

Robust Segmentation of Challenging Lungs in CT using Multi-Stage Learning and Level Set Optimization

Neil Birkbeck¹, Michal Sofka¹, Timo Kohlberger¹, Jingdan Zhang¹, Jens Wetzl¹, Jens Kaftan², and S. Kevin Zhou¹

Abstract

Automatic segmentation of lung tissue in thoracic CT scans is useful for diagnosis and treatment planning of pulmonary diseases. Unlike healthy lung tissue that is easily identifiable in CT scans, diseased lung parenchyma is hard to segment automatically due to its higher attenuation, inhomogeneous appearance, and inconsistent texture. We overcome these challenges through a multi-layer machine learning approach that exploits geometric structures both within and outside the lung (e.g., ribs, spine). In the coarsest layer, a set of stable landmarks on the surface of the lung are detected through a hierarchical detection network (HDN) that is trained on hundreds of annotated CT volumes. These landmarks are used to robustly initialize a coarse statistical model of the lung shape. Subsequently, a region-dependent boundary refinement uses a discriminative appearance classifier to refine the surface, and finally a region-driven level set refinement is used to extract the fine scale detail. Through this approach we demonstrate robustness to a variety of lung pathologies.

1 Introduction

Thoracic CT images are clinically used for screening, diagnosis, and treatment planning of lung diseases [16, 14, 13]. Computer aided diagnosis (CAD) tools built for CT imaging rely on a segmentation of the lung as the first step [24]. For example, algorithms for detection of malignant nodules [15, 1]

¹ Siemens Corporate Research, 755 College Road East, Plainsboro NJ, 08536, USA

² Siemens Molecular Imaging, 23-38 Hythe Bridge Street, Oxford, OX1 2EP, UK
e-mail: neil.birkbeck@siemens.com, michal.sofka@siemens.com,
 timo.kohlberger@siemens.com, jingdan.zhang@siemens.com,
 jens.kaftan@siemens.com, shaohua.zhou@siemens.com




Fig. 1 Examples of how diseases, such as tumors or interstitial lung disease, affect the appearance of lung in CT.

and classification of lung tissue into various diseases [34, 40, 37, 38] restrict processing to within the lung regions only.

Further semantic decomposition of the lung tissue into lobes, e.g., [41, 35], and methods to visualize dynamic motion of lung function through respiratory cycles [8] also rely on a segmentation of the lung. As lung segmentation is an important prerequisite for these applications, it is important for it to be robust to abnormalities. Failures in the segmentation algorithm propagate to failures in further processing. In CAD applications, these failures mean a clinician may not be warned about a potential cancerous tumor [1].

Automatically segmenting healthy lung in CT is relatively easy due to low attenuation of lungs compared to surrounding tissue. Pathologies, such as interstitial lung disease (ILD), pleural effusion, and tumors, on the other hand, significantly change the shape, appearance, or texture of lung parenchyma (Fig. 1). In a clinical setting, pathologies are often no longer an exception but rather the norm, and methods robust to these variations are necessary.

In this article, we address the automatic segmentation of challenging cases described above and illustrated in Fig. 1. We first review the basic algorithms for segmenting healthy lung parenchyma and demonstrate the limitations of these algorithms in the presence of pathologies (§2). We then discuss how texture cues, anatomical information, and statistical shape models, have been leveraged to improve robustness in the presence of disease. This is followed by the presentation of a method that utilizes machine learning with both texture cues and anatomical information from outside the lung to obtain an robust lung segmentation that is further improved with level set refinement (§3). These algorithms are implemented in a modular software framework that promotes reuse and experimentation (§4).

2 Segmenting Healthy Lungs

As healthy lung parenchyma has lower attenuation than the surrounding tissue, simple image processing techniques often achieve good results. The at-




Fig. 2 Healthy lung tissue is easily separable from body tissue with a single threshold.




Fig. 3 Segmentation by simple thresholding often excludes airways and vessels (left). These vessels are easily removed by morphological operations or by removing small connected components on 2D slices (right).

tenuation of healthy lung tissue in CT varies across the lung, depends on the phase of respiratory cycle [25] and image acquisition settings [2], with mean values being around -860 to -700 HU [25]. However, as illustrated in Fig. 2, the intensity distribution within the lung region is often completely disjoint from that of the higher density body tissue. Thus simple image processing operations, such as region growing, edge tracking [10], or simple thresholding [11, 23, 4], can be used to separate lung from body tissue.

For example, thresholding methods first separate body tissue from surrounding background tissue by removing large connected components touching the edges of the images. Lung tissue can then be separated from the body tissue with a threshold. Hu *et al.* propose to use a dynamic threshold, τ , that is iteratively updated to be the mean of the average intensities between the lung and body, $\tau_t = (\mu_{\text{lung}}^{t-1} + \mu_{\text{body}}^{t-1})/2$, where μ_{lung}^t and μ_{body}^t are the mean HU values of the lying below and above τ_t respectively [11]. The initial threshold, τ_0 , is estimated with $\mu_{\text{lung}}^0 = -1000$ and μ_{body}^0 is the average intensity of all HU values greater than 0.

Depending on the application, the final step is to separate the segmented lung tissue into left and right lungs. The thresholded lung regions are con-




Fig. 4 Small juxtapleural nodules like those in the left figure can be included in the final segmentation through the use of morphological operations, but finding a set of parameters to include all possible tumors, such as those in the cases on the right, is challenging.

nected by the trachea and large airways, and can be separated by region growing a trachea segmentation from the top of the segment [11, 39]. Further, the lungs often touch at the anterior and posterior junctions and need to be separated using another post-process that uses a maximum cut path to split the single lung component into the individual lungs [26, 4, 11].

The simple thresholding scheme gives a good gross estimate to the lung volume but often excludes small airways and vessels that a clinician would include with the lung annotation. These small regions can be filled through the use of morphological operations [11, 36], or by performing connected component analysis on the image slices [39] (see Fig. 3).

Another subtle problem comes with nodules that lie on the boundary of the lung [27]. A complete segmentation of the lung is essential for cancer screening applications [3], and studies on computer aided diagnosis have found the exclusion of such nodules to be a limitation of automated segmentation and nodule detection methods [1]. These nodules can be included in the segmentation through the use of special post-processing steps, such as the adaptive border marching algorithm [23], which tries to include small concave regions on the lung slices. However, such algorithmic approaches are bound to fail for larger tumors whose size and shape are unconstrained (Fig. 4).

Even more problematic is interstitial lung disease, which causes a dramatic change in the attenuation and local texture pattern of lung parenchyma. Unlike the clearly separated histograms of healthy lung (Fig. 3), such diseased tissue will often have attenuation values that overlap with the surrounding body tissue (Fig. 5). Although higher thresholds, such as -300 HU [39], can be used to include more lung tissue, simple thresholding methods are incapable of robustly segmenting these challenging pathologies.

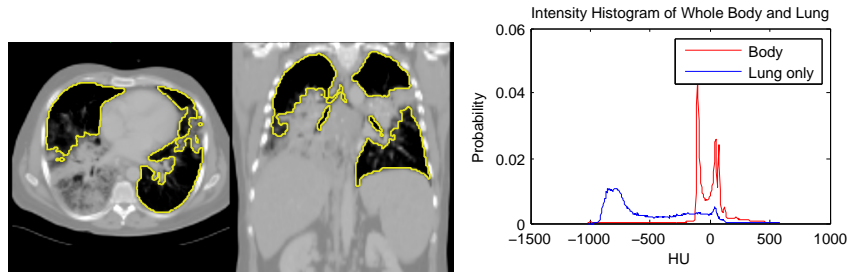


Fig. 5 Lungs with diffuse pulmonary disease has higher density tissue and is difficult to separate with a simple thresholding.

2.1 Cues for segmenting pathological lung

Intensity alone is the single strongest cue for segmenting healthy lung tissue, but in order to address the high density lung tissue associated with interstitial lung disease, and the different shapes associated with varying tumors, cues based on texture, anatomy, or shape priors must be exploited.

2.1.1 Texture Cues

Diseased parenchyma has a different texture pattern which can often be extracted through the use of texture features [20, 39, 37]. Texture, the local structural pattern of intensities, is commonly characterized by measurements obtained from a co-occurrence matrix [9], which records the joint frequency of intensity values between two pixels separated by a fixed offset computed over small volumes of interest around each image voxel. Quantities derived from this co-occurrence matrix, such as entropy, can be used to directly identify diseased tissue [39], or statistical classifiers can be trained to distinguish healthy from pathological tissue using features derived from this matrix [37].

2.1.2 Anatomical Cues

As the appearance of the pathological lung parenchyma can vary dramatically, texture and intensity cues are incapable of capturing all pathologies. These quantities, which are internal to the lung, can be combined with neighboring anatomical context, such as distance and curvature of the ribs. For example, Prasad *et al.*, use the curvature of the ribs in order to adaptively choose a threshold for the lung segmentation [22]. As the lung border should also lie close to the ribcage and spine, distance to these anatomical structures can be combined with intensity features in classification of the lung border [12].

2.1.3 Shape Modeling

In order to avoid segmentations with unlikely shapes, such as the one in Fig. 5, a shape prior can be used as an additional cue for segmenting challenging cases. Although some inconsistencies in the shape can be removed, for example, by post-processing operations [11] or through ensuring smoothness in the resulting segmentations [12], the result should be constrained to be lung-like. Explicit statistical models of shape variability, such as a point distribution model that models the shape variability with a low dimensional linear subspace can be used to constrain the resulting segmentation [6]. These models have been effectively used to overcome pathologies [32, 30].

Segmentation-by-registration are another class of methods that enforce a prior constraint on the resulting segmentations. In such approaches, a reference image complete with annotations is aligned to the target image through the process of image alignment. The target segmentation is then derived from the aligned reference image. To increase generality, multiple reference images can be aligned to the target image and the final segmentation can be taken as the fused result. Depending on the matching score used in the registration, such methods have shown to be effective for pathological lungs [28].

3 Multi-Stage Learning for Lung Segmentation

To ensure robustness, a complete solution to the segmentation of pathological lung in CT has to include components that address shape and appearance variability caused by both tumors and diffuse lung disease. In this section we introduce a robust machine learning method for segmenting challenging lung cases that uses external anatomical information to position a statistical mesh model on the image. A large database of annotated images is used to train discriminative classifiers to detect initial poses of the mesh models and to identify stable boundary landmarks [30]. The boundary of this initialization is then guided by texture and appearance patterns to identify the boundary of the lung. Finally, as pathologies often only occupy a small portion of the lung, a fine-scale level set surface refinement is used to ensure the final segmentation includes all healthy regions as well as preserving the segmentation in the pathological regions [18]. The elements of the system are illustrated in Fig. 6.

3.1 Learning-based Robust Initialization

For the initialization phase, the carina of the trachea is detected in the input volume. Given the carina location, a Hierarchical Detection Network (HDN) [31], is used to detect an initial set of pose parameters of statistical

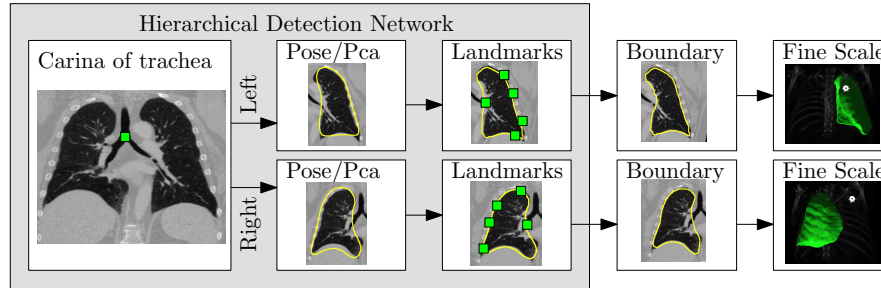


Fig. 6 A hierarchical detection network is used to initialize the lung segmentations from a set of stable landmarks. These segmentations are further refined by trained boundary detectors and finally by a fine-scale level set refinement.




Fig. 7 The variability of right lung shape model by using one of the first three basis vectors to represent the shape.

shape models of each lung (§3.1.2 & §3.1.3). In the next level of the hierarchy, stable landmark points on the mesh surface are refined to get an improved lung segmentation that is less sensitive to appearance changes within pathologies (§3.1.4). As we will see, these landmark points tend to use nearby anatomical context, e.g., ribs, to improve the segmentation. A further refinement comes by displacing the mesh surface so as to maximize the score of an appearance-based classifier (§3.2).

3.1.1 Input Annotation Database

A set of CT images complete with manual annotation of the lungs is used as input to the algorithm. In an offline training step, statistical shape models of the shape variation within the database annotations are learned.

The input annotation meshes are first brought into correspondence and remeshed so that each annotation mesh has the same number of points and a consistent triangulation. Each annotation mesh, $\mathcal{M}_k = (\mathcal{P}_k, \mathcal{T})$, then consists of a set of points, $\mathcal{P}_k = \{\mathbf{v}_{ki} \in \mathbb{R}^3\}_{i=1}^N$, and a single set of triangle indices, $\mathcal{T} = \{\Delta_j \in \mathbb{Z}^3\}_{j=1}^{N_{tri}}$.

The variability in the point coordinates is modeled through a low-dimensional linear basis giving a statistical shape model,

$$\mathcal{S} = (\{\hat{\mathbf{v}}\}_{i=1}^N, \{U_j\}_{j=1}^M), \quad (1)$$

that consists of a mean shape, $\{\hat{\mathbf{v}}\}_{i=1}^N$, and the linear basis shapes, $U_j = \{\mathbf{u}_{ij}\}_{i=1}^N$. The linear basis is empirically estimated by performing PCA on the aligned input annotation meshes. The alignment uses procrustes analysis to remove translation, orientation, and scale variation of the corresponding meshes.

A mesh in the *span* of the basis can be approximated by modulating the basis shape vectors and applying a similarity transform

$$g(\mathbf{v}_i; \{\lambda_j\}, \mathbf{p}, \mathbf{r}, \mathbf{s}) = \mathbf{p} + \mathbf{M}(\mathbf{s}, \mathbf{r}) \sum_j (\hat{\mathbf{v}}_i + \mathbf{u}_{ij} \lambda_j), \quad (2)$$

where $\mathbf{M}(\mathbf{s}, \mathbf{r})$ is the rotation and scale matrix parameterized by rotation angles, \mathbf{r} , and anisotropic scale, \mathbf{s} . \mathbf{p} is a translation vector, and $\{\lambda_j\}$ are the shape coefficients. Figure 7 illustrates the variability encoded by the first three basis vectors for a model of the right lung.

Each training shape can be approximated in the same manner, meaning each training shape has an associated pose and shape coefficient vector. In the following description of the image-based detection procedure, the relationship between these pose and shape coefficients and the image features is modeled with machine learning so that the parameters of the shape model can be inferred on unseen data.

3.1.2 Hierarchical detection network

The Hierarchical Detection Network (HDN) uses an efficient sequential decision process to estimate the unknown states (e.g., object poses or landmark positions) of a sequence of objects that depend on each other [31]. In the case of lung segmentation, the unknown states are the poses and shape coefficients that align a statistical model of the lung to the image, which depend on the detected trachea landmark point. The location of stable lung boundary points are dependent on the pose parameters of the lungs.

The HDN infers multiple dependent object states sequentially using a model of the prior relationship between these objects. Let θ_t denote the unknown state of an object (e.g., the 9 parameters of a similarity transform or the 3D coordinates of a landmark), and let the complete state for $t + 1$ objects, $\theta_0, \theta_1, \dots, \theta_t$, be denoted as $\theta_{0:t}$. Given a d dimensional input volume, $V: \mathbb{R}^d \mapsto \mathbb{R}$, the estimation for each object, t , uses an observation region $V_t \subseteq V$. The complete state is inferred from the input volume by maximizing the posterior density, $f(\theta_{0:t} | V_{0:t})$, which is recursively decomposed into

a product of individual likelihoods, $f(V_t|\theta_t)$, and the transition probability between the objects, $f(\theta_t|\theta_{0:t-1})$.

The recursive decomposition of the posterior is derived by applying a sequence of *prediction* and *update* steps. For object t , the *prediction* step ignores the observation region, V_t , and approximates the posterior, $f(\theta_{0:t}|V_{0:t-1})$, using a product of the transition probability, $f(\theta_t|\theta_{0:t-1})$, and the posterior of the preceding objects:

$$f(\theta_{0:t}|V_{0:t-1}) = f(\theta_t|\theta_{0:t-1}, V_{0:t-1})f(\theta_{0:t-1}|V_{0:t-1}), \quad (3)$$

$$= f(\theta_t|\theta_{0:t-1})f(\theta_{0:t-1}|V_{0:t-1}), \quad (4)$$

as θ_t and $V_{0:t-1}$ are assumed to be conditionally independent given $\theta_{0:t-1}$.

The observation region, V_t , is then combined with the prediction in the *update* step,

$$f(\theta_{0:t}|V_{0:t}) = \frac{f(V_{0:t}|\theta_{0:t})f(\theta_{0:t})}{f(V_{0:t})} = \frac{f(V_t|V_{0:t-1}, \theta_{0:t})f(V_{0:t-1}|\theta_{0:t})f(\theta_{0:t})}{f(V_t|V_{0:t-1})f(V_{0:t-1})} \quad (5)$$

$$= \frac{f(V_t|\theta_t)f(\theta_{0:t}|V_{0:t-1})}{f(V_t|V_{0:t-1})}, \quad (6)$$

where the denominator is a normalizing term, and the derivation assumes V_t and $(V_{0:t-1}, \theta_{0:t-1})$ are conditionally independent given θ_t . The likelihood term, $f(V_t|\theta_t)$, is modeled with a discriminative classifier,

$$f(V_t|\theta_t) = f(y = +1|V_t, \theta_t), \quad (7)$$

where the random variable $y \in \{-1, 1\}$ denotes the occurrence of the t^{th} object at pose θ_t . The posterior, $f(y = +1|V_t, \theta_t)$, is modeled with a powerful tree-based classifier, the Probabilistic Boosting Tree (PBT) [33], which is trained to predict the label, y , given the observation region, V_t and a state, θ_t .

The prediction step models the dependence between the objects with the transition prior. In the case of lung segmentation, each object is dependent on one of the previous objects, meaning

$$f(\theta_t|\theta_{0:t-1}) = f(\theta_t|\theta_j), \quad \exists j \in \{0, 1, 2, \dots, t-1\}. \quad (8)$$

The relation between the objects, $f(\theta_t|\theta_j)$, is modeled using a Gaussian distribution whose parameters are learned from training data. In our application, the poses of each lung are dependent on the initial carina landmark point (Fig. 6). The stable landmarks, which are distributed on the boundary of the lungs, are then dependent on the pose of the lung.

The full posterior, $f(\theta_{0:t}|V_{0:t})$, is approximated through sequential importance sampling, where a set of weighted particles (or samples), $\{\theta_t^j, w_t^j\}_{j=1}^P$, is used to approximate the posterior at state t , and these particles are propagated, in sequence, to the next object (see [31] for more details).




Fig. 8 An illustration showing the spatial filtering process used when selecting stable landmarks. The landmark with lowest error is selected on the bottom left of the lung, then a region within that landmark is ignored for future selection (triangles marked white). The process continues until no more vertices can be selected. Here we used a large filter parameter of $r = 100\text{mm}$ for illustration purposes.

3.1.3 Pose detection

The pose of each lung is represented with a 9 dimensional state vector, $\theta_t = \{\mathbf{p}_t, \mathbf{r}_t, \mathbf{s}_t\}$, containing a translation vector $\mathbf{p} \in \mathbb{R}^3$, Euler angles for orientation \mathbf{r} , and an anisotropic scale, \mathbf{s} . The 9 pose parameters are decomposed into three sequential estimates of the respective 3 dimensional quantities [42, 21],

$$f(\theta_t|V_t) = f(\mathbf{p}_t|V_t)f(\mathbf{r}_t|\mathbf{p}_t, V_t)f(\mathbf{s}_t|\mathbf{p}_t, \mathbf{r}_t, V_t), \quad (9)$$

allowing for efficient estimation by using fewer particles during inference. For position, the classifier in (7) is trained using 3D Haar features, whereas orientation and scale are estimated using steerable features [42].

The pose detection aligns the mean shape to the image, giving an initial segmentation. Additional detail is incorporated into the segmentation by then detecting the first three shape coefficients, $\theta_t = \{\lambda_1, \lambda_2, \lambda_3\}$ from (2), using the HDN.

A uniform sampling of shape coefficients, $\lambda_k^{\min} \leq \lambda_k \leq \lambda_k^{\max}$, is used during detection. The bounds, $[\lambda_k^{\min}, \lambda_k^{\max}]$, are estimated from the shape coefficients of the training data. These uniformly sampled coefficients are augmented with the results of the pose estimation to get a candidate set of particles. Steerable features extracted from the image volume at vertices of the mesh are used to train the classifier in (7).

3.1.4 Refinement Using Stable Landmarks

The alignment of the shape models to the input volume through the above described detection process is still sensitive to the appearance and shape changes of pathologies. For this reason, we utilize a robust improvement of the shape alignment that detects a set of stable landmark points on the surface

of the lung [30]. This subset of landmark points are chosen as a stable subset of mesh vertices by using the reliability of their corresponding detectors.

Training & Identifying Stable Landmarks During training, all the input annotations are brought into correspondence using registration in order to build the statistical shape model (§3.1.1). Each one of the mesh vertices is then a potential landmark. Let \mathbf{p}_{ij} denote the location of landmark, i , in training volume j . A position detector is trained for each landmark independently, and the detected result for each detector gives a predicted landmark location: \mathbf{d}_{ij} . The covariance of the error for each detector is then computed as $\mathbf{C}_i = \sum_j (\mathbf{d}_{ij} - \mathbf{p}_{ij})(\mathbf{d}_{ij} - \mathbf{p}_{ij})^\top$.

Landmarks that are stable should have low uncertainty, meaning that the score

$$s_i = \text{trace}(\mathbf{C}_i), \quad (10)$$

should be low. Further, the selected landmarks should be spatially distributed along the surface of the object. In order to select a subset of stable and well distributed landmarks, a greedy selection scheme with a spatial filtering of radius r is used. First, the landmark with the lowest s_i is greedily chosen. Then the mesh vertices within radius r of this vertex on the mean mesh are no longer considered. The process then continues by iteratively choosing the next best landmark and removing nearby vertices from consideration (see Fig. 8). Let $L \subset \{1, 2, \dots, N\}$ denote the mesh indices of the selected landmarks.

The parameter r gives a trade-off between the spatial distribution of the vertices and the stability of the chosen landmarks. In our experiments we found that a radius of $r = 20\text{mm}$ gives a good balance.

Alignment of Shape to Landmarks During detection, in the HDN framework, the stable landmark positions are first predicted from the pose detection in §3.1.3. The resulting detected position of the landmark gives a constraint on where the corresponding vertex should move. As not all vertices have corresponding stable landmarks, the entire statistical mesh model must then be deformed to fit the detected landmark positions. This is accomplished by finding a smoothly varying non-rigid transformation using a thin-plate spline that transforms the input landmark vertices from their pose estimate \mathbf{d}_i^0 (§3.1.3) to their detected positions, \mathbf{d}_i . The transform takes on the following form

$$f(\mathbf{x}) = \sum_{i \in L} \mathbf{a}_i r_i^2 \log(r_i) + \mathbf{A}\mathbf{x}, \quad (11)$$

where $r_i = |\mathbf{x} - \mathbf{d}_i^0|$ is the distance between a point $\mathbf{x} \in \mathbb{R}^3$ and the pre-image of the landmark point. \mathbf{A} is a 3×4 affine transformation matrix, and \mathbf{a}_i are 3×1 vector weights for each landmark point. The unknowns, \mathbf{A} and \mathbf{a}_i , are obtained by solving the linear system that arises when constraining the transformation such that $f(\mathbf{d}_i^0) = \mathbf{d}_i$.

3.2 Freeform refinement

The estimation of the stable landmarks in the previous section provide a robust initial estimate of the lung surfaces. Although the landmarks provide a more detailed initialization than the first three shape coefficients from the preceding stage, a more detailed estimate is obtained through the use of a freeform surface deformation approach [21] that is similar to the methods used in the active shape model [6].

Given the input volume, V , the freeform surface deformation finds the most likely mesh, \mathcal{M} , within the linear span of the statistical shape model:

$$\max f(\mathcal{M}|V) \quad s.t. \quad \mathcal{M} \in \text{span}(\mathcal{S}). \quad (12)$$

The above posterior is computed as an average on the surface of the mesh, and at each surface point, the posterior is evaluated a discriminative classifier:

$$f(\mathcal{M}|V) = \frac{1}{N} \sum_{v_i} f(\mathbf{v}_i|V) = \frac{1}{N} \sum_{v_i} f(y_i = 1 | \mathbf{v}_i, \mathbf{n}_i, V), \quad (13)$$

where the random variable $y_i = \{-1, +1\}$ takes a value of $+1$ if the surface boundary exists at point \mathbf{v}_i with normal \mathbf{n}_i in the input volume.

The statistical classifier for the boundary, $f(y_i = 1 | \mathbf{v}_i, \mathbf{n}_i, V)$, can take into account such things as raw image intensity, spatial texture, or distance to anatomical structures in order to discriminate between surface points either on or off the lung boundary. We use a classifier that automatically selects the best set of features [33]. If only healthy cases exist, the classifier will pick features like image gradient or raw intensity. However, robustness to pathological cases can be obtained by ensuring pathological cases exist in training.

Instead of maximizing (12) directly using the shape coefficients and pose of the shape model, the freeform refinement performs a sequence of local per-vertex optimizations followed by a projection of the shape back into the span of the shape space. The local optimization for a vertex, \mathbf{v}_i , searches for the best displacement, $d_i \in [-\tau, \tau]$, along the normal, \mathbf{n}_i by maximizing the score of the classifier:

$$d_i = \underset{-\tau \leq d \leq \tau}{\operatorname{argmax}} f(\mathbf{v}_i + d\mathbf{n}_i|V), \quad (14)$$

where τ defines a local search range around the point. The vertex position is then updated with the best displacement, $\mathbf{v}_i \leftarrow \mathbf{v}_i + d_i\mathbf{n}_i$. As this process is done independently for each vertex, the resulting surface may be irregular. The projection of the displaced shape into the shape space regularizes the result and ensures a valid shape.

Several iterations of the above local optimization and regularization steps are performed, with the search range, τ , being reduced at each iteration. In the final iterations, when the search range is small, the projection of the

mesh into the shape space is replaced with mesh smoothing [21]. This form of regularization gives a more detailed shape by allowing the surface to deviate from $\text{span}(\mathcal{S})$.

3.3 Fine-scale Refinement

The detection results from the previous section ensure the lung surface encompasses tumors and pathological regions, but the explicit mesh representation is ineffective to represent fine-scale detail. In the final phase, we use an energy-based level set refinement. In addition to the common data and regularization terms, our energy includes a term to remove overlap between lungs (e.g., at the anterior/posterior junctions), and another term to keep the refined solution close to the detection meshes that are output from the previous phase (§3.2). This energy framework can easily incorporate constraints from other adjacent structures such as the heart or liver to further improve the segmentation accuracy [18].

The detected lung surfaces, \mathcal{C}_i , are first converted to signed distance functions, $\Phi_i(\mathbf{x}) : \Omega \subset \mathbb{R}^3 \mapsto \mathbb{R}$ where $|\nabla \Phi_i(\mathbf{x})| = 1$ and $\mathcal{C}_i = \{\mathbf{x} | \Phi_i(\mathbf{x}) = 0\}$. Further, $\Phi_i(\mathbf{x}) > 0$ if \mathbf{x} is inside \mathcal{C}_i , and $\Phi_i(\mathbf{x}) < 0$ outside \mathcal{C}_i (see [5] and references within for more details). Then for each organ we minimize the energy proposed by Kohlberger *et al.* [18]:

$$E(\Phi_i; \{\Phi_j\}_{j \neq i}) := \underbrace{E_d(\Phi_i)}_{\text{data}} + \underbrace{E_s(\Phi_i)}_{\text{smooth}} + \underbrace{\sum_{j \in \mathcal{N}_j^o} E_o(\Phi_i, \Phi_j)}_{\text{overlap}} + \underbrace{E_p(\Phi_i, \Phi_i^0)}_{\text{prior}}. \quad (15)$$

The standard region-based data energy, E_d , measures deviation of the appearance of the inside and outside regions from their respective distributions [5],

$$E_d(\Phi_i) := -\alpha \int_{\Omega} H(\Phi_i) \log(p_i^{\text{in}}(I(\mathbf{x}) | \Phi_i)) + (1 - H(\Phi_i)) \log(p_i^{\text{out}}(I(\mathbf{x}) | \Phi_i)) d\mathbf{x}, \quad (16)$$

where H denotes the Heaviside function. Here, the probability of intensities belonging to the inside or outside of the object, $p_i^{\text{in/out}}$, are modeled with non-parametric Parzen densities (see [7]). This data energy will push the segmentation surface to occupy regions that look similar to its intensity distribution. On the other hand, when the intensity distributions between foreground (lung) and background (body) overlap, this term will have less influence. This happens for diseased cases with large tumors or diffuse lung disease cause p_i^{in} to contain high density tissue.

The smoothness term in (15) penalizes surface area of the desired segmentation,

$$E_s(\Phi_i) := \gamma \int_{\Omega} |\nabla H(\Phi_i)| d\mathbf{x}, \quad (17)$$




Fig. 9 An illustration of the spatially varying exterior surface constraint weights for the left and right lung on the surface of the detected points, $\{\mathbf{p}_j\}$. The dark regions map to $\omega^{\text{out}} = 0.1$ and the white regions map to $\omega^{\text{out}} = 5$. Larger weights are given in a region near the main bronchi.

and is balanced with the weight coefficient γ .

The overlap term, $E_o(\Phi_i, \Phi_j)$, penalizes overlap between organ i and a set of neighboring organs, N_i^o :

$$E_o(\Phi_i, \Phi_j) := \beta \int_{\Omega} H(\Phi_i) H(\Phi_j) \Phi_j d\mathbf{x}. \quad (18)$$

As intended, this term only affects regions that lie within the intersection of both surfaces, and the constraint is proportional to the depth of penetration into the surface Φ_j .

Finally, the prior shape term prefers the surface to be close to the detected shape, Φ_i^0 :

$$E_p(\Phi_i, \Phi_i^0) := \int_{\Omega} \underbrace{\omega_{l(\mathbf{x})}^{\text{in}} H(\Phi_i^0(\mathbf{x}) - \Phi_i(\mathbf{x}))}_{\text{Inside deviations}} + \underbrace{\omega_{l(\mathbf{x})}^{\text{out}} H(\Phi_i(\mathbf{x}) - \Phi_i^0(\mathbf{x}))}_{\text{Outside deviations}} d\mathbf{x}. \quad (19)$$

The prior term utilizes a set of positive valued spatially varying surface weights $\omega_{l(\mathbf{x})}^{\text{in}}$ and $\omega_{l(\mathbf{x})}^{\text{out}}$ to control the penalization of the surface deviations occurring either inside or outside the detected shape. These surface weights are specified on the surface of the detected mesh, $\{\mathbf{p}_j\}$, which is pre-annotated with corresponding weights, $\{\omega_j^*\}$:

$$\omega_{l(\mathbf{x})}^* = \omega_j^* \quad \text{where} \quad j = \operatorname{argmin}_j \|\mathbf{x} - \mathbf{p}_j\|^2. \quad (20)$$

To maintain a consistent weighting over the energy minimization, the points, \mathbf{p}_j , are also evolved during the minimization [19].

In practice, since our detection result is typically undersegmented, we use a large uniform penalization for the interior deviations, $\omega_{l(\mathbf{x})}^{\text{in}} = 10$. In order to allow the fine-scale refinement to fill in small details, we use a smaller weight for the exterior surface deviations and constrain the solution to not deviate much from the detection in the region of the airways. Fig. 9 illustrates the spatially varying exterior weights, $\omega_{l(\mathbf{x})}^{\text{out}}$.




Fig. 10 Modules making up the pose estimation are connected by a set of pose candidates, with each detector augmenting more information to its output candidates. All *Modules* consume the CT volume as input.




Fig. 11 The full pipeline uses the pose estimation, with the later modules performing consecutive refinements of the mesh.

The contribution of the other terms are balanced by qualitative tuning on a few cases. Here we used a weight of $\alpha = 2$ for the data term, a smoothness weight of $\gamma = 15$, to keep the lung surface smooth, and a large weight on the overlap term, $\beta = 1000$.

3.3.1 Optimization

Each of the lungs (and possibly other adjacent organs) has a corresponding energy in the form of (15), and the minimizers of these energies are coupled by the overlap penalization term (19). In practice, interleaving gradient descents for each of the organs is sufficient to get a desired solution [18].

Specifically, each lung is refined along the negative gradient of the energy term,

$$\frac{\partial \Phi_i}{\partial t} \approx \frac{(\Phi_i^{t+1} - \Phi_i^t)}{\delta t} = -\frac{\partial E(\Phi_i; \{\Phi_j^t\}_{j \neq i})}{\partial \Phi_i}, \quad (21)$$

where the gradient for the coupled overlap term uses the solution from the previous time, Φ_j^t . Each organ is then evolved along (21) in lockstep using a discrete time step δt .

The gradient descent is terminated when either a fixed number of iterations is reached or if the movement of the surface in a given iteration falls below a predefined threshold. For an efficient solution, the level set surfaces are represented with a narrow-band scheme within ± 2 voxels of the zero crossing.

4 A Software Architecture for Detection and Segmentation

In order to implement our lung segmentation algorithm, we use an Integrated Detection Network (IDN) [29], which is a modular design that specifically caters to such detection networks. IDN promotes modularity and reusability through a basic abstraction that decomposes algorithms into the operations, which are called *Modules*, and the *Data* that these modules consume and produce. In contrast to a blackboard model, where information is globally shared through a blackboard, e.g., as done by Brown *et al.* for lung segmentation [4], the connections between modules illustrate the data flow of the entire algorithm. Data is only shared with a module if it is necessary.

Each *Module* provides a set of input slots for incoming *Data* and owns a (possibly empty) set of output *Data* objects. Different *Module* operations may be connected as long as their connecting *Data* object type agrees with the respective input/output types of the *Module*. The connections between these *Modules* through their *Data* forms an acyclic graph.

Modules and *Data* carry all the necessary information to allow both training and detection. In this way, the spatial dependencies between objects in the HDN can be directly encoded with an IDN network, and this same network of *Modules* can be used to train the classifiers as well as to perform the detection and segmentation.

4.1 Pose Detection Modules

Pose detection *Modules* parallel their respective HDN detections and take as input a set of candidate poses and produce an augmented set of candidates, where, for example, the output of the orientation detector contains both position and orientation. Estimation of the object pose and the first few PCA coefficients is achieved by connecting the individual detectors, as illustrated in Fig. 10.

4.2 Segmentation Modules

The main components of the segmentation are also decomposed into independent modules. The stable landmark detector uses an input mesh to predict the locations of the stable landmarks, and outputs a mesh deformed to fit these landmark positions. The learning-based boundary refinement refines a mesh, as does the level set refinement (Fig. 11). All of these elements can

be interchanged provided the input/output types agree. Notice that the final IDN graph parallels the original flowchart of the algorithm in Fig. 6.

5 Experiments

In this section, we perform a qualitative and quantitative analysis at several stages of our algorithm¹. We use a data set of 260 diagnostic CT scans, which has images of varying contrast and slice thickness (ranging from 0.5mm to 5mm). The lung surfaces have been manually annotated in each of these scans by an expert. This data was randomly partitioned into two sets: 192 volumes that were used to train our detectors, and another 68 volumes that were used for testing. The shape models (§3.1.1) computed from the input annotations have a total of 614 vertices for each the left and right lung.

As the stable landmarks are selected by their error rates, we first illustrate the error rates for all the candidate landmarks by computing the accuracy on the training data. We then show how the automatic greedy selection of spatially varying landmarks tends to choose a set of well distributed landmarks that are close to anatomical structures. We then demonstrate on testing data how this landmark-based initialization helps improve the results over the pose+PCA detector on challenging cases. Finally, we illustrate how the fine-scale refinement further improves the accuracy, even in the presence of pathologies.

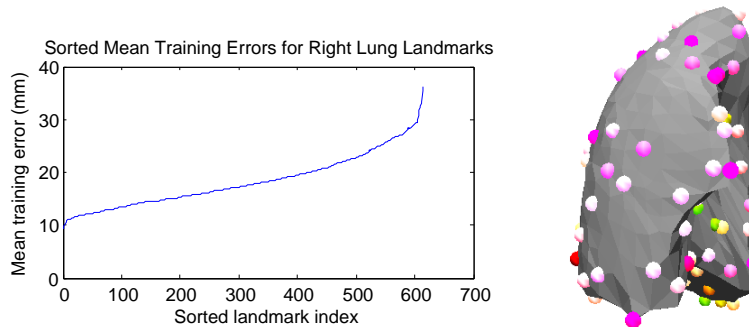


Fig. 12 Sorted mean errors of the 614 landmarks computed from all training volumes for the right lungs, and the 143 selected landmarks obtained using the spatial filtering with parameter of $r = 20\text{mm}$ visualized on the mesh.

¹ The initial analysis and comparison of the performance of the boundary landmarks also appears in our earlier conference publication [30]

Table 1 Testing results of symmetrical point-to-mesh comparisons (in mm’s) between the detected results and annotations for both lungs, with and without stable landmark detection [30], and also with the final level set refinement presented in §3.3.

| Lung | Landmark | Level set | Mean (std.) | Med. | Min | Max | 80% |
|-------|----------|-----------|-----------------|------|------|-------|------|
| right | no | no | 2.35 ± 0.86 | 2.16 | 1.40 | 6.43 | 2.57 |
| right | yes | no | 1.98 ± 0.62 | 1.82 | 1.37 | 4.87 | 2.18 |
| right | yes | yes | 1.30 ± 0.64 | 1.12 | 0.91 | 4.60 | 1.38 |
| left | no | no | 2.31 ± 2.42 | 1.96 | 1.28 | 21.11 | 2.22 |
| left | yes | no | 1.92 ± 0.73 | 1.80 | 1.19 | 6.54 | 2.15 |
| left | yes | yes | 1.34 ± 0.77 | 1.13 | 0.93 | 6.78 | 1.39 |

5.1 Landmark errors

Figure 12 illustrates the landmark scores (computed as the average distance between detected and ground truth) for all 614 landmarks computed on the training set for the right lung (as in §3.1.4). The larger errors on the right of the plot, show that some of the landmarks are less reliable; these landmarks are less likely to be selected. For the landmarks with lower scores, some of the error is attributed to slight misalignments that occur during the correspondence establishment phase used to generate the ground truth landmark positions. When using a spatial radius of $r = 20\text{mm}$, a total of 143 landmarks were chosen for the right lung and 133 for the left lung.

For illustration purposes, by increasing the spatial filtering radius to $r = 70\text{mm}$, we show that the greedy selection is effective in choosing a distributed set of 12 stable landmarks (Fig. 13). The resulting landmarks are often chosen close to anatomical structures, such as ribs and vertebrae. The spatial filtering also ensures that some landmarks are also selected near the top and bottom of the lungs.

5.2 Segmentation Accuracy and Quality

In this next experiment, we analyze the improvement that is attained at several stages of the algorithm on our testing set. First, we ran the algorithm up to the freeform refinement without the use of the stable landmarks (i.e., skipping §3.1.4). We then included the stable landmark stage of the algorithm. And finally, we then ran the full algorithm including the level set refinement. For comparison, the symmetric surface-to-surface distance was computed between the detected and ground truth segmentation.

The results in Table 1 show that the inclusion of the stable landmarks gave a significant ($p < 0.05$) decrease in the surface-to-surface distance for both the left and right lung [30]. In the case of the left lung, notice the decrease




Fig. 13 Using a larger spatial filter radius (70mm), the twelve strongest landmarks selected of the 614 mesh vertices are shown. These landmarks are selected near distinctive anatomical structures such as ribs (3, 4, 5, 12), vertebrae (1, 2) and top (5) and bottom of the lung (9, 10, 11). Reprinted from [30] with kind permission from Springer Science and Business Media.

in the maximum error. This is caused by the stable landmarks correcting a failure case. Finally, notice that the fine-scale level set refinement further reduced this error by filling in the small scale details.

The main advantage of the stable landmarks is the ability to overcome a poor initialization that can occur due to appearance variability in pathologies. Figure 14 illustrates several qualitative examples of pathologies where the robust landmarks lead to an improved result of the free-form refinement.

As the fine-scale surface refinement is constrained by the initial surface estimate, a good initialization is important for final accuracy. Figure 15 illustrates that the final level set refinement fills in the small regions, and that the segmentation still includes the large pathological regions due to the surface prior constraint.

In Fig. 16, we demonstrate that all the terms of the level set energy are necessary. When the *prior* and *overlap* terms are not used (only the data




Fig. 14 Comparison of freeform boundary results without using the stable landmarks (left column of every set) and with stable landmark detection (right column of every set). Reprinted from [30] with kind permission from Springer Science and Business Media.

and regularization terms in (15)), the resulting segmentation excludes the pathological region even though the initialization includes this region. Further, without the *prior* term, the level set leaks into the airways and results in a non-smooth surface near the hilum. Enabling the *prior* term fixes these problems, but the segmentation still has overlap near the anterior junction (middle of Fig. 16). Only when both the *prior* and *overlap* terms are enabled is the desired segmentation achieved.

6 Conclusion

In this article we presented a multi-stage learning method for segmenting challenging lung cases in thoracic CT. The algorithm overcomes the diversity of lung appearance in the presence of pathologies through the use of a statistical shape model and by detecting robust landmarks that use anatomical context outside the lung.

As pathologies may only affect a portion of the lung, it is important to also use a fine-scale refinement to obtain highly accurate segmentations in the healthy regions. Graph-based methods can be used for this refinement [12], but we propose to use a level set approach that uses adaptive region intensity constraints and enforces that the refined surface cannot deviate too far from the robustly initialized detected surfaces. We have illustrated that the resulting segmentation is robust to a variety of pathologies, including high density tissue associated with diffuse lung disease as well as tumors.





Fig. 15 Comparison of the basic boundary detection (top) and fine-scale level set refinement (bottom) on three example cases. The basic boundary detection gives a good coarse estimate that includes the high density tissue, and the fine-scale refinement remains close in these regions but also includes the fine-scale voxel-level details.

Intelligent editing may still be necessary to fix up failures in extreme cases [17]. However, the same principles used in our automatic algorithm, such as the use of machine learning to integrate shape, texture and anatomical cues, can be used to guide the development of such interactive techniques.

References

1. Abe, Y., Hanai, K., Nakano, M., Ohkubo, Y., Hasizume, T., Kakizaki, T., Nakamura, M., Niki, N., Eguchi, K., Fujino, T., Moriyama, N.: A computer-aided diagnosis (CAD) system in lung cancer screening with computed tomography. *Anticancer Research* 25, 483–488 (2005)
2. Adams, H., Bernard, M., McConnochie, K.: An appraisal of CT pulmonary density mapping in normal subjects. *Clin Radiol.* 43(4), 238–42 (Apr 1991)
3. Armato III, S., Sensakovic, W.: Automated lung segmentation for thoracic CT: Impact on computer-aided diagnosis. *Acad. radiol.* 11(9), 1011–1021 (2004)



No overlap term, no prior term No overlap term, with prior term With overlap and prior terms

Fig. 16 In this example, the same data from the right of Fig. 15 is used to demonstrate that both the *overlap* and *prior* terms are necessary to obtain good segmentations in the level set refinement. Without the *prior* term, the segmentation doesn't include the tumor (left), and without the *overlap* term, the resulting segmentations may have overlap (left and middle). With both terms enabled, the desired segmentation, which includes the pathological region and removes overlap, is obtained (right).

4. Brown, M.S., McNitt-Gray, M.F., Mankovich, N.J., Goldin, J.G., Hiller, J., Wilson, L.S., Aberle, D.R.: Method for segmenting chest CT image data using an anatomical model: preliminary results. *IEEE Trans on Medical Imaging* 16(6), 828–839 (1997)
5. Chan, T., Vese, L.: Active contours without edges. *Image Processing, IEEE Transactions on* 10(2), 266–277 (feb 2001)
6. Cootes, T., Hill, A., Taylor, C., Haslam, J.: Use of active shape models for locating structures in medical images. *Image and Vision Computing* 12(6), 355–365 (1994)
7. Cremers, D., Rousson, M., Deriche, R.: A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape. *Int. J. Comput. Vision* 72(2), 195–215 (Apr 2007)
8. Haider, C., Bartholmai, B., Holmes, D., Camp, J., Robb, R.: Quantitative characterization of lung disease. *Computerized Medical Imaging and Graphics* 29(7), 555–563 (2005)
9. Haralick, R.M., Shanmugam, K., Dinstein, I.: Textural features for image classification. *Systems, Man and Cybernetics, IEEE Trans on* 3(6), 610–621 (nov 1973)
10. Hedlund, L., Anderson, R., Goulding, P., Beck, J., Effmann, E., Putman, C.: Two methods for isolating the lung area of a CT scan for density information. *Radiology* 144(2) (July 1982)
11. Hu, S., Hoffman, E., Reinhardt, J.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. *IEEE TMI* 20(6), 490–498 (2002)
12. Hua, P., Song, Q., Sonka, M., Hoffman, E., Reinhardt, J.: Segmentation of pathological and diseased lung tissue in CT images using a graph-search algorithm. *ISBI* (March 2011)

13. International Consensus Statement: Idiopathic pulmonary fibrosis: diagnosis and treatment. American Thoracic Society (ATS), and the European Respiratory Society (ERS)., vol. 161, pp. 646–664 (2000)
14. International Consensus Statement: American thoracic society/european respiratory society international multidisciplinary consensus classification of the idiopathic interstitial pneumonias. *Am J Respir Crit Care Med* 165, 277–304 (2002)
15. Kanazawa, K., Kawata, Y., Niki, N., Satoh, H., Ohmatsu, H., Kakinuma, R., Kaneko, M., Moriyama, N., Eguchi, K.: Computer-aided diagnosis for pulmonary nodules based on helical CT images. *Comput Med Imaging Graph.* 22(2), 157–67 (1998)
16. King, T.E.: Clinical advances in the diagnosis and therapy of the interstitial lung diseases. *Am. J. Respir. Crit. Care Med.* 172(3), 268–279 (2005)
17. Kockelkorn, T., van Rikxoort, E., Grutters, J., van Ginneken, B.: Interactive lung segmentation in CT scans with severe abnormalities. In: *ISBI*. pp. 564–567 (2010)
18. Kohlberger, T., Sofka, M., Zhang, J., Birkbeck, N., Wetzl, J., Kaftan, J., Declerck, J., Zhou, S.: Automatic multi-organ segmentation using learning-based segmentation and level set optimization. In: *MICCAI '11* (2011)
19. Kohlberger, T., Uzunbaş, M.G., Alvino, C., Kadir, T., Slosman, D.O., Funke-Lea, G.: Organ segmentation with level sets using local shape and appearance priors. In: *Proceedings of MICCAI: Part II*. pp. 34–42. *MICCAI '09* (2009)
20. Korfiatis, P., Kalogeropoulou, C., Karahaliou, A., Kazantzi, A., Skiadopoulos, S., Costaridou, L.: Texture classification-based segmentation of lung affected by interstitial pneumonia in high-resolution CT. *Medical physics* 35, 5290 (2008)
21. Ling, H., Zhou, S.K., Zheng, Y., Georgescu, B., Suehling, M., Comaniciu, D.: Hierarchical, learning-based automatic liver segmentation. In: *CVPR*. pp. 1–8. *IEEE Computer Society, Los Alamitos, CA, USA* (2008)
22. Prasad, M., Brown, M., Ahmad, S., Abtin, F., Allen, J., da Costa, I., Kim, H., McNitt-Gray, M., Goldin, J.: Automatic segmentation of lung parenchyma in the presence of diseases based on curvature of ribs. *Acad. radiol.* 15(9), 1173–1180 (2008)
23. Pu, J., Roos, J., Yi, C., Napel, S., Rubin, G., Paik, D.: Adaptive border marching algorithm: Automatic lung segmentation on chest CT images. *Computerized Medical Imaging and Graphics* 32(6), 452–462 (2008)
24. Reeves, A.P., Kostis, W.J.: Computer-aided diagnosis for lung cancer. *Radiol Clin North Am.* 38(3), 497–509 (May 2000)
25. Robinson, P., L, K.: Pulmonary tissue attenuation with computed tomography: comparison of inspiration and expiration scans. *J Comput Assist Tomogr* 3(6), 740–8 (Dec 1979)
26. Silva, A., Silva, J.S., Santos, B.S., Ferreira, C.: Fast pulmonary contour extraction in X-ray CT images: A methodology and quality assessment. In: *SPIE Conf on Medical Imaging: Physiology and Function from Multidimensional Images*. vol. 4321, pp. 216–224 (2001)
27. Sluimer, I., Schilham, A., Prokop, M., van Ginneken, B.: Computer analysis of computed tomography scans of the lung: a survey. *IEEE TMI* 25(4), 385–405 (2006)
28. Sluimer, I., Prokop, M., van Ginneken, B.: Toward automated segmentation of the pathological lung in CT. *IEEE TMI* 24(8), 1025–1038 (2005)
29. Sofka, M., Ralovich, K., Birkbeck, N., Zhang, J., Zhou, S.: Integrated detection network (IDN) for pose and boundary estimation in medical images. In: *Proceedings of the 8th International Symposium on Biomedical Imaging (ISBI 2011)*. Chicago, IL (30 Mar – 2 Apr 2011)
30. Sofka, M., Wetzl, J., Birkbeck, N., Zhang, J., Kohlberger, T., Kaftan, J., Declerck, J., Zhou, S.: Multi-stage learning for robust lung segmentation in challenging CT volumes. In: *MICCAI '11* (18–22 Sep 2011)

31. Sofka, M., Zhang, J., Zhou, S., Comaniciu, D.: Multiple object detection by sequential Monte Carlo and hierarchical detection network. In: CVPR (13–18 Jun 2010)
32. Sun, S., McLennan, G., Hoffman, E.A., Beichel, R.: Model-based segmentation of pathological lungs in volumetric ct data. In: The Third International Workshop on Pulmonary Image Analysis (2010)
33. Tu, Z.: Probabilistic boosting-tree: learning discriminative models for classification, recognition, and clustering. In: ICCV. pp. 1589–1596 Vol. 2 (2005)
34. Uchiyama, Y., Katsuragawa, S., Abe, H., Shiraishi, J., Li, F., Li, Q., Zhang, C.T., Suzuki, K., Doi, K.: Quantitative computerized analysis of diffuse lung disease in high-resolution computed tomography. *Med. Phys* 30, 2440–2454 (2003)
35. Ukil, S., Reinhardt, J.: Anatomy-guided lung lobe segmentation in X-ray CT images. *IEEE TMI* 28(2), 202–214 (2009)
36. Ukil, S., Reinhardt, J.M.: Smoothing lung segmentation surfaces in three-dimensional X-ray CT images using anatomic guidance. *Acad. radiol* 12(12), 1502–1511 (2005)
37. Uppaluri, R., Hoffman, E.A., Sonka, M., Hartley, P.G., Hunninghake, G.W., McLennan, G.: Computer recognition of regional lung disease patterns. *Am. J. Respir. Crit. Care Med.* 160(2), 648–654 (August 1 1999)
38. Uppaluri, R., Mitsa, T., Sonka, M., Hoffman, E.A., McLennan, G.: Quantification of Pulmonary Emphysema from Lung Computed Tomography Images. *Am. J. Respir. Crit. Care Med.* 156, 248–254 (1997)
39. Wang, J., Li, F., Li, Q.: Automated segmentation of lungs with severe interstitial lung disease in CT. *Medical physics* 36, 4592 (2009)
40. Wang, J., Li, F., Doi, K., Li, Q.: A Novel Scheme for Detection of Diffuse Lung Disease in MDCT by Use of Statistical Texture Features. In: SPIE 7260 (2009)
41. Zhang, L., Hoffman, E., Reinhardt, J.: Atlas-driven lung lobe segmentation in volumetric X-ray CT images. *IEEE TMI* 25(1), 1–16 (2005)
42. Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., Comaniciu, D.: Fast automatic heart chamber segmentation from 3D CT data using marginal space learning and steerable features. In: ICCV. pp. 1–8. IEEE (2007)