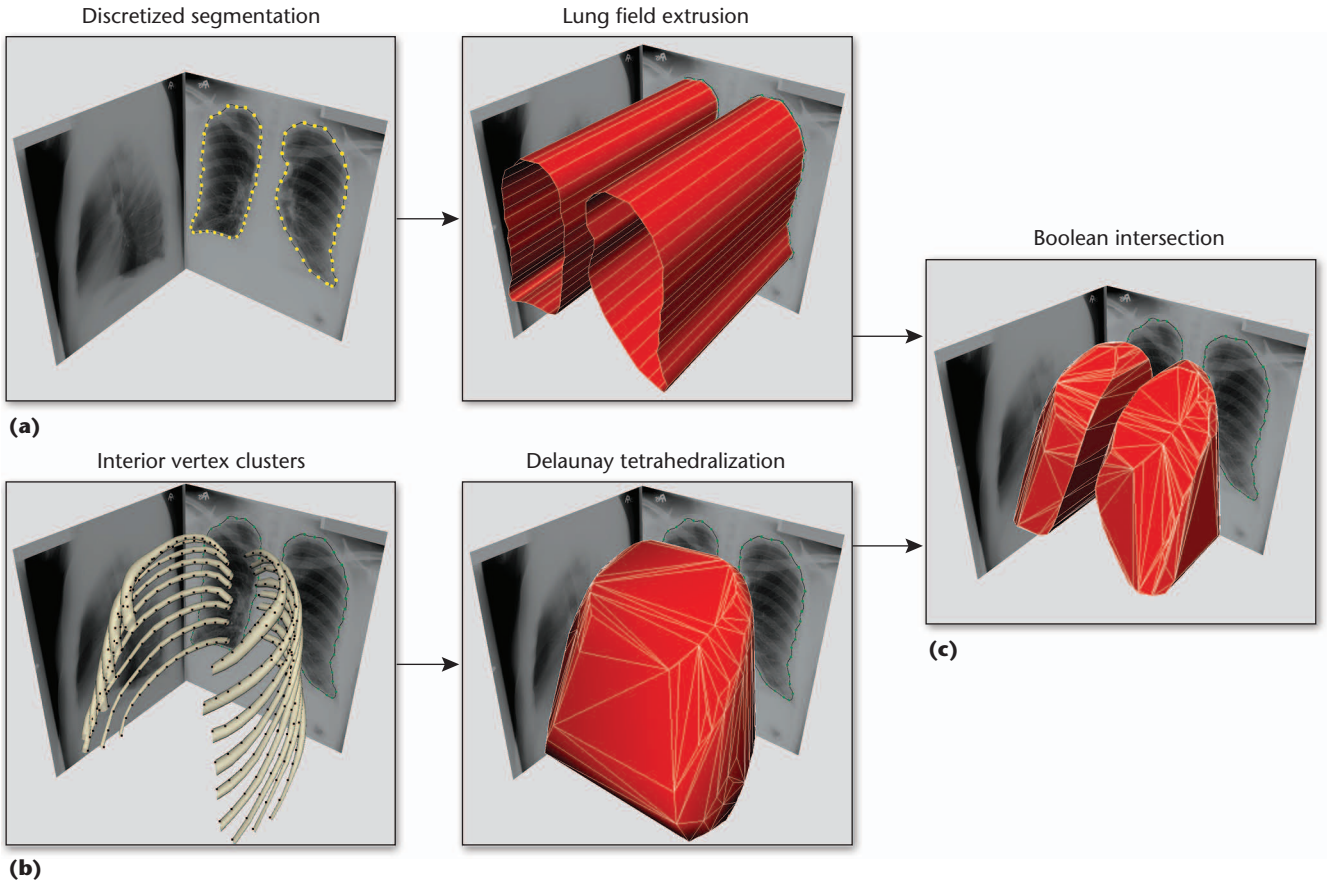

Figure 7. The initial phase of the lung reconstruction: (a) We extrude the discretized lung segmentation into two hollow lung fields. (b) We generate a polygon mesh by taking the Delaunay tetrahedralization of each rib's interior vertex clusters after the rib reconstruction is complete. (c) We create the initial reconstructed lung meshes by taking the Boolean intersection of the extruded lung fields and the Delaunay tetrahedralization. This represents the pipeline of the lung reconstruction algorithm to derive the initial model of the lung.

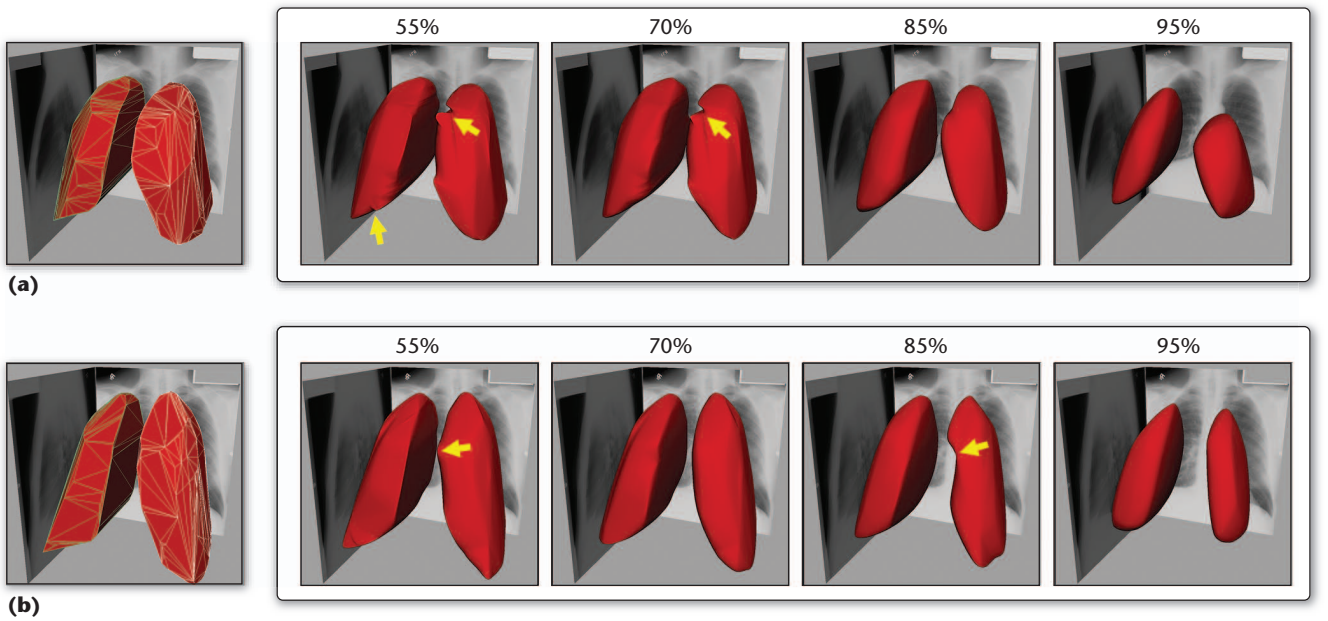

Figure 8. Interactive knowledge-based lung field refinement for two patients with significantly different bone structures: (a) patient 1 and (b) patient 2. The yellow arrows identify geometric artifacts of the reconstruction, which proper user-driven polygon reduction and subdivision smoothing can remove.

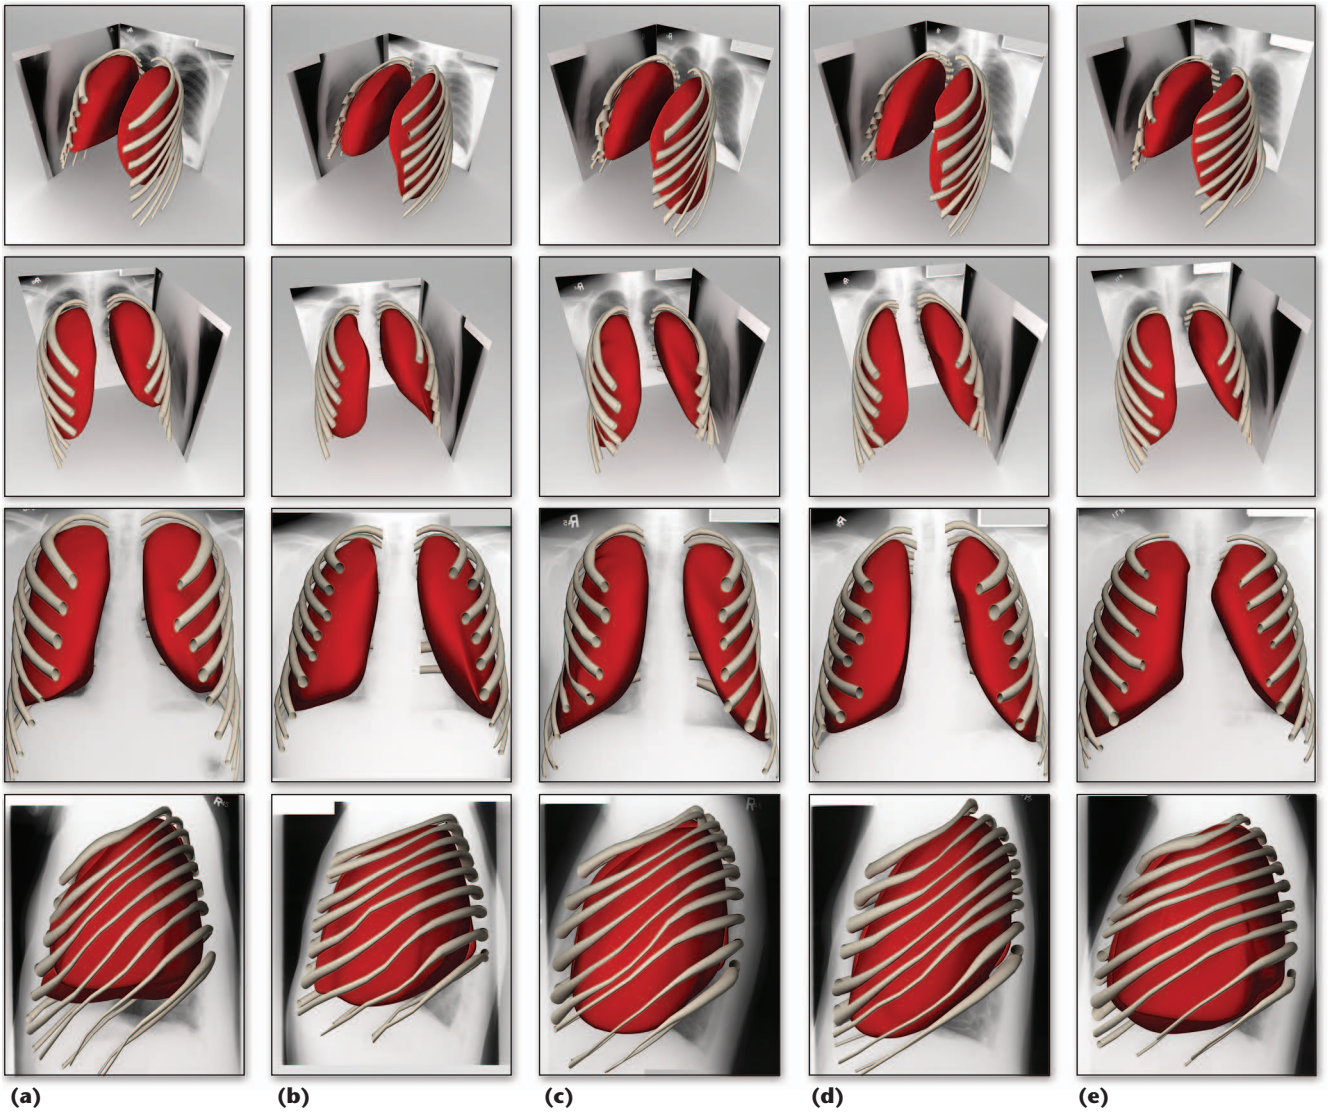

Figure 9. Final output from the rib and lung reconstruction algorithms for five patients with significantly different bone structures, rendered from an anterior-lateral-left view (top row), anterior-lateral-right view (second row), front (third row), and side (fourth row): (a) patient 1, (b) patient 2, (c) patient 3, (d) patient 4, and (e) patient 5. This illustrates the high quality of the resulting ribs and lungs generated by the reconstruction algorithms.

The initial resulting reconstructed lungs looked similar, but they produced geometric oddities at variable polygon reduction percentages. A person driving the process can easily spot abnormal phenomena in the smoothed lungs' surface and make the necessary adjustments in seconds.

We performed the rib and lung reconstructions on the same data sets we mentioned earlier in the lung segmentation section. Figure 9 shows the results for five patients, from several views. Overall, the reconstruction results look good, although some minor distortions appear in the ribs' lateral portion because we're currently unable to segment the ribs reliably using the lateral x-ray images. We hope to further validate the lung reconstruction algorithm's results against a CT scan of a human torso by creating simulated x-rays through volume renderings, performing rib and lung reconstructions, and then comparing the results to the original CT. Future comparisons with CT data should also help us gain knowledge about areas where distortions occur, so that we can mitigate those problems.

Application: Knowledge-Assisted Volume Clipping

One of the best potential applications for the lung reconstruction algorithm is to use the reconstructed lungs' boundaries to clip out the most interesting portion of data from approximate volumetric reconstructions. This will greatly improve a user's ability to visualize the data without manually identifying the lungs in the volume. Several volumetric reconstruction methods, such as

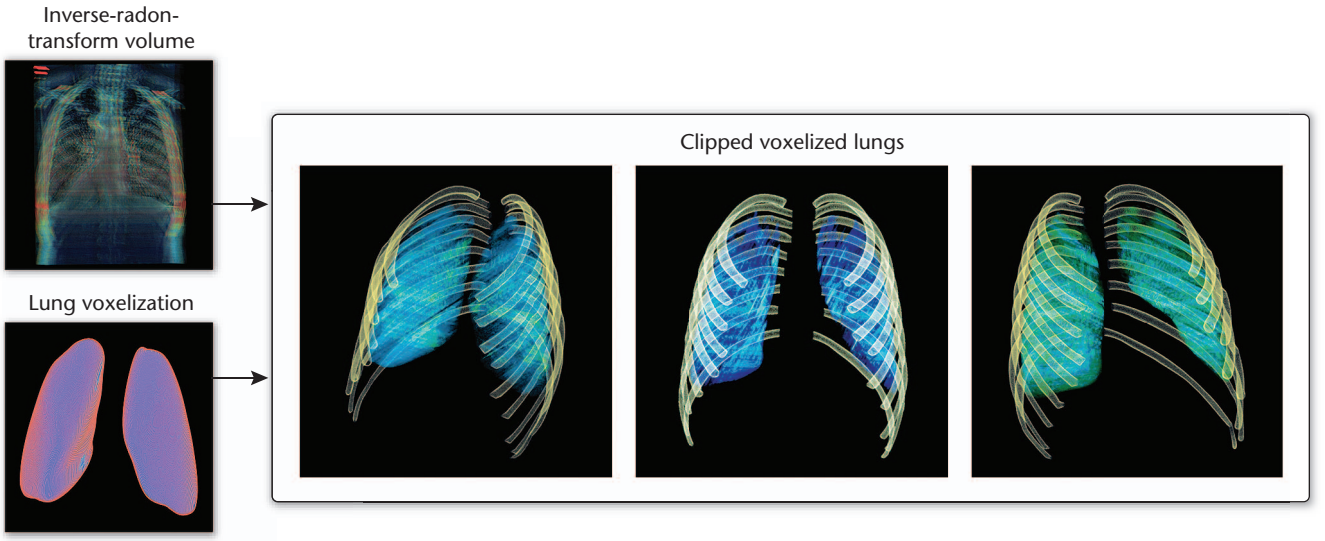

Figure 10. Lung voxelization and volume clipping, along with several renderings of the resulting volume. We clipped the approximate inverse-radon-transform volume on the basis of the mask defined by the lung voxelization that we created using the same set of x-ray images.

the inverse radon transform[5] and epipolar image analysis,[6] require only two input images. However, their output is approximate when used with PA and lateral x-ray image data. Also, the output is typically somewhat boxy because these methods project all the true 3D information onto a plane during the x-ray process. Volume renderings of such data appear very busy, and key features are difficult to discern (which is why the ability to remove voxels that don't correspond to a lung is useful).

In the example images presented here, we used the inverse radon transform on the pair of PA and lateral x-ray images to create the approximate volumetric reconstruction. However, the same technique would apply regardless of how we generated the volume.

To perform volume clipping on the basis of the reconstructed lung fields, we first voxelized these fields using the techniques that Nilo Stolte and Arie Kaufman described.[7] We then recursively filled the lung surface's voxelization to create a mask of the lungs. We gave a 0 intensity value to voxels in the inverse-radon-transform volume if they corresponded to the voxels outside the lungs in the mask. The resulting rendering of the clipped volume presents a far clearer picture of the data corresponding to the lungs (see Figure 10).

Although the resulting clipped volume from our 3D reconstruction techniques is an approximation of the lungs, it provides a new way to visualize the data, which could be helpful for disease detection applications. Visualizations generated by rotating the volumetric lungs while adjusting the transfer functions can help medical doctors notice subtle oddities that they might have overlooked if they had evaluated only the PA x-ray image.

Medical professionals will ultimately evaluate whether such reconstruction techniques can help increase the accuracy of current disease detection methods. Any technique that can even slightly decrease the number of false-positive and false-negative diagnoses is of great value. So, 3D reconstructions in situations in which only 2D x-ray images are currently available is a viable research area in which more work would be helpful.

Acknowledgments

We thank Riverain Medical for providing the data sets used in this article, and we thank the Ohio Department of Development and the Dayton Area Graduate Studies Institute (DAGSI) for providing funding for the Early Lung Disease Detection Alliance to pursue this research.

References

1. M. Chen et al., "Data, Information, and Knowledge in Visualization," *IEEE Computer Graphics and Applications*, vol. 29, no. 1, 2009, pp. 12–19.
2. F. Plourde, F. Cheriet, and J. Dansereau, "Semi-automatic Detection of Scoliotic Rib Borders Using Chest Radiographs," *Studies in Health Technology and Informatics*, vol. 123, 2006, pp. 533–537.
3. C. Koehler, T. Wischgoll, and F. Golshani, "3-D Reconstruction of the Human Ribcage Based on Chest X-Ray Images and Geometric Template Models," to be published in *IEEE Multimedia*;

http://doi.ieeecomputersociety.org/10.1109/MMUL.2009.57.

4. J. Barry, "Construction of Three-Dimensional Delaunay Triangulations Using Local Transformations," *Computer Aided Geometric Design*, vol. 8, no. 2, 1991, pp. 123–142.
5. M. Radermacher, "Radon Transform Techniques for Alignment and Three-Dimensional Reconstruction from Random Projections," *Scanning Microscopy*, vol. 11, 1997, pp. 171–177.
6. I. Feldmann, P. Eisert, and P. Kauff, "Extension of Epipolar Image Analysis to Circular Camera Movements," *Proc. Int'l Conf. Image Processing* (ICIP 03), IEEE Press, 2003, pp. 697–700.
7. N. Stolte and A. Kaufman, "Novel Techniques for Robust Voxelization and Visualization of Implicit Surfaces," *Graphical Models*, vol. 63, no. 6, 2001, pp. 387–412.

Chris Koehler *is pursuing a PhD in computer science and engineering at Wright State University. His research interests include 3D reconstructions, segmentation, and flow visualization. Koehler has a BS in computer science from the University of Akron. Contact him at koehler.11@wright.edu.*

Thomas Wischgoll *is an assistant professor of computer science and engineering at Wright State University. His research interests include scientific visualization, biomedical imaging, and flow visualization. Wischgoll has a PhD in computer science from the University of Kaiserslautern. Contact him at thomas.wischgoll@wright.edu.*