DIGITAL ACCESS TO
SCHOLARSHIP AT HARVARD

Robust Initialization of Active Shape Models for Lung Segmentation in CT Scans: A Feature-Based Atlas Approach

(Article begins on next page)

Research Article

Robust Initialization of Active Shape Models for Lung Segmentation in CT Scans: A Feature-Based Atlas Approach

Gurman Gill,[1,2] Matthew Toews,[3] and Reinhard R. Beichel[1,2,4]

[1] *Department of Electrical and Computer Engineering, The University of Iowa, Iowa City, IA 52242, USA*
[2] *The Iowa Institute for Biomedical Imaging, The University of Iowa, Iowa City, IA 52242, USA*
[3] *Brigham and Women's Hospital, Harvard Medical School, Boston, MA 02115, USA*
[4] *Department of Internal Medicine, The University of Iowa, Iowa City, IA 52242, USA*

Correspondence should be addressed to Reinhard R. Beichel; reinhard-beichel@uiowa.edu

Model-based segmentation methods have the advantage of incorporating a priori shape information into the segmentation process but suffer from the drawback that the model must be initialized sufficiently close to the target. We propose a novel approach for initializing an active shape model (ASM) and apply it to 3D lung segmentation in CT scans. Our method constructs an atlas consisting of a set of representative lung features and an average lung shape. The ASM pose parameters are found by transforming the average lung shape based on an affine transform computed from matching features between the new image and representative lung features. Our evaluation on a diverse set of 190 images showed an average dice coefficient of 0.746 ± 0.068 for initialization and 0.974 ± 0.017 for subsequent segmentation, based on an independent reference standard. The mean absolute surface distance error was 0.948 ± 1.537 mm. The initialization as well as segmentation results showed a statistically significant improvement compared to four other approaches. The proposed initialization method can be generalized to other applications employing ASM-based segmentation.

1. Introduction

Lung segmentation methods are required for automated lung image analysis and to facilitate tasks like lung volume calculation, quantification of lung diseases, or nodule detection. Model-based techniques, such as active shape models (ASMs), have been employed for segmenting lungs in three-dimensional CT scans [1–4] or two-dimensional chest radiographs [5–8] because they incorporate prior knowledge of anatomical shape variation into the segmentation process. Progress has been made over the years in improving the robustness and accuracy of fitting the ASMs, as described in [9] for medical images in general and in [1, 4] for lung segmentation. However, as Ibragimov et al. [9] concluded, the major drawback of ASMs lies in its initialization; for correctly converging to the target image structure, the ASM must be initialized sufficiently close to it.

Several techniques have been proposed for initializing ASMs for lung segmentation. The most straightforward ones place the ASM manually [8] or in a semiautomatic manner by annotating the lung size and position [7]. Van Ginneken et al. [5] and Wang et al. [7] adopt a multiresolution framework [10] to iteratively get closer to the target image structure. However, their method assumes that the target is within a certain distance to the initial model. Several other methods initialize the model based on detecting certain points on or in close proximity to lung anatomy. Iakovidis et al. [6] employ heuristics based on finding salient control points on the spinal cord and rib-cage along with a selective thresholding algorithm. Sofka et al. [4] detect the carina of trachea and use a hierarchical detection network to predict pose parameters of left and right lung. Wilms et al. [3] use heuristics based on the bronchial tree, Sun et al. [1] detect the rib-cage, and Gill et al. [2] predict the location of the carina and lung apex to initialize an ASM. The drawbacks of these methods are that, firstly, they depend on the quality of salient point detection. For example, it is possible for rib-cage and bronchial tree to be incorrectly identified due to local changes

(e.g., disease, imaging artifacts, etc.) and the heuristics employed may not be robust to the incorrect detections. Secondly, they may not work equally well for lung images at different respiratory states such as total lung capacity (TLC) and functional residual capacity (FRC) because the heuristics used may not account for the deformation of lungs [2].

This paper presents a fast and robust method for automatically initializing an active shape model to lung CT scans by learning a feature-based atlas Ψ comprising average lung shape $\boldsymbol{\xi}$ and a set of representative lung features $\boldsymbol{\rho_r}$. The representative features $\boldsymbol{\rho_r}$ can then be matched to features $\boldsymbol{\rho_t}$ in a new image, based on which a mapping is computed to transform the average lung shape $\boldsymbol{\xi}$ to the new image space. This information is then utilized to calculate ASM initialization parameters. Our method is based on correspondences of generic local features identified in the CT volume, rather than a small set of specific salient points.

For computing the feature-based atlas Ψ, we adopt a feature-based alignment (FBA) method [11], which identifies generic lung features in a data-driven fashion and aligns a set of training CT images. The original FBA method operates by identifying a similarity transform between sets of 3D scale-invariant image features. While feature correspondences can be identified, the similarity transform is insufficient for describing the geometrical variation across lung respiratory states (e.g., TLC and FRC). In this paper, we extend the FBA method to affine alignment, which allows us to capture variation in lung shapes and respiratory motion within a single atlas, and demonstrate the advantages compared to the original approach.

We have evaluated our method on 190 images consisting of normal and diseased lungs imaged at different respiratory states. Comparison of results with those provided in [1, 2] shows our initialization method significantly improves lung segmentation accuracy. In order to establish the quality of affine transform produced by our approach, we also compare against a registration-based approach that provides affine alignment based on the robust block matching (RBM) method [12]. The RBM method is a well known registration method that was one of the top performers in the EMPIRE 2010 challenge, which compared 34 registration algorithms [13].

2. Prior Work

The presented approach builds on two main components: a collection of local scale-invariant features for image alignment and a robust active shape model for image segmentation.

Local scale-invariant features are distinctive image patches defined by location, scale, and orientation. Scale-invariant features emerged in computer vision literature as a means of repeatably detecting image structure arising from the same underlying scene or object in different images, despite global changes in translation, scale, and orientation. Feature detection operates by identifying the location and scale of image patches maximizing a function of saliency, for instance, the magnitude of Gaussian derivatives in scale [14] or space [15]. Once identified, distinctive intensity patches

are be encoded and robustly matched between images despite geometrical deformations or missing or occluded structure. Extensive comparisons have shown the gradient orientation histogram (GoH) descriptor to be among the most effective feature encodings [16], in particular rank-ordered variants [17] in terms of achieving correct correspondences. Scale-invariant feature methods have been generalized to 3D medical image context, and this work adopts the approach of Toews and Wells [11] where 3D features are detected extrema in a difference-of-Gaussian scale-space pyramid, which are encoded by a rank-ordered GoH descriptor.

For model-based lung segmentation, we utilized a combination of robust active shape model (RASM) and subsequent graph-based optimal surface finding (OSF) algorithm [1]. The RASM employs a point distribution model (PDM), which is constructed separately for left and right lungs, and the subsequent segmentation steps are carried out separately as well. Below we briefly review the steps employed for initializing the PDM of the RASM.

Given N training CT images $\{C_1, C_2, \ldots, C_N\}$ and pre-segmented training lung shapes (left or right) $\{\mathbf{S_1}, \mathbf{S_2}, \ldots, \mathbf{S_N}\}$, a set of k corresponding points or landmarks are automatically identified for each training shape: $\mathbf{S_i} = \{(x_{i,1}, y_{i,1}, z_{i,1}), \ldots, (x_{i,k}, y_{i,k}, z_{i,k})\}$. A PDM is constructed by aligning the training shapes $\mathbf{S_i}$ using Procrustes analysis and performing principal component analysis (PCA) on them. Let the aligned shapes be denoted by $\mathbf{S_i^a} = \{(x_{i,1}^a, y_{i,1}^a, z_{i,1}^a), \ldots, (x_{i,k}^a, y_{i,k}^a, z_{i,k}^a)\}$. A new shape $\mathbf{S}$ can be now represented in terms of the mean shape, $\boldsymbol{\mu}_{\text{pdm}} = (1/N) \sum_i \mathbf{S_i^a}$, using the linear model:

$$\mathbf{S} = \boldsymbol{\mu}_{\text{pdm}} + \mathcal{P}\mathbf{b}, \tag{1}$$

where $\mathcal{P}$ denotes the shape eigenvector matrix and $\mathbf{b}$ represents the shape coefficients [1]. To start the RASM fitting process, the mean shape $\boldsymbol{\mu}_{\text{pdm}}$ is initialized in the target image space based on pose parameters $\mathcal{T}$ comprising of isotropic scale, location, and rotation. In this paper, we present a novel approach to compute the pose parameters $\mathcal{T}$ for RASM initialization.

3. Methods

In our method, we generate a feature-based atlas based on aligning the training CT images C_i using the extended FBA method (Section 3.3). Utilizing the alignment information, the individual training lung shapes $\mathbf{S_i}$ can be transformed to the atlas space and averaged using existing landmarks. The main purpose of building an atlas is to embed an average lung shape with a set of representative lung features that can then be matched to features in a new CT image.

Our method for automatically initializing a RASM to lung CT images primarily consists of 3 steps: (i) building an atlas $\Psi = \{\boldsymbol{\xi}, \boldsymbol{\rho_r}\}$ comprising average lung shape $\boldsymbol{\xi}$ and a set of representative lung features $\boldsymbol{\rho_r}$ (Section 3.1); (ii) transforming $\boldsymbol{\xi}$ to the new subject image space based on matching features between the new image and $\boldsymbol{\rho_r}$; followed by (iii) computing the pose parameters $\mathcal{T}$ (Section 3.2) for RASM initialization.
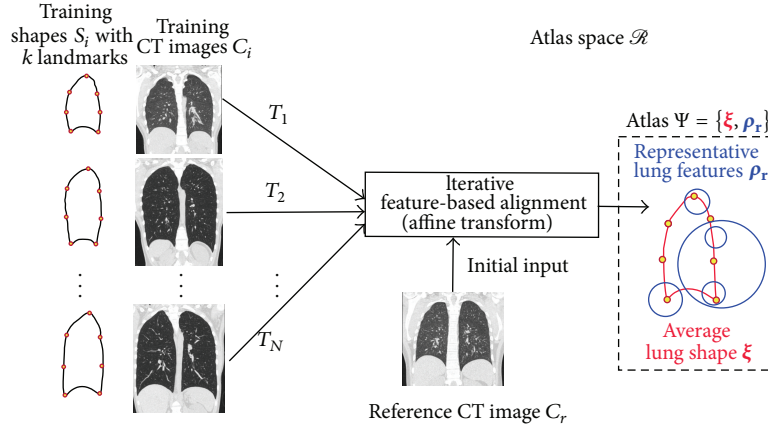
FIGURE 1: Generation of the feature-based atlas Ψ, comprising the average lung shape $\boldsymbol{\xi}$ and representative lung features $\boldsymbol{\rho_r}$ in the atlas space $\mathscr{R}$. Note that the features $\boldsymbol{\rho_r}$ are derived from training CT images and encode both appearance and geometric properties [11]. For sake of clarity, the above process is shown only for the right lung.
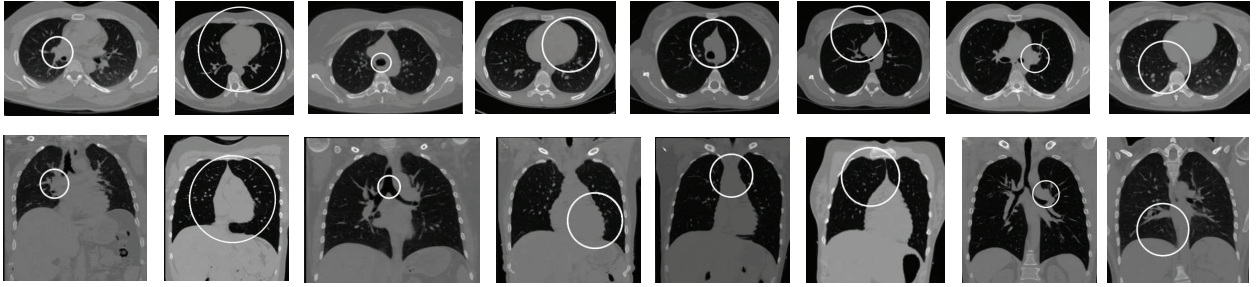

FIGURE 2: Example of representative lung features $\boldsymbol{\rho_r}$. Each column shows a single feature in axial and frontal view.

3.1. Feature-Based Atlas.

To build an atlas Ψ, we derive an average lung shape $\boldsymbol{\xi}$ and associated set of representative lung features $\boldsymbol{\rho_r}$ as follows (Figure 1).

(a) 3D scale-invariant image features $\mathbf{f_i}$ are extracted from all training CT images C_i. One of the training images is chosen as the reference image C_r. Let the features in C_r be denoted by $\mathbf{f_r}$.

(b) The FBA method (Section 3.3) is employed to find an affine transform T_i from a matching set of features in $\mathbf{f_i}$ and $\mathbf{f_r}$.

(c) Training features $\mathbf{f_i}$ are transformed from training to atlas space $\mathscr{R}$ via T_i. Let these be denoted by $T_i\mathbf{f_i}$.

(d) A modeling step [18] computes a new representative set of features $\mathbf{f_r}$ from the set $\{T_i\mathbf{f_i}\}_{i=1}^{N}$.

(e) Steps (b)–(d) are repeated until convergence, that is, when the Frobenius norm of the transformation matrix difference $\max \|T_i^t - T_i^{t-1}\|$ is zero [18]. This yields a representative set of lung features $\boldsymbol{\rho_r} = \mathbf{f_r}$ (Figure 2). Note that the iterative group-wise alignment procedure reduces the dependency on the selection of the reference image [18].

(f) Training shapes $\{\mathbf{S_i}\}_{i=1}^{N}$ are transformed to the atlas space by using the affine transforms T_i. We leverage existing landmarks across shapes $\mathbf{S_i}$ and take their average to produce the average lung shape $\boldsymbol{\xi} = (1/N)\sum_{i=1}^{N} T_i\mathbf{S_i}$. Note that same transformations T_i are used to separately compute the average for left and right lung shapes.

3.2. Computing Pose Parameters.

3D scale-invariant image features $\boldsymbol{\rho_t}$ are extracted from the new CT image C_{new} and matched to the representative lung features $\boldsymbol{\rho_r}$ (Figure 3). The matching set of features output an affine transform T_{new} that describes how C_{new} can be transformed to the atlas space $\mathscr{R}$. Since we already have a representation of the average lung shape in $\mathscr{R}$, it is transformed to the new subject image space using $T_{\text{new}}^{-1}\boldsymbol{\xi}$. Note that the same transformation T_{new}^{-1} is used for transforming both the left and right lung shapes.

RASM model fitting starts by placing the average PDM shape $\boldsymbol{\mu}_{\text{pdm}}$ (1) in the image space. The pose parameters $\mathscr{T}$ are obtained by transforming $\boldsymbol{\mu}_{\text{pdm}}$ to the image space by means of a Procrustes analysis $\mathscr{T}: \boldsymbol{\mu}_{\text{pdm}} \rightarrow T_{\text{new}}^{-1}\boldsymbol{\xi}$.

3.3. Extending the Feature-Based Alignment Approach.

As noted earlier, our work extends the original FBA approach [11] to compute affine alignment. First we briefly describe the key step of feature matching, followed by the new procedure for affine refinement.

Feature matching begins by computing nearest neighbors (NN) between feature descriptors in the image and
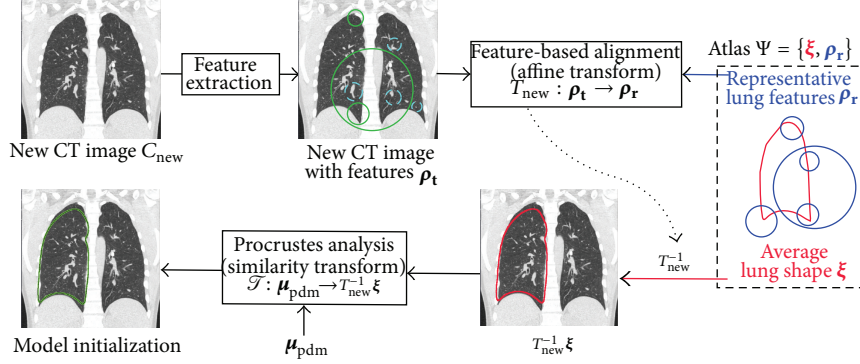
FIGURE 3: Schematic diagram showing the model initialization process for a new CT image. Features $\boldsymbol{\rho}_t$ are shown superimposed on the coronal slice of a CT image. Features matching with representative lung features $\boldsymbol{\rho}_r$ are marked with a solid line while the others are marked with a dashed line. For sake of clarity, the above process is shown only for the right lung.

atlas, using the Euclidean distance similarity measure. This procedure can be computed efficiently via fast approximate nearest neighbor methods. Each NN image-to-atlas feature match provides an estimate of the global image-to-atlas similarity transform, and a highly probable global transform along with inlier image-to-atlas matches is identified via a robust probabilistic voting formulation similar to the Hough transform [11]. Note that incorrect outlier correspondences have no effect on the transform identified, and for this reason alignment is considered robust.

Once the inlier image-to-atlas matches are identified, an affine transform T_i is fitted between them with a least squares approach. This allows our framework to account for a higher degree of shape variability, for example, lungs in different respiratory states, and thus produce a lung shape representative of lung variation across the population. Note that the modeling step (d) in Section 3.1 requires isotropic features (i.e., spherical as against elliptical 3D shape) [11]. For that purpose, we compute the geometric mean from the scale components of the affine transform to transform the feature scale.

The original FBA approach achieves alignment across global image rotation via rotationally invariant features, where a 3D orientation is assigned to individual features based on the structure of the local image gradient. Our lung CT data exhibit only minor orientation differences between subjects due to a similar imaging protocol, and we found that assigning a fixed feature orientation results in a higher number of image-to-image correspondences and improved alignment. This may be because many pulmonary structures exhibit rotational symmetry, for example, airways, in which case orientation is inherently ambiguous. Thus, features of fixed orientation appear to be more effective than rotationally invariant features for identifying correspondences in the case of minor intersubject orientation differences.

4. Evaluation

4.1. Image Data. We selected 190 multidetector computed tomography (MDCT) thorax scans of lungs for testing

from 6 different sets S_{normal}, S_{asthma}, S_{COPD}, S_{mix}, S_{IPF}, and S_{tumor} with no significant abnormalities (normals), asthma (both severe and nonsevere), chronic obstructive pulmonary disease (COPD with GOLD 1 to 4), mixture of different lung diseases, idiopathic pulmonary fibrosis (IPF), and lung cancer, respectively. The total number of scans in sets S_{normal}, S_{asthma}, S_{COPD}, S_{mix}, S_{IPF}, and S_{tumor} was 20, 24, 28, 26, 62, and 30, respectively. The first four datasets contained pairs of TLC and FRC images while the last two were all TLC images. The image sizes varied from $512 \times 512 \times 205$ to $512 \times 512 \times 780$ voxels. The slice thickness of images ranged from 0.5 to 1.25 mm and the in-plane resolution from 0.49×0.49 to 0.91×0.91 mm.

4.2. Experimental Setup. The PDM was built using a separate set of 75 TLC and 75 FRC normal lung scans. The average lung shape $\boldsymbol{\xi}$ was constructed using 50 TLC and 50 FRC scans that were a subset of those used in building the PDM. A 3D SIFT-based feature extractor [18] was used to extract approximately 2500 features in each lung image from which high density (>0 HU) structures (e.g., bones) were automatically removed. The affine feature-based alignment system (Section 3.1) grouped these into 1000 representative lung features, forming one component of our atlas.

We followed the implementation of the RASM-OSF framework for lung segmentation presented in [1] with the exception of RASM initialization. We refer to different methods evaluated in this paper based on their method of initialization: by detecting ribs (M_{ribs}) [1], carina and apex (M_{capex}) [2], and average lung shape $\boldsymbol{\xi}$ in this paper using affine FBA ($M_{FBA^{Aff}}$), using similarity transform provided by original FBA ($M_{FBA^{Sim}}$) [18], and using affine transform obtained from RBM ($M_{RBM^{Aff}}$) [12] (implementation: NiftyReg, http://www.cs.ucl.ac.uk/staff/m.modat).

4.3. Independent Reference and Quantitative Indices. An independent reference standard was generated for all test data sets by first using a commercial lung image analysis software package Apollo (VIDA Diagnostics Inc., Coralville, IA) to automatically create lung segmentations. These were

TABLE 1: Average initialization and final segmentation accuracy obtained by $M_{\text{FBA}^{\text{Aff}}}$ for each test set.

	S_{normal}	S_{asthma}	S_{COPD}	S_{mix}	S_{IPF}	S_{tumor}
D_{init}	0.786 ± 0.054	0.774 ± 0.057	0.741 ± 0.059	0.717 ± 0.067	0.718 ± 0.067	0.783 ± 0.054
D_{final}	0.985 ± 0.006	0.979 ± 0.011	0.981 ± 0.012	0.971 ± 0.017	0.963 ± 0.022	0.981 ± 0.005

TABLE 2: Comparison of median and average of initialization accuracy as well as the P values of a paired Wilcoxon rank test between each method and $M_{\text{FBA}^{\text{Aff}}}$.

	M_{ribs}	M_{capex}	$M_{\text{RBM}^{\text{Aff}}}$	$M_{\text{FBA}^{\text{Sim}}}$	$M_{\text{FBA}^{\text{Aff}}}$
Average D_{init}	0.683 ± 0.092	0.586 ± 0.168	0.527 ± 0.125	0.707 ± 0.116	$\mathbf{0.746 \pm 0.068}$
Median D_{init}	0.7018	0.6287	0.5236	0.7374	$\mathbf{0.7484}$
P value	$2.53e{-}36$	$4.01e{-}58$	$1.36e{-}63$	$6.97e{-}10$	—

TABLE 3: Comparison of median and average of final segmentation accuracy as well as the P values of a paired Wilcoxon rank test between each method and $M_{\text{FBA}^{\text{Aff}}}$. For each method, the number of cases $^{\#}$with Dice values below 0.9, 0.8, and 0.7 is provided.

	M_{ribs}	M_{capex}	$M_{\text{RBM}^{\text{Aff}}}$	$M_{\text{FBA}^{\text{Sim}}}$	$M_{\text{FBA}^{\text{Aff}}}$
Average D_{final}	0.963 ± 0.056	0.944 ± 0.134	0.964 ± 0.058	0.963 ± 0.056	$\mathbf{0.974 \pm 0.017}$
Median D_{final}	0.9784	0.9784	0.9778	0.9786	$\mathbf{0.9792}$
P value	$6.47e{-}05$	$1.18e{-}07$	$6.92e{-}04$	$6.01e{-}04$	—
Average d_a (mm)	1.350 ± 2.456	2.271 ± 6.546	1.398 ± 3.094	1.335 ± 2.493	$\mathbf{0.948 \pm 1.537}$
Median d_a (mm)	0.7342	0.7441	0.7958	0.7144	$\mathbf{0.7108}$
P value	$7.64e{-}04$	$5.85e{-}07$	$1.24e{-}03$	$4.03e{-}02$	—
#($D_{\text{final}} < 0.9$)	23	31	14	23	**3**
#($D_{\text{final}} < 0.8$)	14	23	10	11	**0**
#($D_{\text{final}} < 0.7$)	4	15	7	4	**0**

then inspected by a trained expert under the supervision of pulmonologist, and segmentation errors were manually corrected.

For performance assessment, the dice coefficient D was computed with respect to reference segmentations to measure the accuracy of model initialization (D_{init}) and of the final segmentation after the RASM and OSF segmentation steps (D_{final}). In addition, the mean absolute surface distance d_a was calculated.

5. Results

Table 1 presents the average initialization (D_{init}) and segmentation accuracy (D_{final}) obtained by our method on different test sets. Tables 2 and 3 compare the average initialization (D_{init}) and segmentation accuracy (D_{final} and d_a) between different methods, respectively. Based on the paired Wilcoxon rank test on D_{init}, D_{final}, and d_a, our method $M_{\text{FBA}^{\text{Aff}}}$ shows statistically significant improvement over all other methods. Table 3 also shows the number of cases, for each method, whose final segmentation performance is below a certain Dice value.

Figure 4 shows two examples of model initialization generated using different methods. Our method $M_{\text{FBA}^{\text{Aff}}}$ obtains initializations very close to the target while the other methods are sometimes off. Figures 5(a)–5(f) show 6 different cases where the initialization obtained by other methods causes the segmentation to be inaccurate. In all those cases, our method (Figures 5(g)–5(l)) obtains a closer initialization and good segmentation.

6. Discussion and Conclusion

Compared to the four other approaches, the proposed method delivered significantly better initializations (Table 2), which also translated into significantly better overall segmentation performance (Table 3). In addition, it produces segmentations with $D_{\text{final}} \geq 0.8$ in all cases, while all the other methods have 10 or more segmentations with $D_{\text{final}} < 0.8$. Table 1 shows that initialization performance (D_{init}) is in a close range across test sets. Due to superior initialization, the final segmentations generated by our method converge correctly to the target structure across test sets (Figures 5(g)–5(l)).

The importance of extending FBA to affine transform is underlined by its statistically significant improvement over the version using the similarity transform ($M_{\text{FBA}^{\text{Sim}}}$). Moreover, it is also superior to the affine transforms produced by the RBM method ($M_{\text{RBM}^{\text{Aff}}}$). The reason is that RBM, like most registration methods, typically works well when lung masks are available. For example, 16 out of 20 algorithms discussed in the EMPIRE challenge paper used lung masks [13]. Generating lung masks is the goal of the presented work, and thus we did not employ any masks for the RBM method. Therefore, the affine global alignment produced by RBM can be inaccurate, resulting in low initialization accuracy (Table 2) and subsequent segmentation failures (Figures 5(a) and 5(f)).

Our method produces robust ASM initializations as demonstrated by an experiment on synthetic lung data depicted in Figure 6. Despite different levels of occlusion,
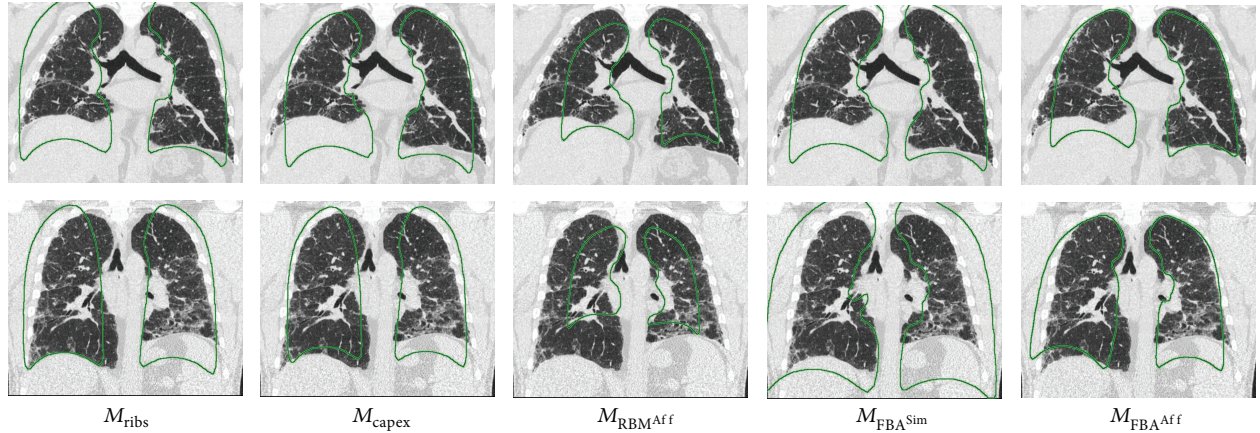
$$M_{\text{ribs}} \qquad M_{\text{capex}} \qquad M_{\text{RBM}^{\text{Aff}}} \qquad M_{\text{FBA}^{\text{Sim}}} \qquad M_{\text{FBA}^{\text{Aff}}}$$

FIGURE 4: Two examples of model initializations based on different methods. The method $M_{\text{FBA}^{\text{Aff}}}$ shows initialization quite close to the target shape whereas the other methods are off in at least one of the examples.
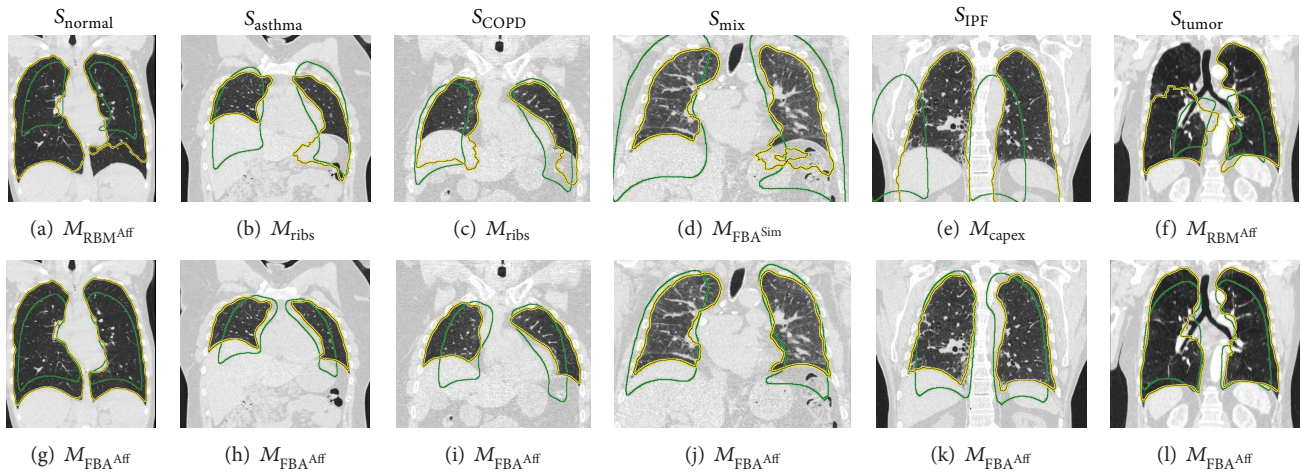

FIGURE 5: ASM initializations (green contour) and final segmentations (yellow contour) across different test sets. (a)–(f) Result of other methods for comparison with results produced with the proposed method (g)–(l).
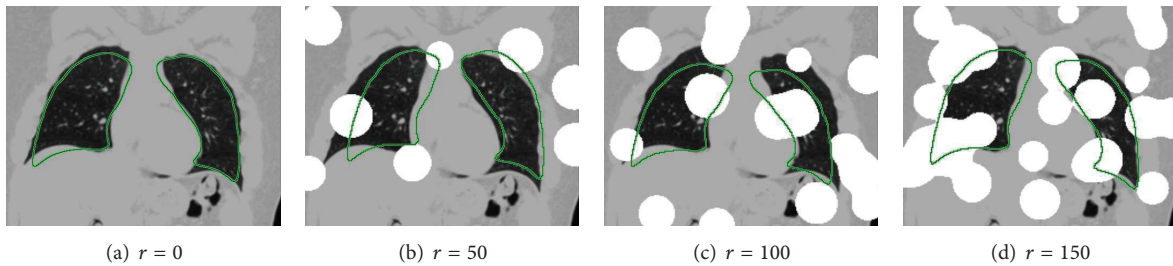

(a) $r = 0$ (b) $r = 50$ (c) $r = 100$ (d) $r = 150$

FIGURE 6: ASM initializations using $M_{\text{FBA}^{\text{Aff}}}$ on an image with varying degree of synthetic occlusion. The synthetic occlusion is generated by placing r spheres, each of radius 25 mm at random locations in the CT image. The voxels of the spheres are assigned a Hounsfield unit value of 500 HU. As can be seen, the method is robust to a large amount of occlusion.

the ASM is placed in close proximity to the actual lung in all cases. In addition, our method is fast and takes about 30 seconds on average, the majority of which is required for feature extraction. In comparison, $M_{\text{RBM}^{\text{Aff}}}$ takes 2.5 minutes, M_{ribs} 2 minutes, and M_{capex} 40 seconds on average.

We note that the mean PDM μ_{pdm} could be directly transformed to the average lung shape ξ during the training phase. However, due to subsequent affine transformation of ξ (Figure 3), the isotropic scale property of the PDM would not be preserved. Future efforts will use anisotropic scale during PDM building and RASM fitting so that all calculations can be done in the affine space.

The proposed method only requires a set of training shapes with landmarks and a set of representative features

learnt from training CT images to initialize an ASM in a new image. Consequently, our method can be generalized to other medical imaging applications that employ an ASM-based segmentation approach. Furthermore, it can be adapted for 2D and 4D applications.

Conflict of Interests

The authors declare that there is no conflict of interests regarding the publication of this paper.

Acknowledgments

References

[1] S. Sun, C. Bauer, and R. Beichel, "Automated 3-D segmentation of lungs with lung cancer in CT data using a novel robust active shape model approach," *IEEE Transactions on Medical Imaging*, vol. 31, no. 2, pp. 449–460, 2012.

[2] G. Gill, M. Toews, and R. Beichel, "An automated initialization system for robust model-based segmentation of lungs in CT data," in *Proceedings of the 5th International Workshop on Pulmonary Image Analysis*, pp. 111–122, 2013.

[3] M. Wilms, J. Ehrhardt, and H. Handels, "A 4D statistical shape model for automated segmentation of lungs with large tumors," *Medical Image Computing and Computer-Assisted Intervention*, vol. 15, part 2, pp. 347–354, 2012.

[4] M. Sofka, J. Wetzl, N. Birkbeck et al., "Multi-stage learning for robust lung segmentation in challenging CT volumes," *Medical Image Computing and Computer-Assisted Intervention*, vol. 14, no. 3, pp. 667–674, 2011.

[5] B. van Ginneken, A. F. Frangi, J. J. Staal, B. M. Ter Haar Romeny, and M. A. Viergever, "Active shape model segmentation with optimal features," *IEEE Transactions on Medical Imaging*, vol. 21, no. 8, pp. 924–933, 2002.

[6] D. K. Iakovidis, M. A. Savelonas, and G. Papamichalis, "Robust model-based detection of the lung field boundaries in portable chest radiographs supported by selective thresholding," *Measurement Science and Technology*, vol. 20, no. 10, Article ID 104019, 2009.

[7] C. Wang, S. Guo, J. Wu, Q. Liu, and X. Wu, "Lung region segmentation based on multi-resolution active shape model," in *Proceedings of the 7th Asian-Pacific Conference on Medical and Biological Engineering (APCMBE '08)*, pp. 260–263, April 2008.

[8] T. Xu, M. Mandal, R. Long, and A. Basu, "Gradient vector flow based active shape model for lung field segmentation in chest radiographs," in *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 3561–3564, 2009.

[9] B. Ibragimov, B. Likar, F. Pernuš, and T. Vrtovec, "A game-theoretic framework for landmark-based image segmentation," *IEEE Transactions on Medical Imaging*, vol. 31, no. 9, pp. 1761–1776, 2012.

[10] T. F. Cootes, C. J. Taylor, and A. Lanitis, "Active shape models: evaluation of a multi-resolution method for improving image search," in *Proceedings of the British Machine Vision Conference*, pp. 32.1–32.10, BMVA Press, 1994.

[11] M. Toews and W. M. Wells III, "Efficient and robust model-to-image alignment using 3D scale-invariant features," *Medical Image Analysis*, vol. 17, no. 3, pp. 271–282, 2013.

[12] S. Ourselin, A. Roche, G. Subsol, X. Pennec, and N. Ayache, "Reconstructing a 3D structure from serial histological sections," *Image and Vision Computing*, vol. 19, no. 1-2, pp. 25–31, 2001.

[13] K. Murphy, B. van Ginneken, J. M. Reinhardt et al., "Evaluation of registration methods on thoracic CT: the EMPIRE10 challenge," *IEEE Transactions on Medical Imaging*, vol. 30, no. 11, pp. 1901–1920, 2011.

[14] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[15] K. Mikolajczyk and C. Schmid, "Scale & affine invariant interest point detectors," *International Journal of Computer Vision*, vol. 60, no. 1, pp. 63–86, 2004.

[16] K. Mikolajczyk and C. Schmid, "A performance evaluation of local descriptors," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 10, pp. 1615–1630, 2005.

[17] M. Toews and W. M. Wells III, "SIFT-rank: Ordinal description for invariant feature correspondence," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR '09)*, pp. 172–177, Miami, Fla, USA, June 2009.

[18] M. Toews, L. Zollei, and W. M. Wells III, "Feature-based alignment of volumetric multi-modal images," in *Information Processing in Medical Imaging*, vol. 7917 of *Lecture Notes in Computer Science*, pp. 25–36, Springer, Berlin, Germany, 2013.