

2D X-ray airway tree segmentation by 3D deformable model projection and registration

Benjamin Irving¹, Tania Douglas², and Paul Taylor¹

¹ Division of Population Health,
University College London, UK

** benjamin.irving@eng.ox.ac.uk

² Department of Biomedical Engineering,
University of Cape Town, Rondebosch, South Africa

Abstract. Chest x-rays (CXR) still play an important role in detection of airway disease. Airway stenosis is a key sign of disease such as paediatric tuberculosis but is challenging to measure in CXR because the airway has relatively low contrast compared to overlying structures. In this study a novel approach to identify airways in CXRs is introduced, using a 3D statistical shape model to guide the segmentation. The 3D model is projected onto the CXR and aligned to the airways using four manual landmarks. The 3D shape model is then fitted to each CXR using an energy function based on image gradient, anatomical shadow and a regularisation term. Anatomical shadow is a novel feature introduced to detect the airway even without a clearly defined boundary. The algorithm achieved a mean error of 6.8 ± 2.6 pixels (0.82 ± 0.31 mm) on the 31 patient test set, a $25 \pm 17\%$ improvement on the initial linear alignment.

1 Introduction

The prevalence of tuberculosis (TB) remains high in many countries and accurate detection paediatric TB is poor [3]; however, there is increasing interest in specialised methods for paediatric airway analysis [1]. Chest X-ray (CXR) examinations are still a key component of pulmonary TB detection, and a common sign of pulmonary TB in children is airway deformation caused by lymphadenopathy [3, 8]. There is potential to automate the detection of the airway boundary, to assist in identification of stenosis, but this is challenging because the boundary is relatively low contrast (or incomplete) compared to better defined overlapping structures such as the heart, rib cage and vertebrae.

Segmentation of the trachea and complete left main bronchus (LMB) and right main bronchus (RMB) in 2D X-rays has not been previously achieved. Previously, Tezoo et al. [9] built a semi-automatic active shape model (ASM) based method to segment the trachea and bifurcation region, with 9 manual landmark points required. However, they do not clearly demonstrate any improvement of

** Benjamin Irving is currently with the Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, UK

the ASM over the initial affine alignment to the manual landmarks, and do not segment the complete LMB and RMB. Complete LMB and RMB segmentation is diagnostically important for the measurement of stenosis in paediatric TB [2].

This study uses a 3D statistical shape model (SSM) of the trachea, LMB and RMB to aid segmentation of the airways in CXRs. Projection of a 3D model onto 2D CXRs has been used for bone structures [4] but not for the airways. The 3D shape model along with novel 2D features are chosen to allow the airways to be detected while excluding overlapping structures with stronger edges and higher contrast. A model developed from CT is more complete than a model from radiographs where the RMB and LMB are often not completely visible.

2 Method

A 3D SSM of the airways derived from CT examinations is used in this study. As described in Figure 1, a silhouette of the mean airway is found and the points are projected onto a 2D plane and aligned to the airway region. This 2D projection is used to generate an energy function to measure the fit of the model to airway features in the CXR. The parameters of the 3D SSM are then optimised to determine the best fit of the projection to the 2D image.

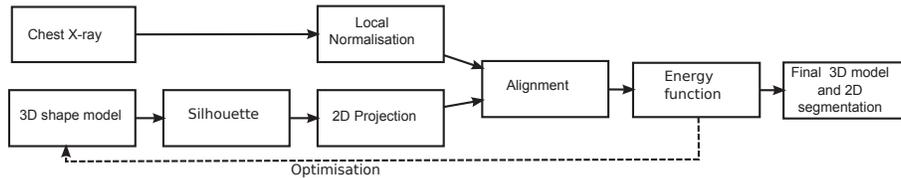


Fig. 1. An outline of the 3D to 2D projection and registration method. The 3D shape model is projected and aligned with a processed 2D CXR. An optimisation of 3D shape model modes of variation is performed using 2D image features.

2.1 Model Projection and alignment

Irving et al. [7] developed a 3D airway SSM consisting of a mesh representation of the mean airway (shown in Figure 3a) and the first 11 PCA derived modes of variation – from TB and non-TB paediatric patients with normal airways and stenosis. It is important to include examples of pathology in the statistical model in order to correctly detect X-ray cases with severe airway stenosis. This model is used in our method.

The silhouette of an object can be used to define the edges of the surface mesh that will be visible in a 2D projection. There are a number of methods to extract a silhouette, although, for small models (meshes with less than 10 000 faces) a brute force approach has almost the same speed as more complex methods [6].

A brute force approach was implemented to detect the airway silhouette vertices (\mathbf{s}_i) of the 3D mesh (see Figure 3a). The brute force approach was used to evaluate each edge in the airway shape mesh. Silhouette edges in a mesh are defined as edges separating a face with a normal facing towards the viewer and a face with a normal away from the viewer [6]. The normal (N) of each face in the airway surface mesh is found using the cross product of two edges³. The inner product of the face normal (N) with the viewing direction (E) is used to detect the orientation of each face. If $N \cdot E > 0$ then the polygon faces frontwards (as shown in Figure 2). If $N \cdot E < 0$ then the polygon faces backwards and if $N \cdot E = 0$ then the normal is orthogonal to the viewing direction.

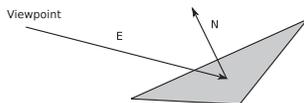


Fig. 2. The viewing direction (E) and outward facing normal (N) of a triangle, which are used to define the direction of the face

The radiographic projection of the shape (represented by the \mathbf{s}_i) was found by calculating the vector $\mathbf{v}_i = \mathbf{l}_0 + d\mathbf{l}$ from X-ray source \mathbf{l}_0 to each \mathbf{s}_i in direction \mathbf{l} using parameters based on the Lodox X-ray system dimensions ($\mathbf{l}_0 = [1000, 0, 0]$ mm, airway model centred at $[0, 0, 0]$ mm). The intersection of \mathbf{v} with the detector plane at $[-100, 0, 0]$ mm can then be calculated. The z-coordinate of each point is kept the same as no magnification occurs in the scanning direction of the system.

Once the silhouette points have been projected onto the detector plane, this algorithm requires some manual interaction to initialise the alignment of the model to the airways in the image. Procrustes analysis [5] was used to align the model with 4 manually annotated landmark points. Figure 3b shows the alignment of the 4 bifurcation points of the projected 3D model with the 4 manual annotations on the image. The optimal scaling, translation and rotation, calculated from the Procrustes analysis of the landmark points, are used to transform the silhouette and skeleton points.

2.2 Image Features

An energy function derived from image features was used to fit the projected shape model to the airway region of the CXR. This function consists of a gradient based component (E_{grad}) and an anatomical shadow component (E_{AS}).

The image intensity of the airway varies due to overlap with the ribs, vertebra and pulmonary structures. In order to reduce this variation, local nor-

³ The meshing procedure produces vertices of a consistent order. Therefore, the cross product of two edges will always lead to a consistent normal direction relative to the volume.

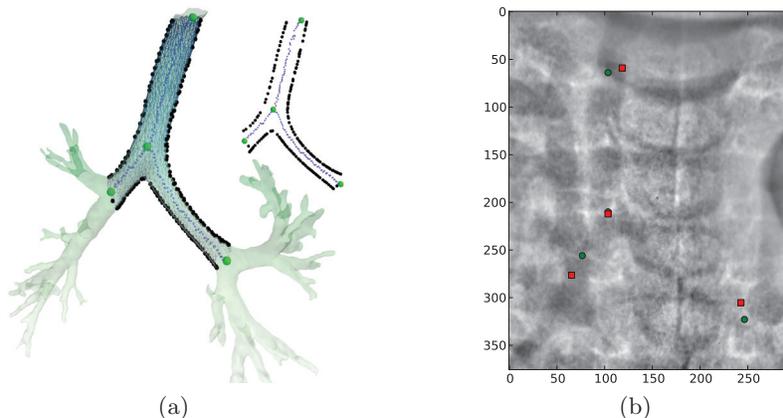


Fig. 3. Projection and linear alignment a) Airway silhouette detection and projection. The detected silhouette vertices (black), skeleton (blue) and branch start and end points (green) are projected. b) Procrustes analysis fitting the projected branch start and end points (red squares) with the 4 manual landmarks (green dots).

malisation, which has been used previously for airway enhancement [9], was applied to the image. Local normalisation can be described by $I_{ln} = \frac{I - G(I)}{\sigma_f}$ where $\sigma_f = \sqrt{G((I - G(I))^2)}$ is the estimate of the standard deviation at each pixel and $G(I)$ is the Gaussian blur function⁴ applied to image $I(x,y)$. Figure 4 demonstrates the effect of local normalisation on a CXR.

A component of E was derived from the image gradient. Vesselness is a common method for detecting tubular structures in images and Canny edge detection is a common technique for edge detection. Both these methods were attempted but failed to provide suitable airway features because of overlapping structures with greater intensity and stronger edges. Instead, the smoothed gradient of the image was used as one set of image features. The Sobel filter was used to extract the gradient and a Gaussian convolution of $\sigma = 5$ was used to smooth the gradient. The gradient derived energy function was calculated as the sum of the gradient at edge silhouette points:

$$E_{grad} = - \sum_{i=1}^p |\nabla I_{ln}(s_i)| \quad (1)$$

where s_i are the p projected silhouette points on image I_{ln} .

Edge detectors provide useful features for airway segmentation but if used alone were not always effective when there are additional strong edges. A second characteristic of the airway is that it appears as a shadow over higher contrast structures. Any structure overlapped by the airway should appear locally darker

⁴ Using the training set, $\sigma = 20$ pixels was used (approximately 80 % of the maximum airway radius). However, similar results on the test set were obtained for the range $\sigma = 10$ pixels to $\sigma = 30$ pixels.

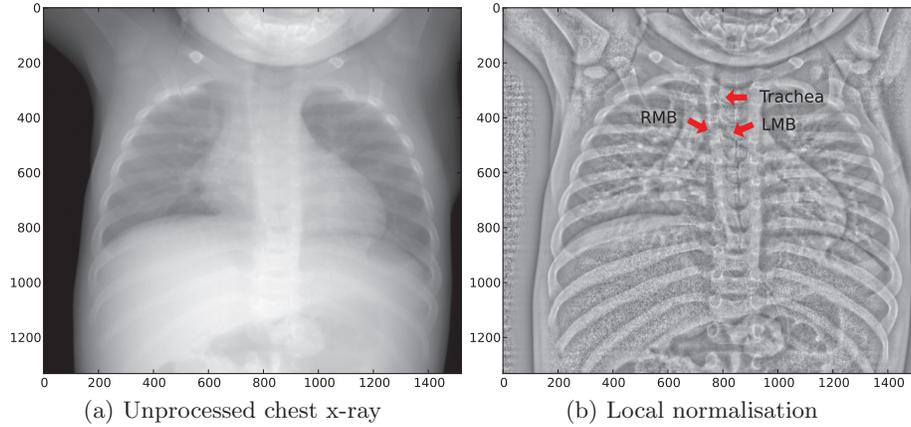


Fig. 4. Improvement in airway visibility after local normalisation

in the airway. Comparing the local greyscale inside and outside the airway, can be used to optimise the fit without well defined edges.

Therefore, *anatomical shadow*, a novel feature is introduced that compares local greyscale inside and outside the model, and updates the location and orientation of the sampling with each change to the 3D shape model. The outward direction of each silhouette point is defined as the vector from the closest skeleton point to the silhouette point (Figure 5). Once the orientation of each silhouette point is found then the mean greyscale value of kernels inside and outside the airway can be compared. The orientation of the line connecting the two circular kernels is always perpendicular to the current airway wall. The energy function was defined as:

$$E_{AS} = \sum_{i=1}^p \frac{c_{i,1} - c_{i,0}}{c_{i,0}} \quad (2)$$

where $c_{i,0}$ is the mean greyscale of the sampled kernel outside the airway and $c_{i,1}$ is the mean greyscale of the sampled kernel inside the airway at silhouette point s_i . The key advantage of this feature over conventional gradient is that the position and orientation is dependant on the shape model, making it a true comparison between regions inside and outside of the model, while using large enough regions to obtain a mean difference of the overlapping airway. The radius each kernel was chosen as 7 pixels and the distance between s_i and each kernel as 9 pixels based on the training set.

2.3 Optimisation

An energy function is used to optimise the fit of the 3D SSM to the airway regions in the CXR. This fit is constrained by the modes of variation of the 3D

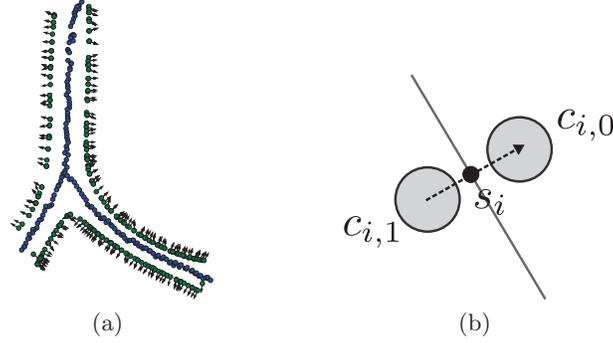


Fig. 5. Anatomical shadow feature a) The silhouette points (green), with unsmoothed centreline (blue) and outward pointing vectors. b) Greyscale difference between $c_{i,1}$ and $c_{i,0}$ for a single silhouette point (s_i).

model and any calculation of this energy function for a particular set of SSM weights involves a projection of the 3D model.

A statistical shape model can be described by variation along each eigenvector from the mean model [5], $\mathbf{x} = \bar{\mathbf{x}} + \Phi \mathbf{b}$ where $\bar{\mathbf{x}}$ is a $n \times 3$ array of vertex positions of the mean model that are deformed by m eigenvectors $\Phi = (\phi_1 | \phi_2 | \dots | \phi_m)$ with weights $\mathbf{b} = (b_1, \dots, b_m)$. Therefore, \mathbf{b} can be used to choose a suitable model that is the best fit to a single airway case.

A total energy $E(\mathbf{b})$ was created from a weighted combination of the energy functions:

$$E(\mathbf{b}) = \alpha_e E_{grad}(\mathbf{b}) + \alpha_g E_{AS}(\mathbf{b}) + r E_{reg}(\mathbf{b}) \quad (3)$$

where α_e and α_g are the weights. $E(\mathbf{b})$ is a function of the 3D shape model weights \mathbf{b} . With any choice of \mathbf{b} , the corresponding 3D shape is determined, this model is projected onto the 2D image and $E_{grad}(\mathbf{b})$ and $E_{AS}(\mathbf{b})$ are calculated from the position of the projected points. The lower the energy, the better the 3D model matches features in the 2D image.

An additional regularisation term (E_{reg}) was included with weight (r). Regularisation is required because the 3D model is poorly constrained by optimising on a 2D projection. E_{reg} is chosen as:

$$E_{reg} = \sum_{j=1}^m \frac{|b_j|}{\lambda_j} \quad (4)$$

where $b_j \in \mathbf{b}$ is the weighting factor of each eigenvector and λ_j is the eigenvalue of the j th eigenvector – equivalent to the variance of each mode within the dataset. This term penalises variation of the model away from the mean model, while penalising modes with larger variation in the dataset less than modes with smaller variation.

The local minimum of $E(\mathbf{b})$ was used to find the optimal \mathbf{b} , and, therefore, the best segmentation, for each chest radiograph. Most optimisation methods require a continuous function. We are working in discrete image space and so linear interpolation was used to determine the features used in $E_{grad}(\mathbf{b})$ and $E_{AS}(\mathbf{b})$ at any required value for s_i . Nelder-Mead downhill simplex optimisation method was used to find the local minimum of $E(\mathbf{b})$ from the initial mean model. At each step in the optimisation, $E(\mathbf{b})$ needs to be calculated a number of times in order to update \mathbf{b} . Each calculation of $E(\mathbf{b})$ requires a reprojection of the 3D model and calculation of $E_{grad}(\mathbf{b})$ and $E_{AS}(\mathbf{b})$ based on the projected silhouette points. Figure 3 shows initial mean model and Figure 6 shows the 3D model after optimisation to case 6 as an example.

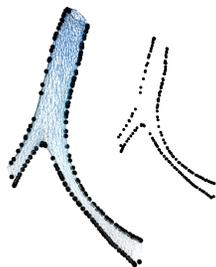


Fig. 6. Case 6 as an example of the optimised 3D shape and projection



Fig. 7. The manually drawn outline of an example airway (Case 21)

In summary, the Nelder-Mead method is used to optimise the energy function $E(\mathbf{b})$ in Equation 3 by finding 3D shape model weights \mathbf{b} that produced an optimal fit to the 2D features, as shown in Figure 1.

2.4 Dataset

A 41 patient CXR dataset was obtained from Red Cross Children's Hospital, South Africa. The images were acquired using the Lodox Statscan radiography unit (pixel size of 0.12 x 0.12 mm) and contained both TB and non-TB patients with ages from 3 months to 60 months. A number of TB patients had severe airway stenosis. The dataset was divided into a 10 patient training set (5 TB and 5 non-TB) and a 31 patient test set (11 TB and 20 non-TB). A small training set was used because only the weights (α_e , α_g and r) for the energy function need to be determined. Two sets of manually drawn airway outlines were used to evaluate the method (See Figure 7). Annotation 1 is used to generate results and Annotation 2, from a less experienced observer, is used for comparison. The observers segmented as much of the airway as they could identify, which varied.

3 Results

Parameters were derived using the 10 patient training set. To evaluate the performance of the method, the fitted projection was compared to the manual outlines of the 31 previously unseen test cases. The shortest distance between each silhouette point (s_i) and the manual outline was found by evaluating each segment of the manual outline and finding the minimum 2-norm distance (d_i). The mean $\mu_{dist} = \frac{1}{p} \sum_{i=1}^p d_i$ for the p projected points on each image was used for evaluation. μ_{dist} was calculated for each airway both after Procrustes alignment and finally after the shape model optimisation. The difference between the fits was calculated as a percentage.

The 10 training cases with manual annotations were used to find the optimal parameters ($\alpha_{edge}=0.4$, $\alpha_{AS}=0.6$ and $r=0.003$) with $\mu_{dist} = 5.4$ for E (Equation 3) by performing a grid search. Optimal parameters without E_{AS} only achieved $\mu_{dist} = 7.4$. The robustness of the algorithm to intra-observer variability of the landmark points was tested by assigning a second set of landmark points to the bifurcation regions in the training images. The mean error for all training cases was 5.4 ± 1.5 pixels (mean and standard deviation). This compared to a fit using the previous landmarks of 5.4 ± 1.9 pixels. This shows that the algorithm is robust to intra-observer variability in the selecting of the landmark points.

Using these parameters the method was applied to the 31 patient test set and the results compared to the manual annotations. Figure 8a shows the mean pixel distance (μ_{dist}) from the manual annotations for each case in the test set and Figure 8b shows the percentage improvement over the initial alignment. The test set achieved a 6.8 ± 2.6 pixel error (8.3 ± 2.0 for TB and 6.0 ± 2.6 for non-TB), a $25 \pm 17\%$ improvement on the initial linear alignment (29 of 31 cases performing better). Taking into account the voxel size of 0.12×0.12 mm, the average error is 0.82 ± 0.31 mm.

Annotations 2 was used to test the variability of the manual annotations and achieved a results of 7.5 ± 2.5 pixel error and $17.5 \pm 22\%$ (26 of 31 performing better). The distance between the two manual segmentations could not be compared because each observer segmented as much of the airway as they could see, which varied. Figure 9 shows four example fits to the test dataset. The algorithm fails to properly detect the LMB in Figure 9d but still improves on the initial alignment (as shown in Figure 8b). The method was implemented using Python and c++, taking 7.5 ± 0.7 minutes per case on a single core of an Intel i7 processor.

4 Discussion

Segmentation of the airways in CXRs is a challenging problem because of poorly defined borders and overlapping high contrast bone and heart structures. This study applies 3D SSM registration to segment the trachea, and complete LMB and RMB in CXR for the first time. The method requires 4 manual landmarks for the initial linear transform and then optimises the shape model using a novel

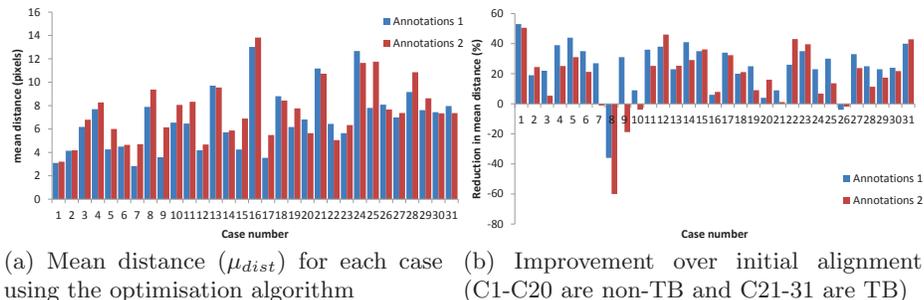


Fig. 8. Accuracy of the airway registration for each case of the test set

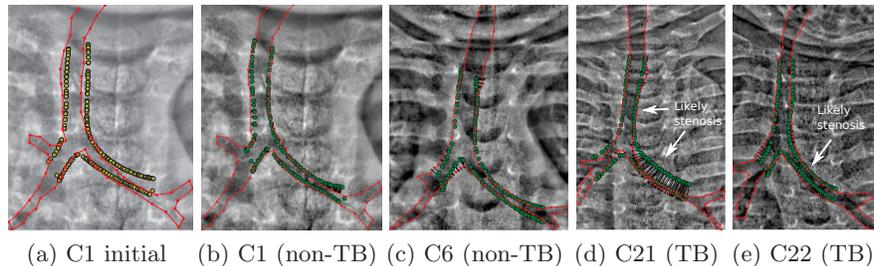


Fig. 9. a) C1 initialisation before fit. b, c, d, e) Comparison between the manual annotations (red line) and the optimised fit (green dots) for cases in the test set.

anatomical shadow feature. The model was able to detect airway shape variation with an error of 6.8 ± 2.6 pixels (or 0.82 ± 0.31 mm), a $25 \pm 17\%$ improvement on the initial similarity transform.

Figure 9a shows that the initial alignment is already a close approximation of the airway leading to a small error but does not capture individual airway changes. However, the model fit shown in Figure 9b captures the true airway shape. This high initial accuracy explains why the improvement in the error is only $25 \pm 17\%$. Figure 8 shows that this method consistently improves on the initial alignment.

The method was evaluated on TB and non-TB cases, where a number of the TB cases have severe stenosis. As illustrated in Figure 9d and 9e, this stenosis was captured in many but not all cases. This resulted in a slightly higher error for TB cases (8.3 ± 2.0 pixels) compared to non-TB cases (6.0 ± 2.6 pixels). The 3D airway SSM was trained using both TB and non-TB cases in order to capture both normal variation and stenosis. However, additional cases of severe stenosis or improved regularisation of the optimisation function might be required to capture all airway stenosis in CXRs.

An alternative method to an SSM could be an active appearance model (AAM) that incorporates both shape and greyscale variation. However, the air-

ways overlap with higher contrast anatomical structures such as the ribs, vertebrae and the heart. The position of these structures relative to the airway is sensitive to small changes in the position of the patient making AAM unsuitable. There are also a number of advantages to using a projected 3D shape model instead of a shape model derived 2D landmarks in CXRs. CT allows a more reliable model of the airway branches which are not consistently visible in X-rays. The problem could also be reformulated in 2D using the same CT dataset, by projecting the normalised 3D training set to 2D and using this to develop a 2D active shape model. However, the 3D optimisation also provides a personalised 3D airway model, and allows further improvements by incorporating additional X-ray views (e.g. Lateral CXR) to better constrain the optimisation.

Detection of the airway boundary allows automatic measurements of bronchial cross sections to be made to aid detection of stenosis. This method also produces a 3D personalised airway model, which could aid visualisation of stenosis.

5 Acknowledgements

The authors would like to thank the Commonwealth Scholarship Commission (CSC) for funding this research.

References

1. Andronikou, S., Irving, B., Hlabangana, T., Pillay, T., P.Taylor, Goussard, P., Gie, R.: Technical developments in post processing of paediatric airway imaging. *Pediatr Radiol* 43, 269–284 (2013)
2. Andronikou, S., Joseph, E., Lucas, S., Brachmeyer, S., Toit, G., Zar, H., Swingler, G.: CT scanning for the detection of tuberculous mediastinal and hilar lymphadenopathy in children. *Pediatr. Radiol.* 34, 232–236 (2004)
3. Andronikou, S., Wieselthaler, N.: Modern imaging of tuberculosis in children: thoracic, central nervous system and abdominal tuberculosis. *Pediatr. Radiol.* 34, 861–875 (2004)
4. Benameur, S., Mignotte, M., Destrepes, F., De Guise, J.: Three-dimensional bi-planar reconstruction of scoliotic rib cage using the estimation of a mixture of probabilistic prior models. *IEEE Trans. Biomed. Eng.* 52, 1713–1728 (2005)
5. Cootes, T., Taylor, C., Cooper, D., Graham, J., et al.: Active shape models-their training and application. *Comput. Vis. Image. Und.* 61, 38–59 (1995)
6. Hartner, A., Hartner, M., Cohen, E., Gooch, B.: Object space silhouette algorithms. In: *Theory and Practice of Non-Photorealistic Graphics: Algorithms, Methods, and Production System SIGGRAPH* (2003)
7. Irving, B., Goussard, P., Gie, R., Todd-Pokropek, A., Taylor, P.: Identification of paediatric tuberculosis from airway shape features. In: *Proc. MICCAI*. vol. 14, pp. 133–140 (2011)
8. Sanders, V.M., Pitcher, R.D., Douglas, T.S., Kibel, M.A., Daya, R.B., van As, A.B.: Digital radiographic measurement of the main bronchi: A pilot study. *Ann. Trop. Paediatr.* 29, 209–216 (2009)
9. Tezoo, T., Douglas, T.: Interactive segmentation of airways from chest x-ray images using active shape models. In: *Proc. EMBS*. pp. 1498–1501 (2012)