Preprint of the accepted manuscript in the Journal of Medical Image Computing.
Doi:10.1016/j.media.2014.05.007

Computer Assisted Detection of Abnormal Airway Variation in CT Scans related to Paediatric Tuberculosis

Benjamin J. Irving[a,*], Pierre Goussard[b], Savvas Andronikou[c], Robert Gie[b], Tania S. Douglas[d], Andrew Todd-Pokropek[e], Paul Taylor[a]

[a] *Centre for Health Informatics and Multiprofessional Education, UCL, London, UK*
[b] *Desmond Tutu TB Center, Department of Paediatrics and Child Health, Stellenbosch University, Tygerberg, South Africa*
[c] *Department of Radiology, Faculty of Health Sciences, University of Witwatersrand*
[d] *MRC/UCT Medical Imaging Research Unit and the Biomedical Engineering Programme, University of Cape Town*
[e] *Medical Physics, UCL, London, UK*

Abstract

Airway deformation and stenosis can be key signs of pathology such as lymphadenopathy. This study presents a local airway point distribution model (LA-PDM) to automatically analyse regions of the airway tree in CT scans and identify abnormal airway deformation. In our method, the airway tree is segmented and the centreline identified from each chest CT scan. Thin-plate splines, along with a local mesh alignment method for tubular meshes, are used to register the airways and develop point distribution models (PDM). Each PDM is then used to analyse and classify local regions of the airway. This LA-PDM method was developed using 89 training cases and evaluated on a 90 CT test set, where each set includes paediatric tuberculosis (TB) cases (with airway involvement) and non-TB cases (without airway involvement). The LA-PDM was able to accurately distinguish cases with airway involvement with an AUC of the ROC classification (and 95% confidence interval) of 0.87 (0.77-0.94) for the Trachea-LMB-RMB region and 0.81 (0.68-0.90) for the RMB-RUL-BI region – outperforming a comparison method based on airway cross-sectional features. This has the potential to assist and improve airway analysis from CT scans by detecting involved airways and visualising affected airway regions.

*Corresponding author
Email address: benjamin.irving@eng.ox.ac.uk (Benjamin J. Irving)

Keywords: point distribution model, airway, CT, computer assisted detection, classification, feature selection

1. Introduction

Tuberculosis is still prevalent – particularly in developing countries. The WHO estimates the annual number of new cases of TB to be 343 / 100 000 in Africa, higher than any other region (WHO, 2012). Childhood tuberculosis represents a large proportion of TB cases and is particularly difficult to diagnose; the disease is confirmed in less than 40% of cases (Schaaf et al., 1995). While most TB in adults is confirmed from sputum samples, childhood TB has a low sputum bacilli count, and, therefore, the diagnosis is based on a number of factors, with imaging playing a major role. Other indicators include: contact with an infected adult (often difficult to determine); symptoms and signs (vague and common); and tuberculin skin test (positive test only indicates exposure and false negatives can be attributed to HIV or other immune disorders). Each of these tests is imperfect but a combination is used to diagnose TB (Gie, 2003). Recent studies have highlighted increasing clinical interest and the need for further research into the diagnosis and treatment of paediatric TB (Sandgren et al., 2012).

Lymph node involvement is a key indicator of childhood TB, and as the lymph nodes enlarge the airways are compressed or deformed. This is because paediatric patients have smaller airways with less well developed cartilage – predisposing them to compression. This is known as lymphobronchial TB and is a current clinical focus in childhood TB research (Lucas et al., 2012; Goussard et al., 2013). Airway involvement due to lymphadenopathy is relatively common in children (between 29% (Andronikou et al., 2004) and 38% (Theart et al., 2005) of cases). This deformation can occur due to one enlarged lymph node, or compression can be due to an enlarged lymph node on both sides or a lymph node and a vessel, which can be used as an indicator of TB (Goussard and Gie, 2007; Andronikou et al., 2004). Figure 1 shows a MinIP projection with airway stenosis caused by lymphadenopathy in a paediatric pulmonary TB patient. The most affected airways include the bronchus intermedius (BI), left main bronchus (LMB), trachea and right main bronchus (RUL) (Lucas et al., 2012; Goussard et al., 2013).

Determining airway involvement and cause is important for the treatment of paediatric TB and other diseases affecting the airways (Andronikou et al.,

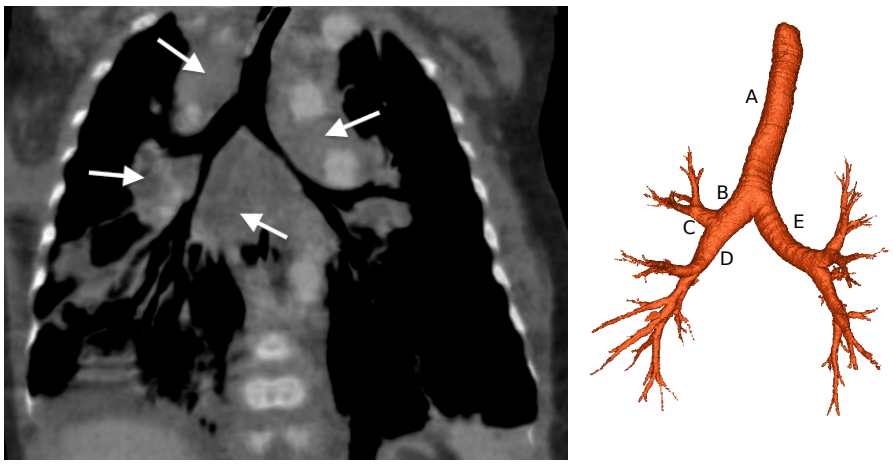

(a) (b)

Figure 1: **a)** Coronal minimum intensity projection (MinIP) image with arrows indicating lymphadenopathy. Stenosis of the left main bronchus (LMB) and bronchus intermedius (BI) is visible. Arrows indicate locations of lymphadenopathy. **b)** A diagram of the airway showing: *A*) Trachea, *B*) right main bronchus (RMB), *C*) right upper lobe bronchus (RUL) *D*) bronchus intermedius (BI) and *E*) left main bronchus (LMB)

2013). Bronchoscopy is the "gold standard" for determining airway involvement but is invasive, general anaesthesia is often required and the external cause of the airway involvement can not be seen. Recent studies suggest that CT with volume rendering can be used as an alternative to bronchoscopy (du Plessis et al., 2009) – offering the benefits of bronchoscopy while also allowing visualisation of the external cause of the airway involvement. This method allows manual measurements of airway cross-sections from the 3D rendering of the region, but requires considerable manual interaction in the form of setting thresholds, viewing parameters, and manual assessment of the airways, and a more automated approach to monitor airway involvement would be beneficial.

Thus, there is considerable value in automatically detecting airway involvement from lymphadenopathy to assist in the detection and assessment of paediatric patients with tuberculosis (and potentially other diseases). In this study we developed a method we call the local airway point distribution model (LA-PDM) to assess normal and pathological variation in local regions of the airway. Point distribution models are effective for capturing variation that is more complex than airway narrowing – where more complex variation

is related to airway involvement from lymphadenopathy.

Section 2 discusses current methods used for airway analysis and Section 3 outlines our proposed LA-PDM method. This method performs a 3D airway segmentation using chest CT images, and then a branch labelling of the segmented airway (Section 3.1). Surface point correspondence is developed between the dataset of segmentations (Section 3.3 and 3.4) and used to create point distribution models for local regions of the airway (Section 3.6). The LA-PDM can then be used to distinguish normal and abnormal variation in local airway regions of a new CT image. Section 4 applies the LA-PDM to a 90 patient test dataset of paediatric chest CT scans. The results of the LA-PDM on the test set is shown in Section 5 and compared to simpler branch diameter based features, and show promising results for detection of airway pathology related to paediatric TB.

2. Airway shape analysis

A number of studies have developed methods that automate branch diameter measurements as well as the broncho-aterial ratio from segmented adult airways (Kiraly et al., 2008; Tschirren et al., 2005; Palágyi et al., 2006). Applications include detection of chronic obstructive pulmonary disease (COPD) and asthma identification (Petersen et al., 2010; Fetita et al., 2010; Wiemker et al., 2004). These methods are useful but rely on the identification of pathology from variation in local airway diameter.

External changes such as lymphadenopathy produce more global changes in airway shape that are difficult to identify from an analysis of local features. Point distribution models (PDM) allow more complex airway shape variation to be modelled. Deligianni et al. (2006) use a PDM of an airway phantom to model breathing but do not consider inter-patient variability or detection of pathology. Pinho et al. (2011) used a PDM to build individualised models of the healthy trachea for patients with stenosis. The difference between the models can be used to assist with stent implants. Given the global variability of the airway shape and the complexity of a branching airway tree, this method is not suitable for the analysis of airway bronchi (Pinho et al., 2011).

Initial steps in the development of our LA-PDM method were introduced in Irving et al. (2011) which provided a proof of concept by using the trachea and main bronchi to distinguish cases with TB (acquired from South Africa) from cases without TB (acquired from the UK). This paper presents

a more detailed account of a refined method, extended to analyse multiple regions of the airway tree with improved region selection and correspondence detection. A full evaluation of the method is also presented, on a previously unseen dataset of cases with and without airway abnormalities that have been acquired from the same hospital. We analyse the shape model characteristics and compare the performance to features derived from airway cross sections. The authors are not aware of any previous methods that have detected abnormal airway variation from lymphadenopathy or paediatric TB.

3. Local airway point distribution models for abnormality detection

Figure 2 shows our proposed method that uses a PDM of local regions of the airway to detect airway abnormalities. Airways are segmented and the branches are labelled. Surface mesh correspondence is developed using a reference airway, and PDMs with trained classifiers are used to detect airway abnormalities.

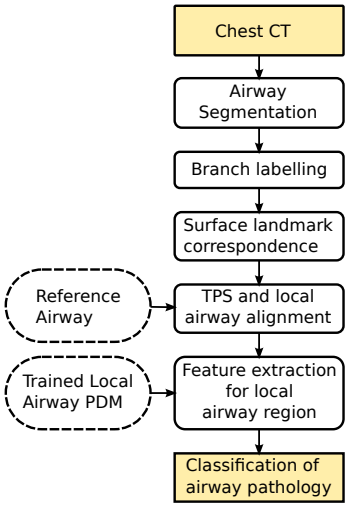

Figure 2: The steps for airway shape analysis using local airway point distribution models. The airway is segmented from a chest CT and the airway structure is extracted. Surface landmarks are generated and a representative airway from the training set is warped onto the airway to create vertex-vertex correspondence. Features are generated using the airway PDM and are used to classify airway pathology.

3.1. Segmentation and branch labelling

Initially the airway tree is segmented from each chest CT scan and each branch is labelled. There are a number of well established methods for airway segmentation and labelling, and we implement existing algorithms (Lo et al., 2012; Palágyi et al., 2006). The first axial slice containing lung volume was used as a reference position from which the segmentation was initialised. The glottis could be a more reliable landmark on the trachea but was not present in the majority of chest CT volumes used in this study because of its position in the upper airway (Shi et al., 2006). In this study we use an existing morphology-based airway segmentation method (Irving et al., 2009; Lo et al., 2012). In this method, the trachea is detected in the initial slice using location, size and circularity. Morphological closing and reconstruction is used to enhance airway cross sections in 3 dimensions in the volume. A threshold is then applied to the filtered airways and a region growing method initialised from the trachea is used to segment the airway.

The Palágyi et al. (2006) iterative thinning approach is used to extract the centreline, identify branch points and label branches from the segmented airway. This method iteratively removes voxels from the segmentation that do not affect the topology to extract a 1-voxel thick branching centreline. A pruning method based on branch length and location is then used to remove false branches. Finally, starting from the trachea, voxel connectivity of the branching centreline is used to iteratively identify branch points and label each branch. Branch points are found when a voxel of the centreline has more than two 26-neighbours. The segmentation was then labelled using a distance transform from the labelled centreline. Original voxel size was used during skeletonisation, and voxel dimensions were taken into account when the skeleton was converted to smooth branch centrelines in the following steps.

Figure 3 shows the labelled airway centreline before and after pruning of false branches and the labelled airway segmentation. Figure 4 shows the segmentation and centreline in relation to the original CT volume. In some cases, due to the size of the airway and scan quality, the third generation of bronchi were not fully segmented. This does not affect our shape analysis method because the Trachea, LMB, RMB and BI are used for analysis in this study. These branches have been reported to contain 218 of all 259 compressions in a 98 patient study (Lucas et al., 2012). Improved image resolution and reduction in artefacts would allow further generations of bronchi to be accurately segmented. However, the segmentation does affect labelling

and, therefore, 39 of the 179 cases required minimal user interaction (a single mouse click) to identify or remove a branch point from the centreline. The interaction is minimal, and fully automated branch labelling is not the focus for this study.

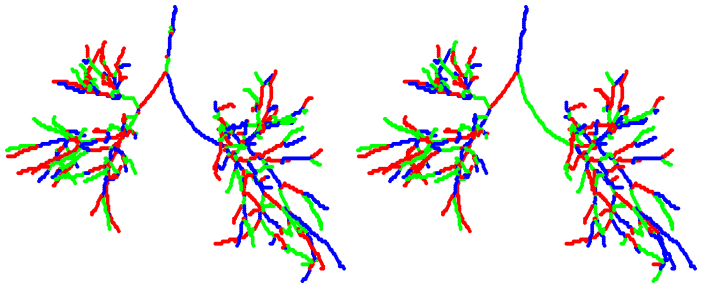

(a) Centreline before pruning (b) Centreline after pruning

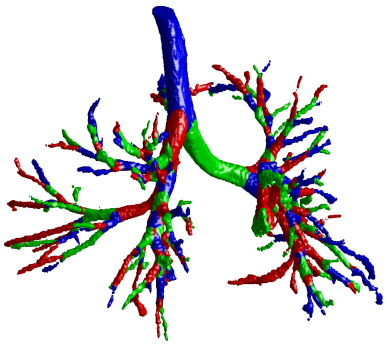

(c) Labelled airway segmentation

Figure 3: Centreline identification of the a segmented airway tree using a chest CT of a 20 month patient. Red, green and blue are used to illustrate distinct airway branches.

3.2. Surface point projection

A mesh of the airway surface was used to represent each airway and each face was labelled according to the branching structure of the airway. This triangular surface mesh was constructed from the voxel segmentation, and then smoothed using an implicit fairing method (Desbrun et al., 1999). Implicit fairing uses curvature flow to remove noise and uneven edges while preserving the mesh geometry. The smoothing was used to remove noise due to the resolution of the CT scan. Voxel resolution has more of an impact in

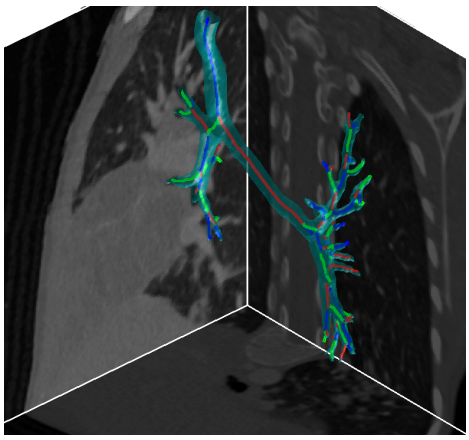

Figure 4: Rendering of a paediatric airway segmentation and branch centrelines with intersecting slices of the original CT volume from the dataset. Red, green and blue are used to illustrate distinct airway branches. Each branch centreline is individually smoothed.

the segmentation of paediatric airways compared to those of adults because the smaller volume means that airways are represented by fewer voxels.

Corresponding landmarks are required to align the airways in the dataset in order to develop the PDM. However, the airways do not have clear landmarks between bifurcations, and instead pseudo-landmarks were projected onto the airway surface mesh using the topological structure of the airway and the tubular shape of the branches. This was performed using the following steps: the centreline was smoothed and resampled equidistantly and pseudo-landmark points were projected orthogonally to the centreline at these resampled centreline points.

The centreline of the airways is represented as a one-voxel thick branching medial line. Linear interpolation between each set of neighbouring voxels was used to achieve a more detailed representation of the centreline. A moving average smoothing filter was applied to the interpolated points to remove noise added to the centreline from the voxel resolution. B-Spline interpolation of the centreline voxels is an alternative but tends to show more sensitivity to voxel offset (as shown in Figure 5). Nevertheless, either algorithm can be used for centreline smoothing. The branch is then resampled evenly m times to form landmark points (p_i) on the centreline.

At each resampled point (p_i), the tangent $T(p_i)$ to the centreline was estimated:

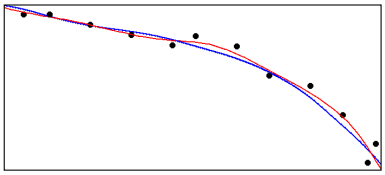

Figure 5: A subsection of a smoothed centreline of a single branch with voxel coordinates (black) and two examples of smoothing: a 5 voxel moving average smoothed centreline (blue) and B-Spline interpolated centreline of degree 13 (red).

$$T(p_i) = \frac{p_i'(t)}{|p_i'(t)|} \tag{1}$$

$$p_i'(t) = \frac{p_i(t+h) - p_i(t-h)}{2h} \tag{2}$$

where each $p_i(t)$ can be represented as a function in the vicinity of the landmark and $p_i(t+h)$ is a small progression along the centreline. A set of orthogonal vectors ($\mathbf{r}_i$) is then translated to each of the centreline points (p_i) and rotated to be orthogonal to the centreline (see Figure 6). Euler angles (x-convention) were used to construct a rotation matrix that can be used to map each of these orthogonal vectors onto each branch centreline. The Euler angles are defined by three rotations (ϕ, θ, ψ). By setting $\phi = -\psi$ the orthogonal vectors are mapped to any orientation of the centreline while removing the ambiguity of rotation around the centreline.

The intersection of the vectors (r_{ij}) with the triangular surface mesh was found using Möller and Trumbore (2005) ray/triangle intersection method. This method determines if the vector falls inside each face when intersecting a plane (with the plane normal defined by the face). Each face must be checked individually. To improve the speed of the search, only faces that fall within a local sphere of intersection are analysed. The radius of the sphere is chosen to be larger than the branch radius, guaranteeing that the point of intersection falls in the sphere (see Figure 8a). This representation allows a set of corresponding surface points to be mapped onto each branch of each segmented airway. Figure 7 and 8 show the mapping of surface points onto the Trachea-LMB-RMB and RMB-RUL-BI regions, respectively, of example airway segmentations.

The benefit of this representation is that the surface landmarks are based on the most consistent features of the airway tree: the bifurcation points and

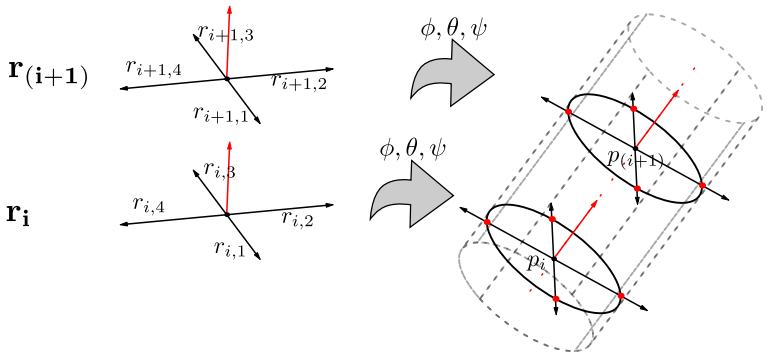

Figure 6: Mapping landmark points to the airway surface. Each set of vectors is translated to p_i and Euler angles are used to align the central vector with the centreline tangent. The intersection between the vectors and the airway surface is then found by ray/triangle intersection.

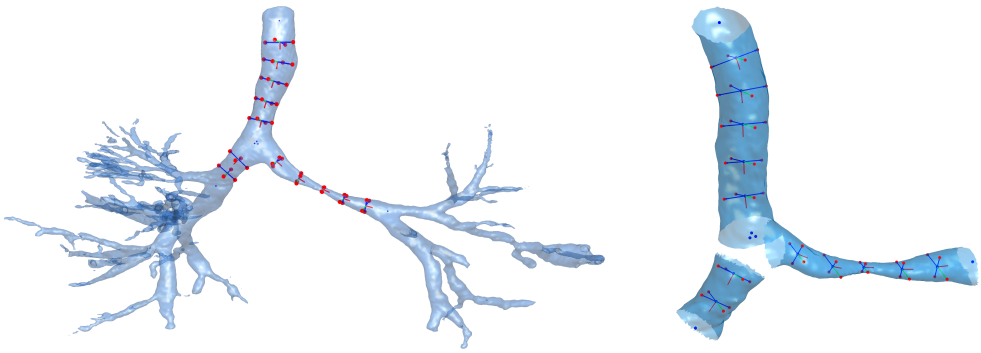

(a) Airway segmentation with landmark points for the Trachea-RUL-BI region

(b) Enlarged region showing the point mapping to individual branches

Figure 7: Mapping corresponding points onto the surface mesh of an airway

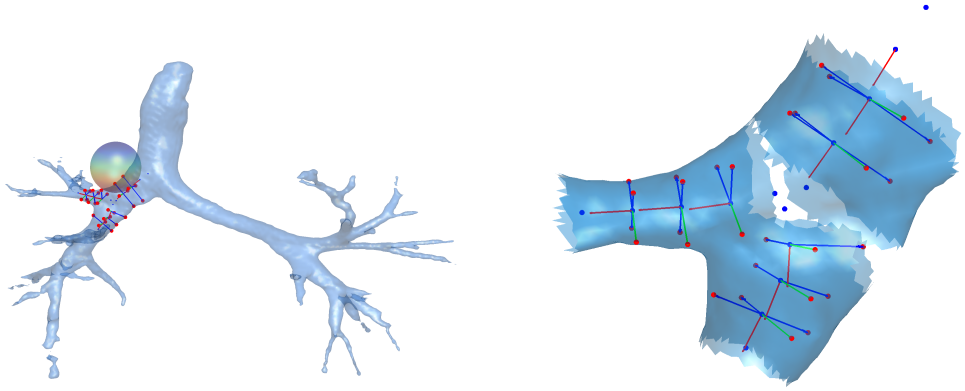

(a) (b)

Figure 8: Mapping corresponding points onto the surface mesh of an airway. RMB-RUL-BI region (with a single sphere showing the area of search for an example vector)

centreline. Statistical analysis methods could be applied directly to these corresponding landmark points. This would simplify the method because registration of an airway surface mesh template is not required. However, by only representing an airway with landmark points projected from the centreline, concave regions will be represented by a larger number of landmarks and convex regions will be represented by fewer landmarks; this is not ideal. In addition, if a large number of landmarks are chosen there is also a risk that the orthogonal radii will overlap, due to the curvature of the centreline, causing folding in the surface representation. As an example, the first two projections of the Trachea in Figure 7 show a region of larger curvature. Using a small number of landmarks to align a template mesh overcomes these challenges and a surface mesh provides a more detailed representation of the airway.

3.3. Mesh alignment step 1: Thin-plate-spline warp

One airway from the control dataset was selected as a template and a thin-plate-spline (TPS) warp was applied to align the template with each airway in order to represent each airway in the dataset with a corresponding mesh. Figure 9a shows a template warped by TPS to the RMB-RUL-BI region of the airway.

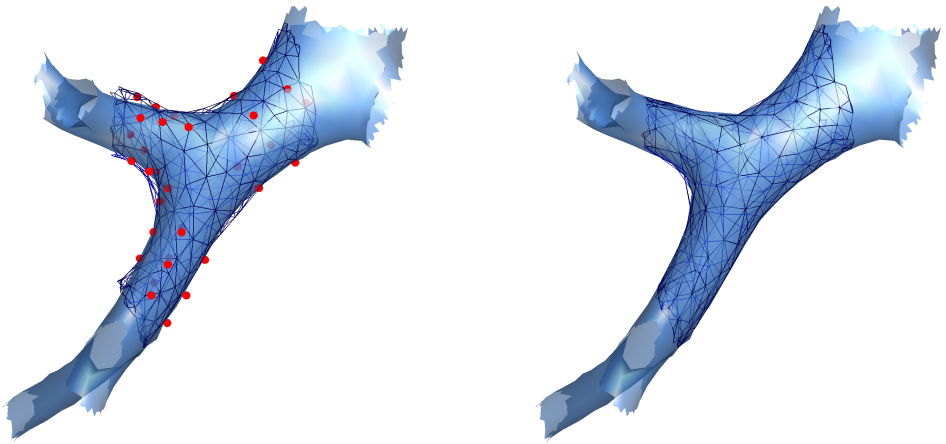

(a) Thin-plate spline warp by aligning landmark points (red) of a template mesh to the landmark points of a region of an airway

(b) Local alignment (performed after TPS warp) of each template vertex to the airway surface

Figure 9: Registration of a template mesh to an airway mesh for the RMB-RUL-BI region

TPS is a method of interpolation that minimises the bending energy of the surface (Bookstein, 1997, 1989). This method is also useful as a method of non-rigid registration; a surface with one set of landmarks is warped to a corresponding set of landmarks in a physically realistic way. The thin-plate spline function in 3D that transforms a point $\mathbf{p} = (x, y, z)$ is given by:

$$\mathbf{f}_j(\mathbf{p}) = \sum_{i=1}^{k} w_{ij} U(|\mathbf{p} - \mathbf{P}_i|) + a_0 + a_x x + a_y y + a_z z \tag{3}$$

where $\mathbf{f}(\mathbf{p}) = [f_x(\mathbf{p}), f_y(\mathbf{p}), f_z(\mathbf{p})]$ is the new position of the point $(\mathbf{p})$. $\mathbf{P}_i$ are the k landmark points on the shape, and w_{ij}, a_0, a_x, a_y and a_z are the weighting factors. $U(r) = |r|$ for the 3D case, where $r = \mathbf{p} - \mathbf{P}_i$. The weighting factors are found by writing Equation 3 as a system of linear equations $H = LW$, where the landmark and transformed coordinates are input as p and $\mathbf{f}(\mathbf{p})$, respectively. The weights, $W = L^{-1}H$, can then be found (as explained in detail in Bookstein (1997)).

3.4. Mesh alignment step 2: local alignment

The TPS warp will only be exactly aligned at the landmarks and further matching is required to exactly align the template mesh with each target mesh (as shown in Figure 9). The simplest method is to project the template mesh to the closest point on the target mesh (Paulsen et al., 2002; Hutton et al., 2003) but this can lead to unrealistic deformations and movement to the closest point while not representing small deformations.

Approaches such as Meller and Kalender (2004) and Kaus et al. (2003) optimize the fit based on the distance between the meshes while an additional force preserves the mesh structure. Our method is based on these methods but is enhanced for deformed tubular objects by adding a third term based on the surface orientation. For each vertex (t_i) on the template mesh, a force ($F_{i,tot}$) is calculated to direct the warp. This consists of a component based on the closest point on the object mesh (r_i) (Equation 4) and an internal forcing component is included to preserve the size of the faces. The force is calculated by the change in the distance of each of the m neighbouring vertices to a vertex t_i from the initial distance v_{ij}^{init} (Equation 5). These two forces are not enough to match to small indentations or protrusions found in the airways. Therefore, a local *inflation/deflation* force is also applied based on the normal of each vertex $\hat{n_i}$ (calculated from the normal of the surrounding faces). The amount of inflation and deflation is controlled by the distance and direction of the target mesh to $F_{i,1}$ (Equation 6).

$$F_{i,1} = \vec{r_i} - \vec{t_i} \tag{4}$$

$$F_{i,2} = \sum_{j}^{m} \hat{v}_{ji}(||\vec{v}_{ji}|| - ||\vec{v}_{ji}^{init}||) \tag{5}$$

$$\text{where } \vec{v_{ji}} = \vec{t_j} - \vec{t_i}$$

$$F_{i,3} = \hat{n}_i(\hat{n}_i \cdot F_{i,1}) \tag{6}$$

$$F_{i,tot} = \alpha F_{i,1} + \beta F_{i,2} + \gamma F_{i,3} \tag{7}$$

In Equation 7, the forces are weighted with α, β and γ. This procedure is applied iteratively until stability is reached. An example mesh after this local alignment is also shown in Figures 9b. This method enables any segmented airway to be registered to a template mesh, thus achieving vertex correspondence.

3.5. Bifurcation region relabelling

The previous sections present a method to develop correspondence between the surface mesh of segmented airway trees. This method can be applied to the whole airway tree or to a region of the airway. Figures 8 and 9 illustrate regional based airway analysis and this section discusses a method to refine the branch labelling at bifurcation.

Each face in the mesh was assigned a branch label by assigning the label of the closest voxel of the labelled skeleton. However, this labelling criterion leads to variation in label assignment in the branch bifurcation region. Mesh faces in a bifurcation region are assigned to the connected branches but this assignment is biased towards branches with a smaller diameter and is influenced by branch shape and orientation. This leads to irregular branch end points which adds additional complexity to the statistical shape representation. As an example, Figure 3c shows how most of the bifurcation region is assigned to the trachea.

To produce consistent branch labels at bifurcation, adjustments are made to this label assignment. A face (m) is only labelled as belonging to a particular branch if:
1) The centroid (c_m) of m falls between the two branch end points (creating a well defined branch end by cutting of the branch orthogonal to the centreline at the branch beginning and end bifurcation)

$$\hat{n}_{j1} \cdot \hat{n}_{m1} > \alpha \text{ and } \hat{n}_{j2} \cdot \hat{n}_{m2} > \alpha \tag{8}$$

2) It is near the branch centreline (ignore neighbouring branches)

$$|c_m - e_{j1}| + |c_m - e_{j2}| - |e_{j1} - e_{j2}| < \beta \tag{9}$$

where e_{j1} and e_{j2} are the start and end points of the branch, and therefore derived from the start and end positions of the extracted branch centrelines (shown in Figure 3b). n_{m1}, n_{m2} are the directions from e_{j1} and e_{j2} to c_m (see Figure 10). α and β are constants identified by qualitatively assessing the training set as 0.2 and 25 mm respectively. m that are not assigned to a branch are labelled as the bifurcation region. This method reduces variance due to region selection by creating consistent edges between branch labels as demonstrated in the branch endpoints of Figure 11 and 12 introduced later.

3.6. Airway shape features

Feature vectors are derived from both LA-PDMs of airway regions and cross-sectional sampling of each branch in the airway tree.

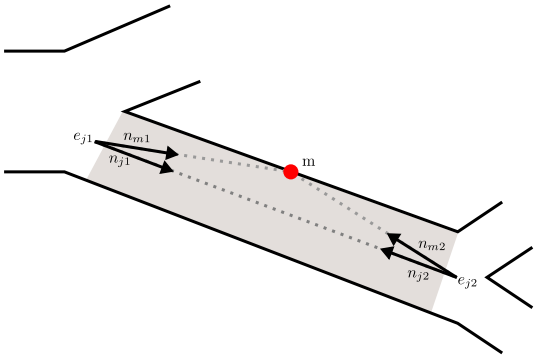

Figure 10: Relabelling the surface mesh where m is the face to be relabelled. m is labelled according to the face's position in relation to the branch end points e_{j1} and e_{j2}.

3.6.1. LA-PDM features

Each airway is represented by a dense surface mesh. The corresponding vertices of each airway surface mesh can be used to build a PDM of the airway and extract a set of principal modes that represent the variation of each airway.

Some of the variation between airways in the dataset is due to size, position and rotation, which are related to patient age and scan position, and not of interest for the detection of airway pathology. Therefore, the airways are first aligned by generalised Procrustes analysis (GPA). GPA is used to iteratively align objects represented by a set of corresponding points using translation, scaling and rotation.

Principal component analysis (PCA) is an integral part of developing a point distribution model (PDM). PCA applies a linear transform that projects the vertices onto an uncorrelated space and can be used to extract relevant features (Cootes et al., 1995). PCA modes are ordered by the variance and, therefore, by selecting the subset of the modes with the most variance, each shape can be represented by a feature vector of lower dimensionality than the input feature vector.

As input into PCA, each object is represented as a $3n$ dimensional stacked vector of mesh vertices where n is the number of vertices in the mesh; for the Trachea-RMB-LMB regions $n \approx 1700$. Therefore, the GPA aligned 3D landmark points of the object $\mathbf{x} = [(x_1, y_1, z_1)...(x_n, y_n, z_n)]$ are now represented as a single vector $\mathbf{x_i} = (x_1, y_1, z_1, ..., x_n, y_n, z_n)^T$ for object i.

PCA is computed from the covariance matrix and it can be shown that the eigenvectors (Φ) of the covariance matrix can be used to project an airway

($\mathbf{x}$) into uncorrelated space, where the dimensions of the uncorrelated space are defined by the orthogonal eigenvectors. This projection is:

$$\mathbf{b} = \Phi^T(\mathbf{x} - \bar{\mathbf{x}}) \tag{10}$$

where $\mathbf{b}$ is the representation of the airway in the new space. Therefore, an airway can be represented in terms of the mean shape and a displacement along each principal component by $\mathbf{b}$:

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi\mathbf{b} \tag{11}$$

Of particular interest is that PCA results in an ordered set of eigenvectors, where the contribution of the variance of each eigenvector (ϕ_i) is represented by the eigenvalues (λ_i). A shape can, therefore, be approximated by a set of m eigenvectors $\Phi = (\phi_1|\phi_2|...|\phi_m)$ and $\mathbf{b} = (b_1, ..., b_m)$. This means that a dense surface mesh model of a shape is reduced to a much shorter feature vector represented in terms of the principal components of shape variation.

A common way of choosing the number of principal components (eigenvectors) that will represent the shapes is based on the amount of variance they represent. Typical choices range from 90% to 99.5% of the variance (van Ginneken et al., 2002). The variance represented by the first m eigenvectors can be calculated from the sum of the first m eigenvalues over the sum of all $3n$ eigenvalues:

$$f = \sum_{i=1}^{m} \lambda_i \Big/ \sum_{i=1}^{3n} \lambda_i \tag{12}$$

Each object is now represented by a vector $\mathbf{b}$ of length m. These vectors represent the object shape (with an accuracy of up to the chosen variance f) and can be used to distinguish and classify airway shapes.

3.6.2. Local region analysis

Variation in airway shape could potentially be modelled at a number of levels from the variation in a single cross-section of a branch to global variation of the entire airway. In this study, local regions of the airway are modelled individually, where each region consists of a parent branch, the bifurcation region and two child branches. Modelling each branch individually is not as effective because variation in the position of a branch relative to neighbouring branches is important; lymphadenopathy can lead to branch

deformation as well as stenosis (an example of tracheal displacement by lymphadenopathy is presented in (Andronikou and Wieselthaler, 2004)).

Alternatively, shape models could be applied to the entire airway. However, PCA is a linear projection to a lower dimensional space and with each additional generation of a tree structure the variation of the airways becomes increasingly non-linear. While there are benefits to modelling the relationship between parent and child branches, there is limited value to modelling the relationship between the shape changes of second order connected branches. There are non-linear formulations of PCA (including Kernel PCA in (Schlkopf et al., 1997)), but modelling the airway as a set of 3-branch structures using PCA seems the most promising approach.

Therefore, this method creates a number PDMs of the airway for variation 3-branch regions. Each 3-branch region overlaps the previous region. For example, the RMB is represented in both the trachea-RMB-LMB and RMB-RUL-BI regions. These regions are automatically identified, a point distribution model is built from the dataset and each 3-branch region is assigned a feature vector ($\mathbf{b}$) that can be used for classification.

In paediatric pulmonary TB, the most common locations of lymphadenopathy are subcarinal (90%), hila (85%), anterior mediastinum (79%) and paratracheal (63%) (Andronikou et al., 2004). Bronchial compression caused by lymphadenopathy is most apparent in the BI, LMB, RMB and trachea (Andronikou et al., 2004; Lucas et al., 2012). Therefore for identifying TB, two regions were modelled (trachea-RMB-LMB and RMB-RUL-BI). This could be extended to the rest of the airway if useful for modelling additional airway pathology. The model is aligned using the branching structure and extracted mesh, and because three branches are modelled at any one time, no more complexity is added by analysing further generations.

3.6.3. Radius based features as an alternative to LA-PDM

LA-PDM provides a solution to a novel application and there are no comparison methods available. However, manual measurements of airway diameter are used clinically as an indicator of airway involvement and are generally assessed in terms of proximal and distal involvement in each bronchi (du Plessis et al., 2009). The method that we have proposed in the previous sections can implicitly measure airway cross-sections as part of the landmark placement step (Section 3.2). Therefore, as a more conventional comparison to the LA-PDM, feature vectors based on cross-sectional measurements of the airway bronchi were also derived. As discussed in Section 3.2, vectors

orthogonal to the branch centreline were projected to the surface and the intersection between the mesh and the vector was found. This provided radius measurements in four orthogonal directions for n equidistant samples along each branch. In this comparison we divide each bronchi into three regions: proximal, middle and distal and derive a single mean diameter ($\bar{d}$) for each region:

$$\bar{d} = \frac{1}{2(q-p+1)} \sum_{i=p}^{q} \sum_{j=1}^{4} ||\mathbf{r}_{ij}|| \tag{13}$$

where p and q are the scalars defining the start and end points of the region from the n centreline samples in the branch, and $\mathbf{r}_{ij}$ is the radius at position i and orientation j (as illustrated in Figure 6). These features are normalised across airways by the branch length (l). As with the PDM derived features, the same 3-branch regions were used for this method, and therefore, 3 branches each with 3 mean region diameters, resulting in 9 features, were used for classification, and compared to classification using the LA-PDM derived features.

Linear discriminant analysis is used to classify these airway regions using these derived features with leave-one-out cross validation.

4. Application of the LA-PDM to Clinical Cases

This section describes the application of the LA-PDM method to the paediatric Chest CT dataset.

4.1. Paediatric Chest CT Cases

A set of 89 paediatric non-TB and TB cases were acquired in 2010 from Tygerberg Hospital (South Africa) and Great Ormond Street Hospital (U.K.). These cases were used to develop the method and determine the parameters used in the algorithm. In 2012, a dataset of 90 patients with and without TB were all collected at Tygerberg Hospital. These were previously unseen and used only for evaluation in this study. Several cases with artefacts from tubes or severe movement were excluded from the study.

The test dataset included 42 non-TB patients with a mean age of 3.1 ± 3.8 years and 48 TB patients with a mean age of 2.4 ± 2.8 years. Patients under 5 years have more malleable airways and are, therefore, more predisposed

to airway deformation. These images were acquired using the Siemens SOMATOM Sensation 40 and the Siemens SOMATOM Emotion 6. The pixel size in each slice varied from 0.21 - 0.54mm and the slice thickness varied from 0.6 to 1.5 mm.

Patients with suspected TB only undergo a CT scan when there are signs or symptoms of airway involvement (du Plessis et al., 2009) and therefore, all TB cases in this dataset have suspected airway involvement. Suspected airway involvement is diagnosed from signs of obstruction of the large airways including stridor and wheezing or tracheal cough, or radiologically. Furthermore, all TB cases used in this study were classified as having probable or definite TB, where definite is culture confirmed and probable is not culture confirmed but combines clinical grounds, radiological assessment, Mantoux skin test and contact history.

In order for airway involvement to be compared to controls, a pulmonologist selected cases without airway involvement or tuberculosis. These cases include children with the following conditions: parenchymal lung disease such as interstitial lung disease, congenital lung malformations which affect the parenchymal tissue and not the airways; suspected lung metastases; and infective conditions other than TB. These will include cases of bronchiectasis, cystic fibrosis and pleural disease.

4.2. LA-PDM implementation

This section provides additional details on the experimental set-up and parameter settings.

For landmark generation, linear interpolation between the centreline voxel coordinates was performed and 100 points were sampled between each voxel. This was chosen to be suitably dense, from which, to derive resampled equally spaced landmark points. The centreline was then smoothed using a moving average. Five voxels were used to calculate the moving average at each point as this removed the appearance of individual voxel noise. Therefore, 500 points were used to calculate the smoothed centreline from the interpolated points.

A single case with no visible pathology (from the training set) was used as the template for registration. To improve computational efficiency, the airway mesh was reduced to 20% of the original number of vertices while maintaining vertex-to-vertex correspondence using a surface simplification method (Garland and Heckbert, 1997). Therefore, the Trachea-RMB-LMB

region was represented by 3286 faces and 1733 vertices, and the RMB-RUL-BI regions were represented by 1093 faces and 618 vertices.

The number of resampled centreline points used to generate surface landmark points were: Trachea, 5; LMB, 5; RMB, 2; BI, 3 and RUL, 3. These were chosen to provide similar spacing for each branch. Each centreline point was used to generate four surface landmarks, resulting in 48 and 32 landmarks used to represent the Trachea-LMB-RMB and RMB-RUL-BI regions, respectively (as shown in Figure 7 and 8), and were used to register the template mesh. These values were chosen so that the landmarks were suitably far apart that the orthogonal radii never intersected – which would cause folding in the surface representation. The number of resampled centreline points used, for radius based measurements, was [100 100 50] (number of points in each branch of a 3-branch section). At each of these equidistant points, 4 points were projected to the surface.

The accuracy of the mesh alignment was evaluated by comparing volumes generated from the template mesh and airway region of interest using the Jaccard distance (the complement of the Jaccard index). The proportion of error in the registration (V_{dif}) and a opening of the error using a 6-connected structuring element (V_{open}) as a function of α, γ and β (Equation 7) were used for evaluation using a grid search. The opening removes single voxel errors that may be introduced during re-voxelisation of the mesh for comparison, and, therefore, V_{open} represents larger local errors that have more impact on accuracy. For stability, only $\alpha < 1.0$ and $\gamma < 1.0$ were considered when choosing optimal values. A step size of 0.1 for α and γ, and $\beta = 0.5, 1.0, 1.5$ were used in the search. Optimal values were found to be $\alpha = 0.1$ and $\gamma = 0.8$ with a mean $V_{dif} = 0.021 \pm 0.009$ and $V_{open} = 0.0022 \pm 0.0049$. Just using closest point mapping (Equation 4), the errors were $V_{dif} = 0.029 \pm 0.019$ and $V_{open} = 0.009 \pm 0.015$. Using closest point mapping and the mesh structure term (Equations 4 and 5), the errors were $V_{dif} = 0.027 \pm 0.018$ and $V_{open} = 0.006 \pm 0.011$. Therefore, including the *inflation/deflation* term improves the registration of the template and reduces the variance in the accuracy of the fit. This is particularly noticeable for V_{open} (larger local errors) and agrees with qualitative observations that adding the third term improves the registration for narrowed and stenosed branches. This optimisation could be performed on an individual basis to choose parameters for each case. However, in this work the same fit parameters are used throughout the dataset.

Figure 11 and 12 show alignment of a template mesh to the trachea-RMB-LMB and RMB-RUL-BI regions of 24 example cases from the 90 patient

test dataset. These cases are used to illustrate the typical variation found in the dataset and the template mesh alignment (the darker region is the aligned template mesh). As can be seen, differences between cases with airway involvement and controls are sometimes not clearly apparent.

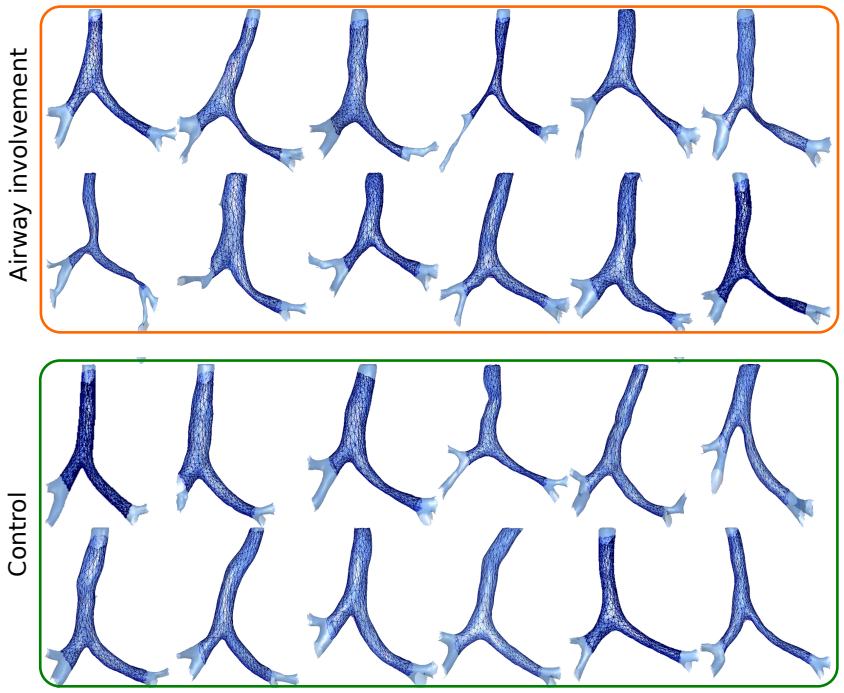

Figure 11: Example cases of registration of the template mesh to the trachea-RMB-LMB region. The first two rows are patients with airway involvement and the remainder are controls.

4.3. Model Variation and feature extraction

PCA modes of variation were extracted from the PDM model. The variation of the first 5 PCA modes for the trachea-RMB-LMB and RMB-RUL-BI regions are shown in Figure 13 and Figure 14. The centre column shows the mean airway ($\bar{\mathbf{x}}$) with variation ($\mathbf{x} = \bar{\mathbf{x}} \pm \mathbf{\Phi}_i b_i$). Using $b_i = \pm 3\sqrt{\lambda_i}$ shows the variation of one mode three standard deviations from the mean in the dataset (Cootes et al., 1995). These principal components exhibit a range of changes including branch narrowing, local stenosis, deformation, length and angular changes. These modes also include variation that is considered pathological such as deformation and stenosis of branches. For example, the

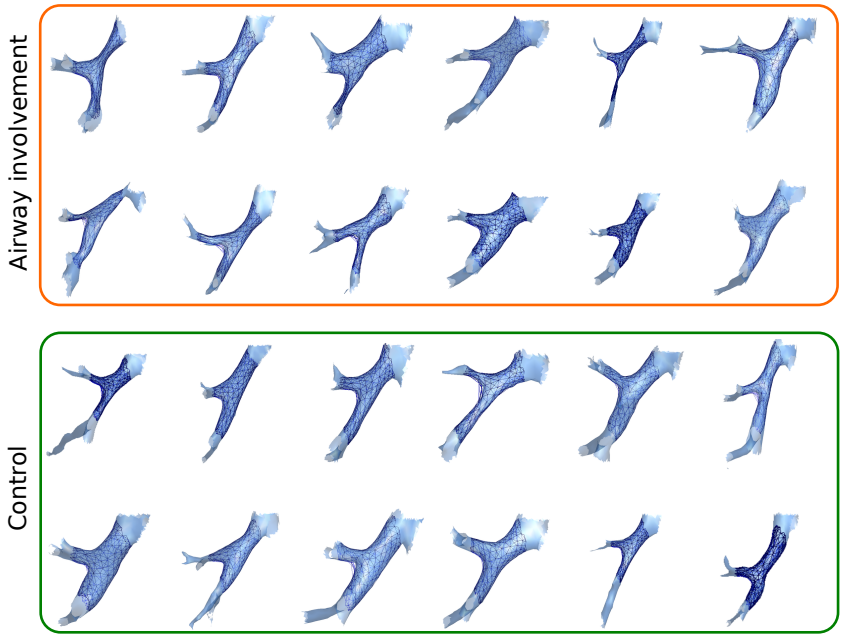

Figure 12: Example cases of registration to the template mesh to the RMB-RUL-BI region. The first two rows are patients with airway involvement and the remainder are controls.

fifth mode of Figure 13 shows central narrowing of the LMB while the first and second modes of Figure 14 show types of region based narrowing. Note that due to the perspective some of the 3D variation is not visible in the figures.

Figure 15 shows two example airway segmentations from a TB patient and control. In this example there is narrowing of the LMB and BI for the TB case. The representation of the airway variation in this example by the LA-PDM is shown later in Figure 21. Small differences are also visible between the 11 component representation of the trachea-LMB-RMB region (dark blue) and the original airways (light blue). These differences are represented by the remainder of the principal components that are excluded from the model.

4.4. Model evaluation

Before analysis of the classification accuracy, the characteristics and performance of the point distribution model were analysed. Measures of compactness, generalisation and specificity, that are outlined by Styner et al. (2003), were used to evaluate the model.

A compact model can represent the variance within the dataset with only a few parameters. The variance can be calculated as a function of the number of shape parameters (modes of variation).

The model generalisation represents a model's ability to represent unseen cases (Styner et al., 2003). This was evaluated using leave-one-out reconstruction of the dataset. For each case in the dataset, the PDM was built without the case, and used to reconstruct the case using M modes of variation. The mean absolute distance (MAD)[1] was then used to approximate the error of the PDM representation of each case. The mean and standard deviation of the MAD values for the entire dataset was calculated. Confidence intervals were found by the Central Limit Theorem.

Model specificity evaluates the ability of the model to only generate shapes that are similar to those found in the dataset (Styner et al., 2003). A thousand new instances of the feature vector $\mathbf{b_i}$ were generated for each PDM. Monte Carlo simulations were used to randomly generate $\mathbf{b_i}$ from a multivariate normal distribution with standard deviation $\sigma_i = \sqrt{\lambda_i}$ and mean

[1]The mean absolute distance is defined as the mean distance between the vertices of the model representation and the vertices of the corresponding case (Styner et al., 2003)

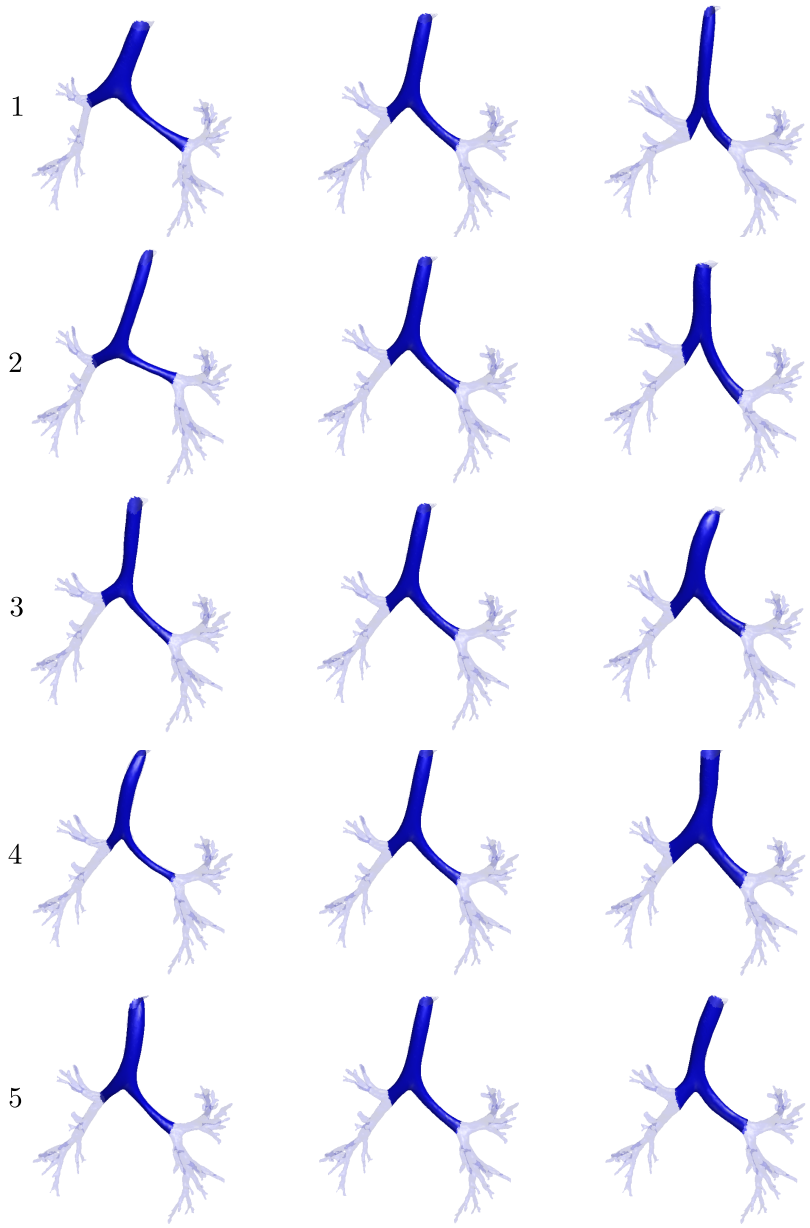

Figure 13: Modes of variation 1-5 for the trachea-LMB-RMB ($b_i = 0$ is the mean model and $b_i = \pm 3\sqrt{\lambda_i}$ represents the shape 3 standard deviations along each mode). Dark regions in the figures show the local PDM model and the light region is an example airway that is registered to the model to aid visualisation.

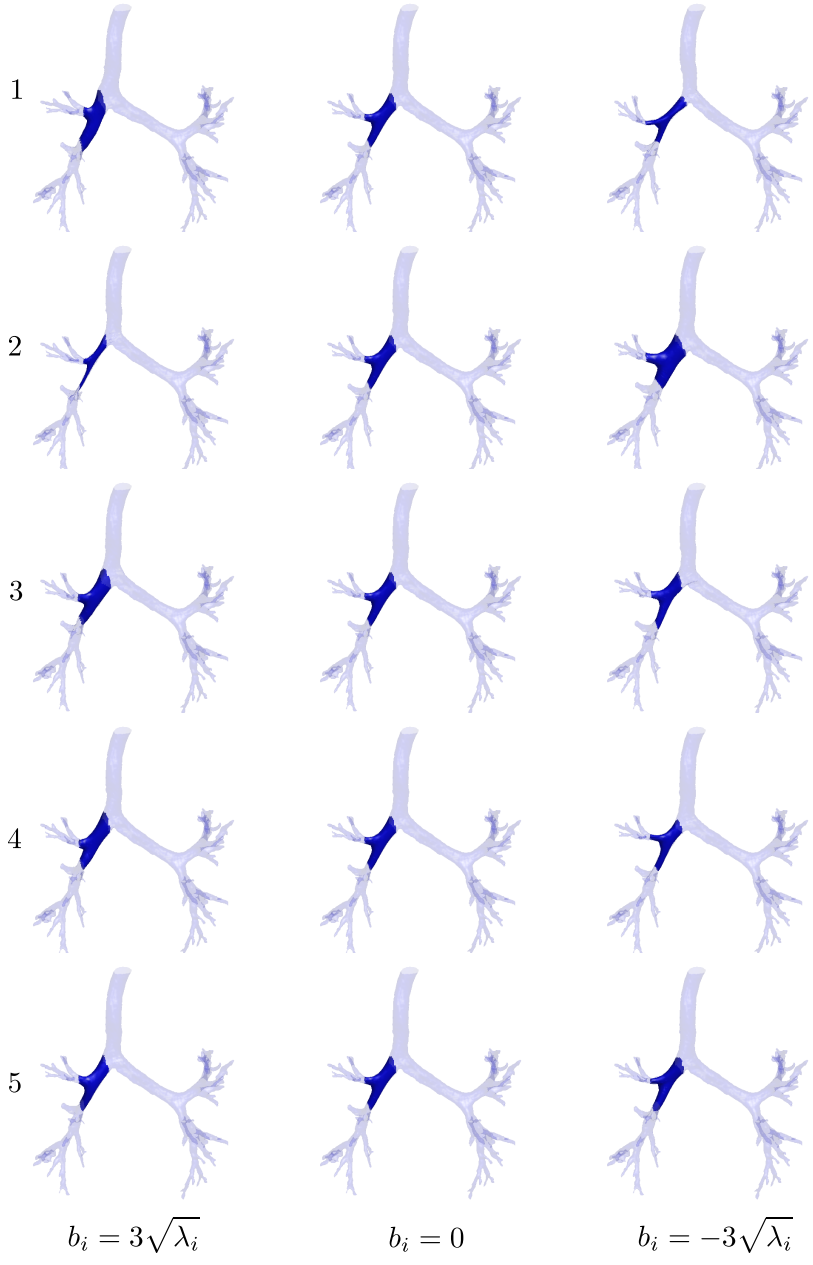

Figure 14: Modes of variation 1-5 of the RMB-RUL-BI ($b_i = 0$ is the mean model and $b_i = \pm 3\sqrt{\lambda_i}$ represents the shape 3 standard deviations along each mode)

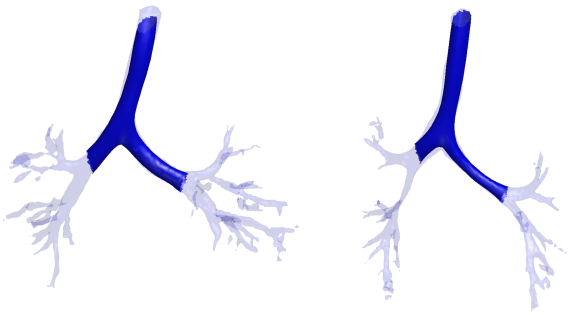

(a) non-TB case (b) TB case

Figure 15: Example airways from the non-TB and TB datasets with overlaid 11 component representation of the first 3 branches

$\mu_i = 0$ (Hu et al., 2010). MAD from each instance to each case in the dataset was calculated, and the minimum distance was recorded, which represents to closeness of the instance to the nearest case in the dataset. The mean and standard deviation of the minimum MAD distances were calculated.

PCA produces modes of variation that are ordered by the contribution to the total variance within the dataset. Figure 16 shows the contribution of each mode to the total variance. In both cases approximately 90% of the airway variation in the dataset is represented by the first 11 modes.

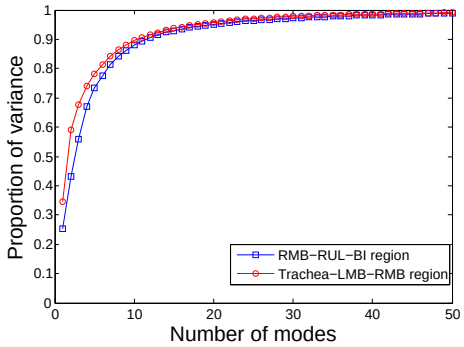

Figure 16: The proportion of the total variance in the dataset that is represented by n principal modes for the Trachea-LMB-RMB and RMB-RUL-BI regions. Principal modes are ordered by variance and, therefore, higher modes will make smaller contributions. In this dataset, the first 11 modes represent $\approx 90\%$ of the variation.

The model generalisation and model specificity for both PDMs are shown in Figure 17 as a function of the number of modes of variation (M). As ex-

pected the model generalisation – the ability of the model to represent unseen cases– performs better (decreases) with increasing number of modes (using 11 modes G(M) ≈ 0.5 mm). Also as expected, the specificity (the ability of the model to only represent similar cases) performs worse (increases) with increasing number of modes. The Trachea-LMB-RMB model outperforms the RMB-RUL-BI model both in terms of generalisation and specificity. This is probably because there is less variability in the latter region.

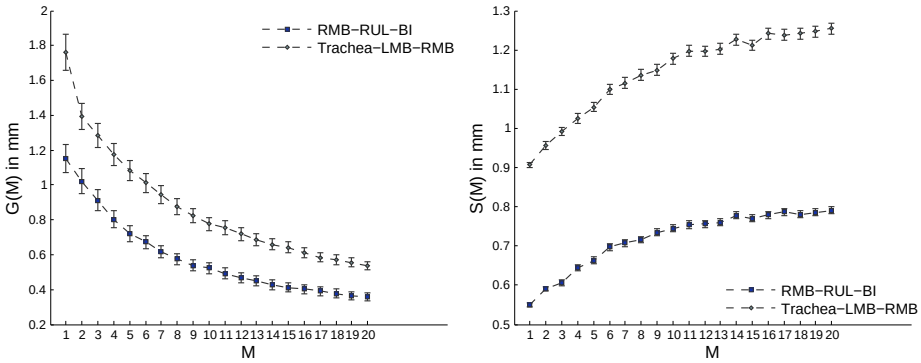

(a) Generalisation of the model (b) Specificity of the model

Figure 17: Generalisation G(M) and specificity S(M) of the model with confidence intervals, as a function of the number of modes (M), for the Trachea-LMB-RMB and RMB-RUL-BI regions.

5. Detection of cases with airway abnormalities

The CAD pipeline that we have developed is able to detect airway deformation related to paediatric TB. The following steps are used for detection of an unseen case: a chest CT of a paediatric patient with suspected TB is acquired, the airway is automatically segmented, the airway centreline and branch points are identified, and landmarks are placed on the airway surface using the methods described in Section 3.1 and 3.2. The airway template mesh for each 3-branch region is then registered to the segmented airway surface to develop vertex to vertex correspondence with the training set (Sections 3.3 and 3.4). The point distribution model – created using the training set – is then used to project the mesh vertices of the test case, and the first 11 modes of variation are extracted as features. A classifier, trained on features from the training set, is then used to classify each 3-branch region of the airways. This allows automatic detection of airway abnormality in each 3-branch region based on variation of the shape of each bronchi as well as variation with respect to neighbouring bronchi.

Evaluation was performed using leave-one-out cross validation (LOOCV) on the test set where the classifier was tested on each case and trained on the remainder over the entire dataset. The receiver-operating-characteristic (ROC) was used to evaluate accuracy. Area-under-the-ROC curve (AUC) provided a single performance measure. Confidence intervals were used to assess the confidence in both the AUC and ROC curves, and were calculated using bootstrapping with 5000 replications. Bootstrapping was performed using the bias corrected and accelerated (BCa) method as this is the preferred approach (Altman et al., 2005).

Before classification, Fisher mapping was used as a useful visualisation of the linear separability between TB and non-TB cases for a set of features. Figure 18 shows that separability of test cases with and without airway involvement can be achieved by 11 principal modes of variation (90% of the dataset variation) using Fisher mapping. Fisher mapping is an affine mapping that maximises the ratio of the inter/intra class variability (van der Heijden et al., 2004).

The key outcome of this method is detection of airway abnormality in unseen cases. Figure 19 shows ROC curves, for the two regions of interest, distinguishing cases with airway involvement from controls, using linear discriminant analysis (LDA) with leave-one-out cross validation on the test set. 95% confidence intervals for the sensitivity and AUC of the ROCs were

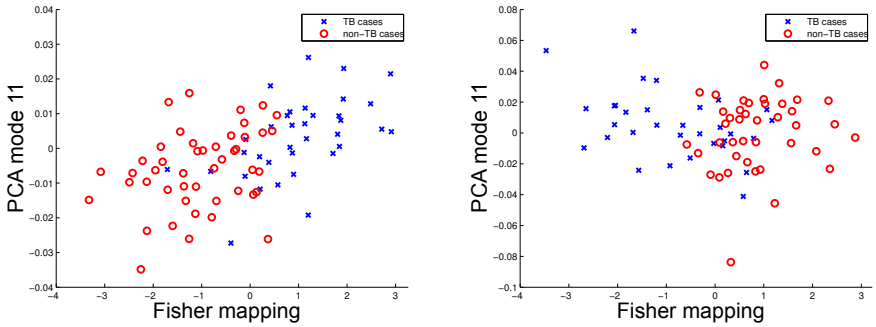

(a) Trachea (b) RMB

Figure 18: Fisher mapping of first 10 modes of dataset versus the 11th mode on the test set

estimated using bootstrap BCa with 1000 replications. The figure shows a comparison between classification using the LA-PDM derived features that we introduce and classification using features derived from mean airway cross-section. Using the LA-PDM achieved an AUC of $0.87(0.77-0.94)$ on the Trachea-RMB-LMB region (Figure 19a). Figure 19b shows that second order bronchi can also play a role in TB detection with an AUC of $0.81(0.68-0.90)$. Using the alternative set of features derived from mean airway cross-section an AUC of $0.72(0.59-0.83)$ was obtained for the first region but classification failed on the second region $0.59(0.46-0.73)$.

The significance of the AUC improvement was calculated using bootstrapping (BCa with 5000 replications) with a two-tailed Monte Carlo technique. Assuming a significance level of $\alpha = 5\%$, the PDM derived features were significantly better than the average diameter derived features for both the Trachea-RMB-LMB (p=0.024) and RMB-RUL-BI regions (p=0.013).

This suggests that branch diameter is not enough for the detection of airway shape variation (due to paediatric TB), and that LA-PDM can considerably improve detection by accounting for variation in regions of the airway tree. The flexibility of the PDM to model variation automatically is also a considerable advantage.

So far we have shown the effectiveness of LA-PDM for the detection of airway involvement in unseen cases. However, analysis of more specific involvement could also be obtained from individual modes of variation and identifying the modes that have the most impact on classification. This

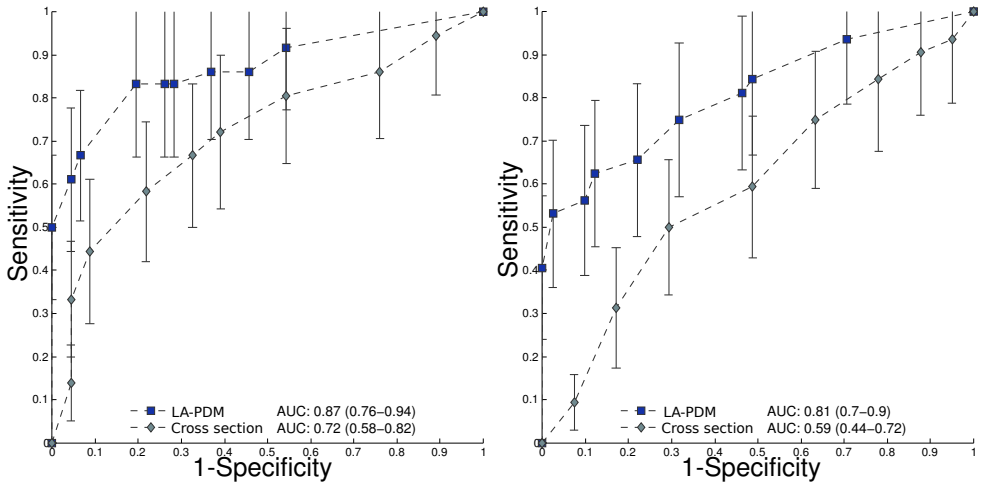

(a) Trachea-RMB-LMB (b) RMB-RUL-BI

Figure 19: ROC curves for accuracy of detection of cases with airway involvement using LA-PDM and mean airway cross section features

method used 11 principal components, representing 90% of the variation for classification. Figure 20 considers the impact of the number of components on the accuracy of detecting airway involvement and shows the key modes of variation that play a role in distinguishing airway involvement from normal variation. The AUC was calculated using LOOCV for each number of features on the validation set.

Interesting to note, is that large improvements in performance are seen when adding certain components while others offer no improvement. The modes that allow airway involvement to be identified include the modes 1, 2, 5 and 10 for the trachea-RMB-LMB region and modes 1, 6, 7, 8, 11 and 14. Other modes may also contribute to the classification but only improve the classification when used in combination with one of the modes described above. Figure 20b shows that the AUC could be improved further if more modes than those contributing to the 90% variance were included. The plots also show that adding certain modes can worsen the performance of the classifier; fitting a classifier to a feature that does not contribute to the distinguishing the datasets has the potential to be penalised by the cross validation.

This section shows that accurate classification of disease with airway involvement can be performed automatically. However, these modes could have

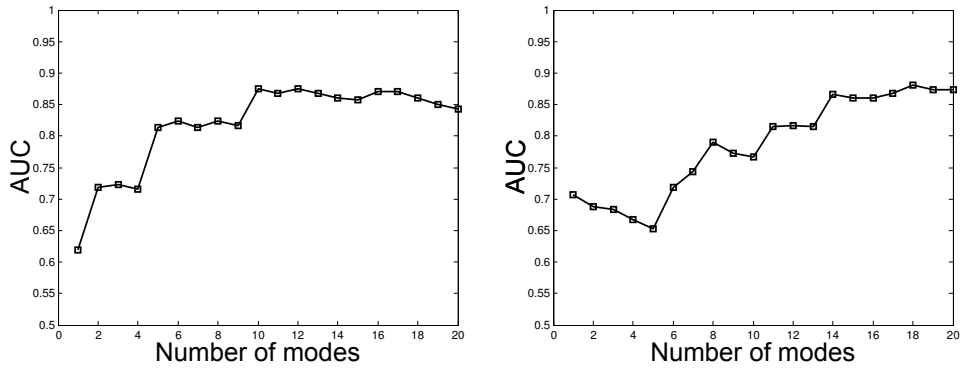

(a) Trachea-RMB-LMB region. Including modes 1, 2, 5 and 10 appears key to improving the accuracy of the classification.

(b) RMB-RUL-BI region. Including modes 1, 6, 7, 8, 11 and 14 appears key to improving the accuracy of the classification.

Figure 20: AUC as a function of number of modes used in the classification.

more use beyond a automatic classification of the airway. By identifying the modes that are most important for detection of airway involvement, these modes could be used to illustrate the specific type of abnormal airway involvement to the viewer. Figure 21 illustrates this point by including the key modes from the two regions along with the parameters of the example airway shown in Figure 15b. The parameters of the example airway are plotted in terms of the standard deviation of each mode. The Figures show that the largest deviations (from the population mean) in the example have clinical importance because they illustrate the types of airway pathology, which, in this example, include narrowing of the bronchus-intermedius (R7) as well as local stenosis of the LMB (T10).

This software was written using a combination of Matlab R2011a and C++. C++ was incorporated to improve the speed at various bottlenecks. The software was evaluated on a system with a 2.80GHz Intel quad-core processor and 6GB of RAM. The mean time for segmentation of each airway was 168 ± 57s. Extraction of the centreline, pruning and branch point detection for each airway took a mean time of 22 ± 18s. Once each airway was segmented and the structure extracted, then mesh generation, registration, training and classification of the entire 179 patient dataset took 1162s on a single run.

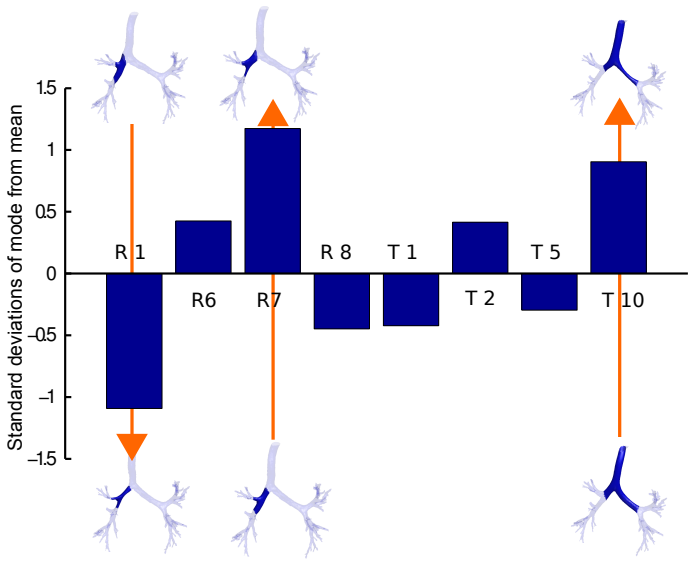

Figure 21: The example airway (Fig. 15b) is shown in terms of the most important modes for classification (from Fig. 20) where R represents the modes from the RMB-RUL-BI and T is trachea-RMB-LMB. R1, R7 and T10 show the largest deviation from the mean airway for this case and, therefore, the types of pathology found in this airway, including narrowing of the RMB and stenosis of the LMB.

6. Discussion

The LA-PDM method introduced in this paper can accurately distinguish between airway involvement (from paediatric pulmonary TB) and normal airways by examining regions of the airway likely to be affected by lymphadenopathy (AUC of $0.87(0.77 - 0.94)$ and $0.81(0.68 - 0.90)$). The LA-PDM derived features show more promise than features derived from the airway diameter, which is probably due to the ability to represent more complex variation in feature space. However, we are not aware of any previous studies that model the effect of lymphadenopathy on airway changes and are, therefore, not able to compare to previous methods. Instead we propose these results as a benchmark for future airway shape analysis [2].

A training set of 89 patients was used during development of the method and determination of landmark point and registration methods. Because

[2]Supplementary data is available at:
http://www.birving.com/data/1/ShapeModel/indexairwayshape.html

the training set was acquired from a number of hospitals, evaluation was performed on the previously unseen 90 patient test set (from a single hospital) using leave-one-out cross validation.

The classification based on shape aims to distinguish pathological shape variation from inter-patient variation, and variation due to age and breathing artefacts. The standard deviation of the patients' age is approximately 3 years. With a larger dataset, it would be possible to divide the dataset into age groups and possibly improve the performance by removing noise caused by age variation. However, studies have shown that the proportions of the airway do not change considerably with age in children (Masters et al., 2006). Breathing artefacts could also add noise to shape-based airway classification as it is not possible to perform a breath hold scan on young patients. The bronchi lengthen and dilate during inspiration but these changes are expected to be distinguishable from TB pathology. The accuracy of the classifier also indicates that breathing variation does not have a considerable effect. With newer 128 and even 256 slice CT, the whole chest volume could be imaged in less than a second, eliminating motion.

These results show that the LA-PDM has promise for detecting airway involvement in chest CT examinations and, thus, speeding up CT assessment. As demonstrated in the results section, the methods can also be used to provide key features of each airway being examined, providing information on the type of deformation from the modes of variation that each airway exhibits. If these key modes of variation were linked to the location or locations of lymphadenopathy then further information could be derived from shape changes.

This method could be used to provide additional automated analysis and visualisation for any patient undergoing a CT examination and might be applied to other diseases affecting the airway. Bronchoscopy is currently the "gold standard" for evaluating airway stenosis and deformation but is an invasive procedure. CT with techniques such as virtual bronchoscopy and virtual rendering, shows potential as an alternative to bronchoscopy. The method presented here has the potential to improve and automate analysis of airway shape deformation. Initial steps to extending the LA-PDM to X-ray examinations for routine X-ray screening are also being assessed (Irving et al., 2013).

Finally, CT scanners are increasingly being deployed, including in areas with a high TB prevalence such as South Africa. These are often not manned by radiologists or by junior trainees due to the shortage of radiologists in de-

veloping countries. Telereading is not always the solution because radiologists from low prevalence regions are not familiar with all the features of TB in children because of the limited number of cases in their countries. So there is a definite role for computer assisted detection in CT, and furthermore for screening in radiography.

7. Acknowledgements

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

8. References

Altman, D., Machin, D., Bryant, T., Gardner, M., 2005. Statistics with confidence. 2nd ed., BMJ Books, London.

Andronikou, S., Irving, B., Hlabangana, L.T., Pillay, T., Taylor, P., Goussard, P., Gie, R., 2013. Technical developments in postprocessing of paediatric airway imaging. Pediatric Radiology 43, 269–284.

Andronikou, S., Joseph, E., Lucas, S., Brachmeyer, S., Toit, G., Zar, H., Swingler, G., 2004. CT scanning for the detection of tuberculous mediastinal and hilar lymphadenopathy in children. Pediatric Radiology 34, 232–236.

Andronikou, S., Wieselthaler, N., 2004. Modern imaging of tuberculosis in children: thoracic, central nervous system and abdominal tuberculosis. Pediatric Radiology 34, 861–875.

Bookstein, F., 1989. Principal warps: Thin-plate splines and the decomposition of deformations. IEEE Transactions on Pattern Analysis and Machine Intelligence 11, 567–585.

Bookstein, F., 1997. Shape and the information in medical images: A decade of the morphometric synthesis. Computer Vision and Image Understanding 66, 97–118.

Cootes, T., Taylor, C., Cooper, D., Graham, J., et al., 1995. Active shape models-their training and application . Computer Vision and Image Understanding 61, 38–59.

Deligianni, F., Chung, A., Yang, G., 2006. Nonrigid 2-D/3-D registration for patient specific bronchoscopy simulation with statistical shape modeling: Phantom validation. IEEE Transactions on Medical Imaging 25, 1462–1471.

Desbrun, M., Meyer, M., Schröder, P., Barr, A., 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. Proc. SIGGRAPH , 317–324.

Fetita, C., Ortner, M., Brillet, P., Hmeidi, Y., Prêteux, F., 2010. Airway shape assessment with visual feed-back in asthma and obstructive diseases, in: Proc. SPIE, pp. 76251E–76251E–12.

Garland, M., Heckbert, P.S., 1997. Surface simplification using quadric error metrics, in: Proceedings of the 24th annual conference on Computer graphics and interactive techniques, ACM Press/Addison-Wesley Publishing Co.. pp. 209–216.

Gie, R., 2003. Diagnostic Atlas of Intrathoracic Tuberculosis in Children: A Guide for Low Income Countries. Paris: International Union against Tuberculosis and Lung Disease.

van Ginneken, B., Frangi, A., Staal, J., ter Haar Romeny, B., Viergever, M., 2002. Active shape model segmentation with optimal features. IEEE Transactions on Medical Imaging 21, 924–933.

Goussard, P., Gie, R., 2007. Airway Involvement in Pulmonary Tuberculosis. South African Medical Journal 97, 986–988.

Goussard, P., Gie, R.P., Kling, S., Andronikou, S., Lucas, S., Janson, J., Roussouw, G., 2013. Bronchoscopic assessment of airway involvement in children presenting with clinically significant airway obstruction due to tuberculosis. Pediatric pulmonology .

van der Heijden, F., Duin, R., de Ridder, D., Tax, D., 2004. Classification, Parameter Estimation and State Estimation: An Engineering Approach Using MATLAB. Wiley, Chichester.

Hu, Y., van den Boom, R., Carter, T., Taylor, Z., Hawkes, D., Ahmed, H.U., Emberton, M., Allen, C., Barratt, D., 2010. A comparison of the accuracy of statistical models of prostate motion trained using data from

biomechanical simulations. Progress in Biophysics and Molecular Biology 103, 262 – 272.

Hutton, T., Buxton, B., Hammond, P., Potts, H., 2003. Estimating average growth trajectories in shape-space using kernel smoothing. IEEE Transactions on Medical Imaging 22, 747–753.

Irving, B., Douglas, T., Taylor, P., 2013. 2d x-ray airway tree segmentation by 3d deformable model projection and registration, in: Proc. of Fifth International Workshop on Pulmonary Image Analysis (Accepted).

Irving, B., Goussard, P., Gie, R., Todd-Pokropek, A., Taylor, P., 2011. Identification of paediatric tuberculosis from airway shape features, in: Proc. Medical Image Computing and Computer Assisted Interventions, pp. 133–140.

Irving, B., Taylor, P., Todd-Pokropek, A., 2009. 3D segmentation of the airway tree using a morphology based method, in: Proc. of Second International Workshop on Pulmonary Image Analysis, pp. 297–307.

Kaus, M., Pekar, V., Lorenz, C., Truyen, R., Lobregt, S., Weese, J., 2003. Automated 3-D PDM construction from segmented images using deformable models. IEEE Transactions on Medical Imaging 22, 1005–1013.

Kiraly, A., Odry, B., Godoy, M., Geiger, B., Novak, C., Naidich, D., 2008. Computer-aided Diagnosis of the Airways: Beyond Nodule Detection. Journal of Thoracic Imaging 23, 105–113.

Lo, P., van Ginneken, B., Reinhardt, J.M., Yavarna, T., de Jong, P.A., Irving, B., Fetita, C., Ortner, M., Pinho, R., Sijbers, J., Feuerstein, M., Fabijanska, A., Bauer, C., Beichel, R., Mendoza, C.S., Wiemker, R., Lee, J., Reeves, A.P., Born, S., Weinheimer, O., van Rikxoort, E.M., Tschirren, J., Mori, K., Odry, B., Naidich, D.P., Hartmann, I., Hoffman, E.A., Prokop, M., Pedersen, J.H., de Bruijne, M., 2012. Extraction of Airways From CT (EXACT'09). IEEE Transactions on Medical Imaging 31, 2093–2107.

Lucas, S., Andronikou, S., Goussard, P., Gie, R., 2012. Ct features of lymphobronchial tuberculosis in children, including complications and associated abnormalities. Pediatric radiology 42, 923–931.

Masters, I., Ware, R., Zimmerman, P., Lovell, B., Wootton, R., Francis, P., Chang, A., 2006. Airway sizes and proportions in children quantified by a video-bronchoscopic technique. BMC Pulm. Med. 6, 5–13.

Meller, S., Kalender, W., 2004. Building a statistical shape model of the pelvis, in: Proc. Computer Assisted Radiology and Surgery, pp. 561–566.

Möller, T., Trumbore, B., 2005. Fast, minimum storage ray/triangle intersection, in: Proc. ACM SIGGRAPH.

Palágyi, K., Tschirren, J., Hoffman, E., Sonka, M., 2006. Quantitative analysis of pulmonary airway tree structures. Computers in Biology and Medicine 36, 974–996.

Paulsen, R., Larsen, R., Nielsen, C., Laugesen, S., Ersboll, B., 2002. Building and testing a statistical shape model of the human ear canal, in: Proc. Medical Image Computing and Computer Assisted Intervention, pp. 373–380.

Petersen, J., Lo, P., Nielsen, M., Edula, G., Ashraf, H., Dirksen, A., De Bruijne, M., 2010. Quantitative analysis of airway abnormalities in ct, in: Proc. SPIE, p. 76241S.

Pinho, R., Tournoy, K., Sijbers, J., 2011. Assessment and stenting of tracheal stenosis using deformable shape models. Medical Image Analysis 15, 250–266.

du Plessis, J., Goussard, P., Andronikou, S., Gie, R., George, R., 2009. Comparing three-dimensional volume-rendered CT images with fibreoptic tracheobronchoscopy in the evaluation of airway compression caused by tuberculous lymphadenopathy in children. Pediatric Radiology 39, 694–702.

Sandgren, A., Cuevas, L., Dara, M., Gie, R., Grzemska, M., Hawkridge, A., Hesseling, A., Kampmann, B., Lienhardt, C., Manissero, D., et al., 2012. Childhood tuberculosis: progress requires advocacy strategy now. European Respiratory Journal 40, 294–297.

Schaaf, H., Beyers, N., Gie, R., Nel, E., Smuts, N., Scøtt, F., Donald, P., Fourie, P., et al., 1995. Respiratory tuberculosis in childhood: the diagnostic value of clinical features and special investigations. Pediatr. Infect. Dis. J. 14, 189–194.

Schlkopf, B., Smola, A., Mller, K., 1997. Kernel principal component analysis, in: Artificial Neural Networks - ICANN. volume 1327, pp. 583–588.

Shi, H., Scarfe, W.C., Farman, A.G., 2006. Upper airway segmentation and dimensions estimation from cone-beam CT image datasets. Int. J. Comput. Assist. Radiol. Surg. 1, 177–186.

Styner, M., Rajamani, K., Nolte, L., Zsemlye, G., Szkely, G., Taylor, C., Davies, R., 2003. Evaluation of 3d correspondence methods for model building, in: Taylor, C., Noble, J. (Eds.), Information Processing in Medical Imaging. Springer Berlin / Heidelberg. volume 2732 of *Lecture Notes in Computer Science*, pp. 63–75.

Theart, A., Marais, B., Gie, R., Hesseling, A., Beyers, N., 2005. Criteria used for the diagnosis of childhood tuberculosis at primary health care level in a high-burden, urban setting. International Journal of Tuberculosis and Lung Disease 9, 1210–1214.

Tschirren, J., Hoffman, E., McLennan, G., Sonka, M., 2005. Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans. IEEE Transactions on Medical Imaging 24, 1529–1539.

WHO, 2012. Summary for TB-HIV estimates for 2005 by WHO regions. `www.who.int/tb`. (Accessed: July 2012).

Wiemker, R., Blaffert, T., Blow, T., Renisch, S., Lorenz, C., 2004. Automated assessment of bronchial lumen, wall thickness and bronchoarterial diameter ratio of the tracheobronchial tree using high-resolution ct, in: Proc. Computer Assisted Radiology and Surgery, pp. 967 – 972.