Published in final edited form as:

IEEE Trans Biomed Eng. 2020 April ; 67(4): 1206–1220. doi:10.1109/TBME.2019.2933508.

A Generic Approach to Lung Field Segmentation from Chest Radiographs using Deep Space and Shape Learning

Awais Mansoor,

Sheikh Zayed Institute for Pediatric Surgical Innovation, Children's National Health System, Washington DC.

Juan J. Cerrolaza,

Sheikh Zayed Institute for Pediatric Surgical Innovation, Children's National Health System, Washington DC.

Geovanny Perez,

Division of Pulmonary and Sleep Medicine, Children's National Health System, Washington, DC.

Elijah Biggs,

Sheikh Zayed Institute for Pediatric Surgical Innovation, Children's National Health System, Washington DC.

Kazunori Okada,

Computer Science Department, San Francisco State University, San Francisco, CA.

Gustavo Nino,

Division of Pulmonary and Sleep Medicine, Children's National Health System, Washington, DC.

Marius George Linguraru

Sheikh Zayed Institute for Pediatric Surgical Innovation, Children's National Health System and the School of Medicine and Health Sciences, George Washington University, Washington DC.

Abstract

Computer-aided diagnosis (CAD) techniques for lung field segmentation from chest radiographs (CXR) have been proposed for adult cohorts, but rarely for pediatric subjects. Statistical shape models (SSMs), the workhorse of most state-of-the-art CXR-based lung field segmentation methods, do not efficiently accommodate shape variation of the lung field during the pediatric developmental stages. The main contributions of our work are: (1) a generic lung field segmentation framework from CXR accommodating large shape variation for adult and pediatric cohorts; (2) a deep representation learning detection mechanism, *ensemble space learning*, for robust object localization; and (3) *marginal shape deep learning* for the shape deformation parameter estimation. Unlike the iterative approach of conventional SSMs, the proposed shape learning mechanism transforms the parameter space into marginal subspaces that are solvable efficiently using the recursive representation learning mechanism. Furthermore, our method is the first to include the challenging retro-cardiac region in the CXR-based lung segmentation for accurate lung capacity estimation. The framework is evaluated on 668 CXRs of patients between 3

month to 89 year of age. We obtain a mean Dice similarity coefficient of 0.96 ± 0.03 (including the retro-cardiac region). For a given accuracy, the proposed approach is also found to be faster than conventional SSM-based iterative segmentation methods. The computational simplicity of the proposed generic framework could be similarly applied to the fast segmentation of other deformable objects.

Keywords

Lung field; chest radiograph; deep learning; space learning; shape learning; statistical shape models

I. Introduction

Despite tremendous advancements in tomographic imaging, chest radiography remains the most commonly used imaging modality for pulmonary analysis mainly due to its low cost, low radiation dosage, and widespread availability. Radiation dosage is of particular concern in pediatric applications, especially in neonatal intensive care units where chest radiographs (CXR) are considered the first option for pulmonary diagnosis. Lung field segmentation is the necessary initial step for image-based pulmonary analysis. Accurate delineation of lung field from CXR, however, is challenging due to ambiguous boundaries, pathologies, occultation of lung field by anatomical structures in thorax, anatomical variation of lung shapes, and size across subjects (Fig. 1). Part of the challenge in developing computer-aided diagnosis (CAD) methods, especially for pediatric cohorts, is also the anatomical shape variation of lung field that occur during growth [1], [2]. As shown in Fig. 1, pediatric cohorts have a more compliant chest wall, small thoracic cage, and relative large abdominal space. Furthermore, the diaphragm of children has smaller apposition area which has a concave shape in the posterior-anterior (PA) view CXR [2]. Therefore, existing approaches to lung field segmentation that are designed primarily for adult cohorts, are not accurate at analyzing the pediatric subjects. Although a few pilot studies such as [1] have been conducted recently to look at the age-related radiological biomarkers in lungs, no comprehensive study of pediatric lung field segmentation exists to the best of our knowledge.

Traditionally, CAD algorithms designed to segment lung field from CXR ignore the retro-cardiac region, i.e., the lung region occluded by heart (Fig. 2a). The segmentation label without the retro-cardiac region provides only partial unobstructed lung field. Accurate delineation including the occluded retro-cardiac region, is necessary for correct diagnosis in diseases related to the change in lung capacity such as atelectasis (lung collapse), hyaline membrane disease, and transient tachypnea. Fig. 2c presents the correlation between the lung volume estimated from computed tomography (CT) scans and the segmented lung field area from CXR (with and without retro-cardiac region) from 108 individuals from CNHS data. The details of the data and the population are provided in Section II. The plot shows a stronger overall correlation between the lung capacity calculated including the retro-cardiac region and the lung volume obtained through CT scans ($R=0.80$ without retro-cardiac region, $R=0.86$ including retro-cardiac region; no inspiration/ expiration information was available. R is the correlation coefficient). It is important to note here that the proposed

approach is not intended to replace a functional test for lung capacity but to produce lung field boundaries that are more meaningful when comparing with boundaries obtained using 3D volumetric data.

The current CXR-based lung segmentation approaches (Table I) can be divided into three major categories: **Rule-based methods** that use predefined knowledge about the lung field to create a set of rules (e.g. intensity, edge information, etc.) for segmentation. These are usually heuristic approaches therefore subsequent refinement steps are generally needed [3]-[6].

Feature classification-based methods that formulate segmentation as a classification problem by learning the probability of every pixel (or region) belonging the lung field. The probability is calculated using a set of features extracted around the pixel being classified [7], [8]. Recently, [9] used adversarial architecture for the lung field segmentation. Adversarial networks are generally harder to train, i.e., large datasets and exhaustive parameter optimization is needed. Furthermore, as demonstrated later in Section IV (Experimental Results) ignorance of the object shape specificity results in the sub-optimal performance even by the most sophisticated feature classification-based methods.

Deformable shape model-based methods that use curves and surfaces defining the lung field that can be moved to the true boundary under the influence of *internal forces* from lung shape and *external forces* from lung appearance [11]-[13], [15], [16].

In addition, hybrid methods such as [17] and [18] cross over multiple categories. Amongst these approaches, deformable statistical shape models (SSMs) have demonstrated superior performance due to their ability to seamlessly integrate low level localized appearance features and high level global features. These models learn patterns of shape deformation from the training data of annotated images. A learned model is subsequently deformed to fit the object of interest within the test image by estimating its shape deformation patterns through an appearance-guided iterative optimization procedure. SSMs remained the workhorse for various medical image analysis applications including the lung segmentation; however, the iterative optimization is generally found to be not robust to initialization, complex background, weak edges, and contrast information. Henceforth, accurate initialization of shape models [20] and various refinements [21] remain topics of active research. In addition, conventional SSMs [22], assume a unimodal Gaussian distribution of training shapes; however, in practice, the assumption of both unimodality as well as Gaussianity may be inaccurate when the training data consist of shapes with large variation obtained from multiple cohorts, e.g., from adult and pediatric subjects (see Fig. 1e, 1d).

Contrary to SSM methods, representation learning techniques, including state-of-the-art deep learning methods, have demonstrated great potential in handling a wide range of variation including non-Gaussian and multi-modal Gaussian distributed data [23]-[25]. These techniques have also found to be robust to intensity variation and minima optimization. However, the cost of performing hypothesis testing at the atomic (pixel/voxel) level prohibits their use for large object segmentation. Furthermore, since final segmentation label using these methods are generally obtained as a concatenation of independent atomic-

level hypotheses, object shape specificity cannot be guaranteed. Shape modeling through representation learning has not garnered much attention in the past, primarily because of two reasons. First, the effective representation of a segmentation (detection+delineation) task as a learning problem is not trivial. Second, hand-crafting representation features for deformable objects is not straightforward and relies heavily on the human ingenuity [23], [24].

Recently, representation learning through deep learning (DL) has shown great promise in expanding the scope of learning algorithms to automated feature extraction. Specific to medical imaging, DL frameworks are extensively being used in various organ detection [26], classification [25], and segmentation [27] tasks. In this paper, we extend the applicability of DL to parametrized shape learning and demonstrate it via an efficient generic solution to lung field segmentation. The main contributions of our work are:

- A generic lung field segmentation framework from CXR, accommodating both adult and pediatric cohorts.
- Segmentation of the lung field including the occluded retro-cardiac region for reliable estimation of capacity and inter/intra subject comparisons.
- A DL-based mechanism for the automated detection of object of interest with large shape variation from images acquired under diverse acquisition protocols. This detection mechanism, dubbed *ensemble space learning* (ESL), also addresses the issue of error propagation to subsequent marginal spaces within the current state-of-the-art detection methods: marginal space learning (MSL) [28], [29].
- A hybrid principal component analysis (PCA)-DL based approach for including shape prior information for deformable object segmentation. This module which we call *marginal shape deep learning* (MaShDL) transforms the iterative approach of the conventional SSM-based segmentation methods to a recursive marginal refinement approach. Specifically, the method begins by learning the mode of shape deformation in the eigenspace of the largest variation and then marginally increases the dimensionality of eigenspaces by recursively including the next largest modes. As demonstrated later in the paper, the transformation allows the SSM to be posed as an efficient parameter estimation problem solvable through representation learning.

The proposed framework is evaluated using a comprehensive CXR datasets to demonstrate its potential for generic applicability.

II. Datasets and Reference Standards

Our experiments are conducted on both publicly available and in-house acquired datasets using a wide range of devices, age groups, and pulmonary pathologies. 247 publicly available radiographs from Japanese Society of Radiological Technology (JSRT; <http://www.jsrt.or.jp>) dataset and 108 from the Belarus Tuberculosis Portal (BTP; <http://tuberculosis.by>) were used. For data acquired in-house, after approval from the Internal

Review Board, 313 posterior-anterior CXRs were collected at Children's National Health System (CNHS). The subjects in the JSRT dataset have ages between 16 to 89 year (58.21 ± 14.02 year). The dataset is a standard digital CXR database with and without chest lung nodules created by the Japanese Society of Radiological Technology. The radiographs had dimensions of 2048×2048 pixels, spatial resolution of 0.17×0.17 mm/pixels, and digital resolution of 12 bits. BTP images, from patients between 18 to 86 year (45.60 ± 16.98 year), had dimensions of 2248×2724 pixels, spatial resolution of 0.16×0.16 mm/pixel and the digital resolution of 12 bits. The dataset consists of CXRs obtained from patients diagnosed with or suspected of multi-drug-resistant tuberculosis (MDR-TB). The CXR findings of these patients include consolidation, cavitory lesions, nodules, pleural effusion, pneumothorax, and fibrotic scars. For CNHS data, patients having ages between 3 months to 18 year (4.75 ± 5.30 year) with viral chest infections were scanned. The dataset consists of radiographs collected from individuals having or suspected of having either Human metapneumovirus (hMPV) or rhinovirus. The radiological symptoms to these viruses include acute respiratory infections, chronic lung conditions, chest wall deformities, cardiovascular anomalies. The radiographs have dimensions within the range $(660 - 4240) \times (987 - 4240)$ pixels with spatial resolution ranges between 0.1×0.1 mm/pixel to 0.14×0.14 mm/pixel and a digital resolution of 12 bits. The paired CT data has a resolution of $0.62 \times 0.62 \times 0.62$ mm/pixel and a digital resolution of 16 bits. For CNHS data, patients having ages between 3 months to 18 year (4.75 ± 5.30 year) with viral chest infections were scanned. The dataset consists of radiographs collected from individuals having or suspected of having either Human metapneumovirus (hMPV) or rhinovirus. The radiological symptoms to these viruses include acute respiratory infections, chronic lung conditions, chest wall deformities, cardiovascular anomalies. For consistency of training data, all scans from the three datasets were resized to 2048×2048 pixels using B-spline interpolation. The ground truth labels both including and excluding the retro-cardiac region were prepared by two fellows using the ITK-SNAP interactive software under the supervision of two expert pulmonologists. For ground truth labels including the retro-cardiac region, an overall interobserver agreement of 0.95 ± 0.03 was observed; specifically, 0.94 ± 0.02 for CNHS data and 0.96 ± 0.03 for the JSRT and BTP data was estimated. Ground truth labels excluding the retro-cardiac region were prepared for comparative purposes with the state-of-the-art methods. To construct the statistical shape model, 144 boundary points (72 per left/right lung) with anatomical correspondences are annotated. Specifically, six manually annotated primary landmarks were initially obtained for each lung based on their distinctive anatomical appearance and ability to roughly define the shape of lung. Subsequently, equidistant secondary landmarks were estimated along the lung contour using interpolation between the primary landmarks. In order to make sure that no loss in the segmentation label accuracy has occurred due to the interpolation, the accuracy of the proposed interpolation method was evaluated using the Dice coefficient score (DCS) between the manual ground truth and the landmark-based interpolated contour. A mean DCS of 0.9942 ± 0.0013 was obtained for our dataset. Manual landmarks are needed only for the training data, no manual landmarking is required for test or validation data.

III. Methods

A. Overview

Fig. 3 shows the flow diagram summarizing the proposed framework. The segmentation of a deformable object (lung field) is performed by learning space (localization) and shape parameters using two separate DL architectures. As demonstrated later in the manuscript, the presented DL-based approach for shape parameters learning is theoretically equivalent to the one adopted by conventional SSM techniques: estimating the shape parameters of the object of interest under constraints on shape model and appearance. However, unlike the iterative convergence approaches of conventional SSMs that optimize the entire shape parameter space simultaneously, the proposed method transforms the parameter space into linearly independent subspaces and employs a battery of DL classifiers to learn the shape parameters individually. This marginal learning of independent parameter subspaces makes our approach both computationally tractable as well as significantly more accurate compared to the state-of-the-art SSM approaches. Herein, we introduce a generic method for space and shape parameters learning of deformable objects, which we apply later to the lung field segmentation from CXR.

B. Theoretical Background

Among the approaches for deformable shape representation presented in the literature [30], PCA-based SSM [22] has been found to be most successful due to their simplicity, performance, and compact representation. These models have been widely used to deform an initial estimate of shape (mostly the mean shape of the object of interest obtained using training data) under the guidance of appearance-based image evidence (*external forces*) and shape priors (*internal forces*). SSM uses an explicit point-based representation in which each shape is described by M points (or landmarks) distributed across the contour. Given a set of N aligned shapes $\{\mathbf{x}_n\}_{n=1}^N$ in 2D, the SSM is defined using a mean shape $\bar{\mathbf{x}} \in \mathbb{R}^{2M \times 1}$, a set of K eigenvectors $\{\mathbf{p}_k\}_{k=1}^K$, and a set of corresponding eigenvalues, $\{\lambda_k\}_{k=1}^K$, obtained by applying PCA to the aligned shapes. The magnitude is proportional to the shape variance explained by the particular eigenvector. K is generally chosen to be the smallest number of modes such that their cumulative variance explains a sufficiently large proportion (normally 95% – 98%) of the total variance explained by all M eigenvectors (usually $K \ll M$). Subsequently, any shape $\mathbf{X} \in \mathbb{R}^{2M \times 1}$ in the non-aligned image space can be approximated using the anisotropic similarity transform parameters (presented below), the aligned mean shape $\bar{\mathbf{x}}$, and the weighted sum of K largest modes (eigenvectors).

$$\mathbf{X} \approx A_{\text{space}}(\bar{\mathbf{x}} + \mathbf{pb}), \quad (1)$$

where $A_{\text{space}} = \begin{bmatrix} \mathbf{S}\mathbf{0} & \mathbf{T} \\ \mathbf{0}^T & 1 \end{bmatrix}$ is an invertible matrix called the anisotropic similarity transform matrix. The matrix transforms the mean shape from the aligned shape space to a non-aligned image space, using specifically, position: $\mathbf{T} = \{\mathbf{T}_x, \mathbf{T}_y\}$, orientation: $\boldsymbol{\theta}$, and anisotropic scale: $\mathbf{S} = \{\mathbf{S}_x, \mathbf{S}_y\}$. $\mathbf{b} \in \mathbb{R}^{K \times 1}$ is the shape weight matrix. Given A_{space} and $\bar{\mathbf{x}}$, the weight parametrization of new target shape can be obtained using (1) as $\mathbf{b} = \mathbf{p}^T (A_{\text{space}}^{-1} \mathbf{X} - \bar{\mathbf{x}})$. The

legitimacy of the estimated shape is generally guaranteed by imposing individual constraints on each weight. [22] demonstrated that the suitable constraints on the weights are typically of the order of $-3\sqrt{\lambda_k} \leq \mathbf{b}_k \leq +3\sqrt{\lambda_k}$. After initializing the similarity parameters described by A_{space} , SSM iteratively adjusts the deformable shape until convergence, causing the points of $\mathbf{X}$ to move under the influence of object model and image evidence. The weight vector $\mathbf{b}$ after t iterations is

$$\mathbf{b}^{(t)} \leftarrow \mathbf{b}^{(t-1)} + d\mathbf{b}^{(t)}, \quad (2)$$

where $d\mathbf{b}^{(t)} = \mathbf{P}^T d\mathbf{X}$ is the change in the model parameters at the iteration t . Using eq. (1) and (2), the shape parameter vector Ω for any deformable shape in 2D can be written as

$$\Omega = \left\{ \underbrace{\mathbf{T} \in \mathbb{R}^2, \boldsymbol{\theta} \in \mathbb{R}, \mathbf{S} \in \mathbb{R}^2}_{\text{Space Parameters}}, \underbrace{\bar{\mathbf{x}} \in \mathbb{R}^{2M \times 1}}_{\text{Mean Shape}}, \underbrace{\mathbf{b} \in \mathbb{R}^{K \times 1}}_{\text{Shape Parameters}} \right\} \quad (3)$$

C. Parametrized Shape Representation and Learning

Given a shape parameter vector Ω and a set of 2D training images $\{f_n\}_{n=1}^N$, each having dimensions (N_x, N_y) : $f(nx, ny) \in \mathbb{R} : 0 \leq nx \leq N_x - 1, 0 \leq ny \leq N_y - 1$ (where $f(nx, ny)$ denotes intensity at location (nx, ny)), a representation classifier can be learned that can estimate the correct parameter vector, Ω , by maximizing the following posterior probability over a valid parameter space:

$$\hat{\Omega} = \{ \underbrace{\hat{\mathbf{T}}, \hat{\boldsymbol{\theta}}, \hat{\mathbf{S}}}_{\Omega_{\text{Space}}}, \underbrace{\bar{\mathbf{x}}}_{\Omega_{\text{Mean Shape}}}, \underbrace{\hat{\mathbf{b}}}_{\Omega_{\text{Shape}}} \} = \underset{\Omega}{\operatorname{argmax}} p(\Omega | f). \quad (4)$$

However, due to the large number of testing hypotheses in eq. (4), learning a full space classifier with efficiency comparable to the traditional SSM-based iterative segmentation methods is challenging and requires a large amount of training data [23]. A few attempts have been made in the past for efficient parameter learning by partitioning the parameter space into linearly or marginally independent subspaces. For instance, [23] proposed an efficient method, MSL, for object detection by training classifiers to learn Ω_{Space} . Since its introduction, MSL has been successfully applied in various medical imaging applications such as segmentation of heart [23], left ventricle detection [31], mid-sagittal plane detection [32], and standard echocardiographical plane detection [33]. MSL learns classifiers in marginally independent parameter subspaces. Their work suggested that the dimensionality of effective parameter space can be significantly contracted by separating conditionally independent parameters into semigroups (translations, scales, and orientations). A semigroup is an algebraic structure consisting of a set with an associative binary operation. According to MSL, the object detection approach can be expressed as the maximization of posterior probability of semigroup Ω_{Space} ,

$$p(\Omega_{\text{Space}} | f) = \underset{\Omega_{\text{Space}}}{\operatorname{argmax}} p(\mathbf{T} | f) p(\boldsymbol{\theta} | \mathbf{T}, f) p(\mathbf{S} | \mathbf{T}, \boldsymbol{\theta}, f). \quad (5)$$

Extending the concept, we propose that the posterior probability of the semigroup Ω can be similarly approximated as the maximization of the marginal probabilities of its semisubgroups: Ω_{Space} and Ω_{Shape} ,

$$\arg\max_{\Omega} p(\Omega \mid f) = \arg\max_{\Omega_{\text{Shape}}} p(\Omega_{\text{Shape}} \mid \Omega_{\text{Space}}, f) \times \arg\max_{\Omega_{\text{Space}}} p(\Omega_{\text{Space}} \mid f), \quad (6)$$

However, in contrast to eq. (5) and the MSL framework proposed in [23] that does not impose any commutativity constraints, our proposition in eq. (6) is subject to an assertion that the parameter vector Ω can be estimated marginally only as a nowhere-commutative semigroup: $\{\Omega_{\text{Space}}, \Omega_{\text{Shape}} \mid \Omega_{\text{Space}} \cdot \Omega_{\text{Shape}} \neq \Omega_{\text{Shape}} \cdot \Omega_{\text{Space}}\}$. A nowhere commutative semigroup is any semigroup S , such that for all a and $b \in S$ if $ab = ba$ then $a = b$. The nowhere-commutativity is enforced since, as discussed before, within the context of SSM-based methods (eq. (1)), image-aligned mean shape ($\bar{\mathbf{X}} = A_{\text{space}} \bar{\mathbf{x}}$) serves as the initialization for the shape deformation. During the iterative process of (2), this initial estimate is continuously refined till convergence.

The marginal parameter space simplification introduced in eq. (6) is mainly intended to improve the computational cost of a classifier-based approach to parameter estimation. However, despite this simplification, the proposed classifier-based framework meet or exceed the iterative refinement-based SSM alternatives in terms of segmentation accuracy as demonstrated later in Section IV-B. The semisubgroups Ω_{Space} and Ω_{Shape} can be further partitioned till the trivial semisubgroup level, i.e., a semigroup with one element only,

$$\begin{aligned} \widehat{\Omega}_{\text{Space}}^c &= \{\widehat{\mathbf{T}}, \widehat{\boldsymbol{\theta}}, \widehat{\mathbf{S}}\} = \arg\max_{\Omega_{\text{Space}}} p(\Omega_{\text{Space}} \mid f) \\ &= \underbrace{p(\mathbf{T} \mid f)}_{\text{Position}} \times p(\mathbf{S} \mid \mathbf{T}, f) \times p(\boldsymbol{\theta} \mid \mathbf{T}, \mathbf{S}, f), \end{aligned} \quad (7)$$

Position-Scale

and $\widehat{\Omega}_{\text{Shape}}$

$$\begin{aligned} \widehat{\Omega}_{\text{Shape}}^{nwc} &= \{\widehat{\mathbf{b}}_k : \forall k = 1, \dots, K\} = \arg\max_{\Omega_{\text{Shape}}} p(\Omega_{\text{Shape}} \mid \Omega_{\text{Space}}, f) \\ &= \underbrace{p(\mathbf{b}_1 \mid \Omega_{\text{Space}}, f)}_{\text{Largest mode of shape deformation}} \times \dots \times p(\mathbf{b}_K \mid \mathbf{b}_{K-1}, \dots, \mathbf{b}_1, \Omega_{\text{Space}}, f), \end{aligned} \quad (8)$$

K largest modes of shape deformation

where c and nwc denote commutative and nowhere-commutative semisubgroups respectively. Also note that $p(\mathbf{b}_0 \mid \Omega_{\text{Space}}, f) = \bar{\mathbf{x}}$. Eq. (7) and (8) suggest that by splitting the semigroups Ω to commutative and non-commutative nontrivial semisubgroups, a $5 + K$ dimensional learning space can be approximated to a concatenation of $5 + K$ one dimensional subspaces; therefore, reducing the computational complexity of the manifold

[28], [29]. Individual classifiers can be trained subsequently for independent subspaces, thus simplifying training and reducing the amount of data needed to train the classifier.

D. Deep Learning Network for Space and Shape Parameter Estimation

Our proposed DL framework for learning the parameters Ω consists of two main layers: an unsupervised stacked denoising autoencoder (SdAE) layer for pre-training to initialize the weights of feed forward deep neural network (DNN) and a supervised DNN layer for fine-tuning. Unsupervised pre-training to initialize the weights of DNN has demonstrated to have better convergence properties especially if the labeled training data is not very large [34].

Once the layers are pre-trained using SdAE, the weights and biases of the encoder layer are used to initialize the feed forward DNN. This network architecture is subsequently used for learning space and shape parameters in our DL framework. In our experiments, we do not find any significant difference in performance with and without pre-training. However, pre-training will make significant difference both in terms of accuracy and convergence when new data acquired using significantly different protocol will be included.

For greater details on the training of DNNs and SdAE, readers are encouraged to review [35]. Specific details of the network configuration pertaining to learning space and shape parameters are presented in Sections III-E and III-G respectively.

E. Space Parameters Estimation

MSL and MSDL, the current state-of-the-art learning-based techniques for space parameters estimation ($\hat{\Omega}_{\text{Space}}$), has been found to be very successful in various medical imaging applications [24], [36]. Both approaches solve the same classification problem using two different classification techniques. MSL uses the probabilistic boosting tree classifier while MSDL adopts the deep neural network for the parameter estimation. Both MSL and MSDL are initialized using a bounding box of arbitrary parameters (Fig. 4a). Later these parameters are marginally refined (translation followed by orientation followed by scale). The marginal refinement transforms the arbitrary bounding box into a minimum area bounding box enclosing the object of interest. The sequential parameter learning within MSL, however, results in the propagation of estimation error to successive stages. Specifically, the error in the translation estimation propagates to orientation and scale estimations. Consequently, the cumulative estimation error at a given stage is lower-bounded by the cumulative error at the previous stages,

$$\begin{aligned} \inf_{\mathbf{T}} \left(\Omega_{\text{Translation}} - \arg\max_{\mathbf{T}} p(\mathbf{T} | f) \right) &\rightarrow 0, \\ \inf_{\mathbf{T}, \boldsymbol{\theta}} \left(\Omega_{\text{Translation-Orientation}} - \arg\max_{\mathbf{T}, \boldsymbol{\theta}} p(\mathbf{T}, \boldsymbol{\theta} | f) \right) &\rightarrow \mathbf{T} - \hat{\mathbf{T}}, \\ \inf_{\mathbf{T}, \boldsymbol{\theta}, \mathbf{S}} \left(\Omega_{\text{Space}} - \arg\max_{\mathbf{T}, \boldsymbol{\theta}, \mathbf{S}} p(\mathbf{T}, \boldsymbol{\theta}, \mathbf{S} | f) \right) &\rightarrow \boldsymbol{\theta} - \hat{\boldsymbol{\theta}}. \end{aligned} \quad (9)$$

From eq. (9), the domain normalized error $\mathbf{S} - \hat{\mathbf{S}} \geq \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \geq \mathbf{T} - \hat{\mathbf{T}}$ (further explanation on the propagation of error is provided in Section IV-A). Moreover, since MSL and MSDL are

based on using a minimum area bounding box, deciding the optimal initialization values of similarity transform parameters (Ω_{Space}) for the bounding box is generally not trivial, especially in data with large variation. To address these challenges, we propose ESL that learns Ω_{Space} by transforming it from being a marginally independent semigroup of parameters (as described in MSL, eq. (9)), to a linearly independent semigroup of surrogate parameters. Specifically, instead of estimating Ω_{Space} using the minimum area bounding box, ESL estimates them as a function of four linearly independent vertices of two sets of parallel lines bounding the object of interest. Fig. 4 graphically illustrates the methodological differences between MSL and ESL for the specific application of lung field segmentation. The number of classifiers used by the ESL and MSL are the same; however, the motivation behind ESL is to stop the error by one classifier affect the performance of the second classifier. Specifically, given a pair of parallel bounding lines $l_{\{1,2\}} \in \mathbb{R}^2$, and a second pair of lines $l_{\{3,4\}} \in \mathbb{R}^2$ perpendicular to $l_{\{1,2\}}$, the four intersecting vertices provide the estimation of translation ($\mathbf{T}$) and the scale ($\mathbf{S}$) of the minimum area bounding box enclosing the object of interest (lung field). The box of estimated translation and scale is subsequently used to estimate the orientation (θ). Unlike the MSL, no assumption on the initial values of parameters is needed in the ESL. Moreover, since the parameters of ESL are linearly independent (i.e., $p(l_i | f) = p(l_i | f)$, $\forall i, j$; in MSL ($p(\mathbf{S} | \mathbf{T}, f) = p(\mathbf{S} | f)$)) therefore, $p(\Omega_{\text{Space}} | f) = \mathcal{L}\{\prod_{i=1}^4 p(l_i | f)\} \times p(\theta | \mathbf{T}, \mathbf{S}, f)$, where $\mathcal{L}\{\}$ denotes the sequence of geometrical operations to extract $\mathbf{S}$ and $\mathbf{T}$ from the four estimated vertices. Similar to eq. (9), the lower-bounds on cumulative estimation error for the space parameters using the ESL is,

$$\begin{aligned} \inf \left(\Omega_{\text{Translation}} - \arg\max_{\mathbf{T}} p(\mathbf{T} | f) \right) &\rightarrow 0, \\ \inf \left(\Omega_{\text{Translation-Scale}} - \arg\max_{\mathbf{T}, \mathbf{S}} p(\mathbf{T}, \mathbf{S} | f) \right) &\rightarrow 0, \\ \inf \left(\Omega_{\text{Space}} - \arg\max_{\mathbf{T}, \theta, \mathbf{S}} p(\mathbf{T}, \theta, \mathbf{S} | f) \right) &\rightarrow l_i - \hat{l}_i. \end{aligned} \quad (10)$$

Since the orientation is estimated independently and PA CXRs are acquired under a position protocol (upright), pairs $l_{\{1,2\}}$ and $l_{\{3,4\}}$ can be assumed to be parallel to the horizontal axis and the vertical axes of the image respectively for simplicity (Fig. 4b). Therefore, for the pairs of lines bounding the object of interest parallel to the horizontal ($i \in \{1, 2\}$) and vertical ($i \in \{3, 4\}$) axes, the bounding line estimation problem is reduced to estimating two pairs of x-intercepts (lines bounding the object vertically: l_1, l_3) and y-intercepts (lines bounding horizontally: l_2, l_4).

1) Bounding Lines Estimation:

Training: Four separate DL classifiers are trained for the four bounding lines. Since the boundary line can take only take position from a discrete set of pixels; therefore, we are treating our problem as a classification task. If the intent would have been to estimate the location at sub-pixel level, in that case it would have been more appropriate to use regression framework. To provide contextual information to the classifier, an image patch of size $(2r + 1, N_y)$ or $(N_x, 2r + 1)$ are extracted around each line: $\{l_i\}_{i=1}^4 \pm r$ (see Fig. 4b). A positive

hypothesis for a line l_i is formulated to find the horizontally (or vertically) oriented bounding box centered at position $(\{l_i\}_i = \{1, 3\}, \frac{N_y}{2})$ (or $(\frac{N_x}{2}, \{l_i\}_i = \{2, 4\})$) respectively,

$$|l_i - \hat{l}_i| \leq 1 \text{ pixels}, \quad (11)$$

where l_i and $\hat{l}_i$ denote the ground truth and the hypothesized position of the line i respectively. Similarly, a negative sample satisfies:

$$|l_i - \hat{l}_i| \geq 5 \text{ pixels}, \quad (12)$$

The separation in the positive and negative hypotheses is intended to provide a clean split between the training hypotheses.

Classifier Architecture: The intensities of set of positive class image patches (satisfying eq. (11)) and negative class image patches (satisfying eq. (12)) are first normalized to [0, 1] range and then stacked together to train the framework presented in Sec. III-D. As mentioned in Section II, the digital resolution in all three datasets used in our experiments are 12 bits (4096 gray-levels, DICOM tag= (0028,0101); unsigned, DICOM tag= (0028,0103)); therefore, the CXR intensities are divided by 4096 to achieve normalization. For training datasets acquired under different protocols, corresponding DICOM tags can be used to decide the normalization. Moreover, in our experiments, r is set to 7 based on performance accuracy and efficiency ($r = 3 \dots 13$ was tested); therefore, each training patch has dimensions of $N_x \times (2 * r + 1) = 2048 \times 15$ or $(2 * r + 1) \times N_y = 15 \times 2048$. The dimension of the multiple layer SdAE is $30720 \times 800 \times 400$. For SdAE, we use the sigmoid activation function, learning rate of 0.001, batch size of 1000, and 100 epochs. For DNN we use the sigmoid activation, batch size of 1000, learning rate of 0.1, and 100 epochs. Again, the parameters of the network are empirically estimated to minimize the reconstruction error. Furthermore, the number of layers has been decided empirically: the layers are added to the network until the reconstruction error stops decreasing. The proposed deep learning architecture for line estimation is shown in Fig. 5a.

Hypothesis Testing: Each pixel row (or column) along the axis (as shown in Fig. 5a) is tested for the line position using the trained classifiers. Similar to the practice adopted in [28], the position of the line is determined by taking the mean of the top candidates (10 in our experiments) with the highest score in order to make the framework robust to classification noise. Both mean and median of the top candidates were tried in our experiments; however, we do not find any significant advantage in performance of using one approach over the other. The four intersecting vertices of bounding lines are used to extract $\hat{\mathbf{T}}$ and $\hat{\mathbf{S}}$ using a sequence of well-known geometrical operations.

2) Orientation Estimation:

Training: The orientation estimation hypothesis in ESL is formulated as: finding the object of interest with centroid at position $\mathbf{T}$ having anisotropic scale $\mathbf{S}$ and orientation θ . Using a bounding box with position and anisotropic scale already estimated, the hypotheses for

orientation estimation is generated by rotating the bounding box of size (**S**) around (**T**). The position-scale-orientation (anisotropic similarity) hypothesis is positive if, in addition to satisfying (11) for all four lines, it also satisfies: $|\theta - \hat{\theta}| \leq \Delta\theta^+$ rad, where θ denotes the orientation of the bounding box encapsulating the ground-truth label and $\hat{\theta}$ is the hypothesized orientation. A negative hypothesis satisfies $|\theta - \hat{\theta}| \geq \Delta\theta^-$ rad ($\Delta\theta^+ \neq \Delta\theta^-$). In our experiments, we use $\theta^+ = 0.017$ rad (1 degree) and $\theta^- = 0.034$ rad.

Classifier Architecture: For computational efficiency and feature uniformity, the extracted patches using the oriented bounding box are resized to 64×64 pixels using B-spline interpolation. The proposed architecture and its configuration for orientation estimation are shown in Fig. 5b. The architecture for SdAE and DNN uses same hyper parameters as the bounding line classifiers (Section III-E1).

Hypothesis Testing: A bounding box with estimated position and anisotropic scale from Section III-E1 is rotated with a step-size of 0.0017 radians. Subsequently, trained orientation classifier is used to calculate the similarity scores for each rotated hypothesis. The final estimate is obtained using the average of top 10 candidates.

For difficult to detect objects with inconsistencies in acquisition protocols, classifier used for space parameter estimation can be applied multiple times in a cyclic manner till convergence to obtain the minimum area bounding box. However, since in our application, the CXR were acquired under a standard protocol, this was not required. Moreover, the subsequent module for shape estimation (MaShDL) that is based on the theory statistical shape models have demonstrated to be robust to slight variation in space parameters in our experiments and in the literature [37].

F. Optimal Mean Shape Determination

For optimal performance using SSM-based segmentation methods, the mean contour shape has to be initiated as close to the true boundary as possible [37]. As the anatomical structure of the lung evolves with age, resulting in shape variation amongst various age groups; therefore, we evaluate a multiple shape modeling approach for our generic framework. Although cluster sizes were tried in our experiments, based on maximum likelihood estimation of the Gaussian mixture model (GMM) clustering of the aspect-ratios ($\hat{S}_x / \hat{S}_y$; Fig. 1d); the optimal number of shape models for our dataset is determined to be two for our training dataset based on shape variation and the size of training data. This is also evident through Fig. 1d. The aspect-ratio is also found to be strongly correlated with the shape variation of modes weighted by eigenvalues ($R=-0.945$). However, depending upon the size and diversity of the training data, the number of models can be considered a hyper-parameter of the proposed framework. Finally the use of GMMs to define the number of models can be generalized to other unsupervised clustering methods specific to other intended applications and datasets of the proposed framework. As demonstrated in Fig. 1e, patient information such as age can be used for clustering; however, such information is not readily available for anonymized and publicly available data. Therefore, the training data is partitioned into two groups based on the aspect-ratio.

Training: Let $\{\mathbf{x}_n^{(i)}\}_{n=1}^{N_i}$ define a set of N_i training shapes for the group $i \in \{1, 2\}$ then the optimal mean shape $\bar{\mathbf{x}}_{\text{opt}}^{(i)}$ for that group are obtained iteratively by minimizing the following residual error after generalized Procrustes alignment of N_i group training shapes,

$$\hat{\mathbf{x}}_{\text{opt}}^{(i)} = \underset{\bar{\mathbf{x}}^{(i)}}{\operatorname{argmin}} \sum_{n=1}^{N_i} \sum_{m=1}^M \|\mathcal{J}_n(\bar{\mathbf{x}}^{(i)}(m)) - \mathbf{x}(m)_n^{(i)}\|^2, \quad (13)$$

where $\mathcal{J}_n()$ denotes the generalized Procrustes transformation from the mean shape $\bar{\mathbf{x}}^{(i)}$ to a training shape $\mathbf{x}_n^{(i)}$.

Hypothesis Testing: The appropriate shape model for the test image is chosen based on the estimated aspect-ratio ($\hat{\mathbf{S}}_x / \hat{\mathbf{S}}_y$). A threshold of 1.22 for the aspect-ratio is empirically determined to decide between appropriate shape model for the test image.

G. Shape Parameter Estimation

The concept of using representation learning methods for SSM is not novel. A few attempts have already been made in the literature such as [28] where irregular sampling patterns were used to capture the shape deformation followed by Haarwavelet feature extraction. However, the need for extracting optimal hand-crafted features, the amount of training data needed to learn shape parameters simultaneously, as well as the computational complexity of the multi-parameter classifier made representation learning methods a less attractive choice compared to the conventional iterative optimization techniques. Our proposed approach, Marginal Shape Deep Learning (MaShDL), attempts to address these challenges. To learn Ω_{Shape} , MaShDL adopts a recursive rather than an iterative approach adopted by the conventional SSM (eq. (2)) [22]. Specifically, instead of estimating and optimizing all K modes collectively (eq. (1)); MaShDL refines the aligned mean shape by recursively adding finer modes. This modification simplifies the hypothesis space by letting separate classifiers trained for each mode. From eq. (1), (2), and (8), the estimated aligned shape $\mathbf{x}_k$ using the k largest modes can be recursively written in terms of the aligned shape $\mathbf{x}_{k-1}$, obtained using the $(k-1)$ largest modes,

$$\mathbf{x}_k = \mathbf{x}_{k-1} + \mathbf{p}_k \mathbf{b}_k, \quad \mathbf{x}_0 = \bar{\mathbf{x}}^{(i)}, \quad (14)$$

where $\mathbf{p}_k$ is the k^{th} eigenvector and $\mathbf{b}_k$ is the corresponding weight. It is important to mention here that modes and mean shape in eq. (14) are based on the grouping performed in section (III-F), the superscript (i) is dropped for the ease of reading. Eq. (14) transforms (1) from block parameter estimation of the modes (as performed in eq. (2)) to recursive estimation by successively adding the next lower order mode. Moreover, eq. (8) and (14) imply that Ω_{Shape} is nowhere commutative subgroup since the k^{th} largest mode of variation has to be estimated prior to $(k+1)^{\text{th}}$ mode. Therefore, parameter estimation through representation learning in MaShDL starts with the most informative (highest) mode and sequentially adds lower variability modes.

Training: MaShDL begins by learning the highest mode through deforming the mean shape (or the zeroth mode: $\mathbf{x}_0$) that is subsequently deformed by the second highest mode and so on.

Positive Shape Hypotheses: The positive shape hypothesis for all modes are the same. The positive hypothesis corresponds to extracting $q \times q = 15 \times 15$ patches around these M landmarks (=144 (72 per lung)) of manually delineated ground truth shape.

Negative Shape Hypotheses: The negative hypotheses for the k^{th} mode are fabricated as follows:

- a. Use eq. (1) and the mean shape $\bar{\mathbf{x}}^{(i)}$ to estimate the K “true” modes of variation of n_i^{th} shape in the training set.
- b. To generate negative hypotheses for the k^{th} mode of the n_i^{th} shape, generate a set of synthesized shapes by keeping the $(k - 1)$ largest estimated modes from (a) constant and varying only the k^{th} mode within the range $(\hat{\mathbf{b}}_k \leq \lfloor 3\lambda_k \rfloor)$. Henceforth, a negative hypothesis should satisfy eq. (7) and

$$\begin{aligned} \hat{\mathbf{x}}_k &= \hat{\mathbf{x}}_{k-1} + \mathbf{p}_k \hat{\mathbf{b}}_k, \\ |\mathbf{b}_k - \hat{\mathbf{b}}_k| &\geq 0.25, \forall \hat{\mathbf{b}}_k \in [-3.0, \dots, +3.0]\lambda_k \end{aligned} \quad (15)$$

$\mathbf{b}_k$ is the value of the k^{th} mode obtained using eq. (2). $|\mathbf{b}_k - \hat{\mathbf{b}}_k| \geq 0.25$ translates, in our application, to a minimum landmark-to-landmark distance of 2 pixels.

- c. Extracting $q \times q$ patches (shown as squares in Fig. 5c) around M points of the shape synthesized.

Fig. 6 shows examples of positive (in green) and negative hypotheses (in red) for the four highest modes. Each hypothesis corresponds to a shape depicted by the concatenation of patches of size $q \times q$, each extracted around the M (red or green) landmark points. The extracted hypotheses are subsequently used to train a classifier for the k^{th} deformation mode. Similar to conventional SSM, our framework uses local appearance information to move the object boundary to the optimal position.

Classifier Architecture: In our experiments, identical patch sizes ($q = 15$) are used for training classifiers for all modes ($q = 3 \dots 21$ were tested). Smaller patch sizes are found to be prone to noise while the higher sizes tend to miss subtle shape deformations. The image patches extracted around every landmark point are subsequently stacked together in a specific order (Fig. 5c) to form a single hypothesis. Each training hypothesis has dimensions of 172×225 pixels. A multiple layer ($38700 \times 1600 \times 800$) SdAE followed by DNN is used (shown in Fig. 5c). For SdAE, sigmoid activation function, learning rate of 0.001, batch size of 1000, and 100 epochs are used.

Hypothesis Testing: The optimal mean shape (obtained in Section III-F) is first aligned to the detected object in the test image using $\hat{\Omega}_{\text{Space}}$. Next, the trained classifier for the largest mode of shape variation is used to deform the aligned mean shape $\bar{\mathbf{X}}$ followed by classifier trained for the second highest mode and so on. The process is iterated to estimate the next highest variation mode until a cumulative energy of 95%, which, in our application, is equivalent to using the largest fifteen modes of variation are included. Limiting the number of modes is a common practice when creating PCA-based statistical shape models [22]. Although there is no theoretical limitation on learning all M modes using MaShDL, a larger training dataset is generally needed to train classifiers for lower-ranked modes due to the increasing subtlety between the positive and negative hypotheses with the number of modes. Moreover, it is also predicted that both the total number of estimate-able modes as well as the machine-discernibility of adjacent modes are correlated with both digital and spatial image resolution.

H. Data Augmentation

Since the number of positive hypotheses in our training routine is smaller than the number of negative hypotheses, a data augmentation approach similar to the one presented in [38] along with random sampling is adopted to balance the samples prior to training. Specifically, we used two forms of data augmentation: (1) geometrical augmentation, and (2) appearance augmentation. The geometrical augmentation consists of generating horizontal and vertical reflections of hypotheses while the appearance augmentation consists of slightly altering the intensities in the training images. For intensity alteration we first perform PCA over the entire training dataset ($\{f_n\}_{n=1}^N$). Subsequently, to the n^{th} training normalized image (0, 1), we add the multiples of the extracted principal component with magnitude proportional to the corresponding eigenvalue times a random variable drawn from a Gaussian distribution with zero mean and 0.1 standard deviation, i.e.,

$$[\mathbf{p}_1^f, \mathbf{p}_2^f, \dots, \mathbf{p}_N^f][\alpha_1 \lambda_1^f, \alpha_2 \lambda_2^f, \dots, \alpha_N \lambda_N^f]^T$$

where $\mathbf{p}_n^f$ and λ_n^f denotes the n^{th} eigenvector and eigenvalue respectively. The superscript f is added to denote the training image data and to differentiate them from eigenvalues and eigenvectors of training shape data defined in section III-C. The drawn random variable $\alpha_n \sim \mathcal{N}(0, 0.1)$ is applied to every pixel of the n^{th} training image. Same augmentation scheme is applied to the hypotheses for every classifier in our framework.

IV. Experimental Results

The performance of the proposed framework and its individual modules (ESL, MaShDL) was evaluated using five-fold cross-validation. All three datasets (JSRT, BTP, CNHS) were evenly divided into five sets for training and validation then the results were averaged over the five validation rounds.

A. Space Parameter Estimation: MSL vs. ESL

The performances of ESL and MSL methods were compared using the DL extension of MSL [24]. Furthermore, parameters in the original MSL were reordered from eq. (5) for a more meaningful comparison with ESL: translation followed by scale and orientation estimation respectively. Fig. 7 presents the estimation error in translation and scale using MSL and ESL. The estimation error in translation was 4.42 ± 2.25 mm with the ESL compared to 5.62 ± 3.62 mm using MSL (p -value < 0.001; Wilcoxon rank sum test). For scale estimation, an average error of 3.99 ± 2.97 mm using ESL was obtained compared to 28.09 ± 10.77 mm using MSL (p -value < 0.001). Although both ESL and MSL follow the same mechanism for orientation estimation, as predicted in eq. (9), due to the accumulation of error from translation and scale, MSL achieves an orientation error of 0.11 ± 0.09 radians, which was significantly worse than the one obtained using ESL (0.06 ± 0.07 radians, p -value < 0.001). In our experiments, the average time to perform detection using the ESL pipeline was 5.9 seconds per CXR, compared 10.3 seconds for the MSL using the same computational hardware. To demonstrate the localization accuracy, mean shape ($\bar{x}$) was aligned using both ESL and MSL, Fig. 7c demonstrated the DSC of the alignment of mean shape with the ground truth labels. ESL clearly outperforms the state-of-the-art MSL both in time as well as accuracy. Implementations were done in Matlab (The MathWorks, Inc., Natick, MA) and ran using CPU only.

B. Shape Parameter Estimation: MaShDL vs. ASM

Fig. 8(a) shows the boxplots of DSC for the lung field segmentation using just the mean shape (baseline), SSM-based ASM [22], and the MaShDL (using single and two SSMs). A single model was created using the training data from all three datasets. Two separate shape models were created using the clustering criteria described in Section III-F. Mean shape initialization was performed using ESL (Section III-E). The best results were achieved with the two SSMs; however, in both cases, MaShDL significantly outperforms the conventional ASM (p -value < 0.001 for single SSM, p -value < 0.001 for two SSM). A mean DSC of 0.85 ± 0.04 was obtained using just the mean shape alignment through ESL, 0.92 ± 0.03 using ASM, and 0.96 ± 0.03 using MaShDL. The results in the boxplot are reported using the modes carrying 95% cumulative energy ($K = 15$).

Fig. 8(b) shows the performance of ASM and MaShDL as a function of cumulative modes of variation (two SSMs). The DL mechanism adopted by MaShDL to extract the local appearance features deforms the shape contour to the true object boundary using less number of modes than ASM. Also from eq. (2) and (14), each atomic unit within ASM and MaShDL have the same order of computational complexity; therefore, MaShDL was demonstrated to be faster than ASM in our experiments for a given performance accuracy. Specifically, MaShDL framework was also found to be at least four times faster on average than SSM for a given accuracy.

C. Quantitative Comparison with State-of-the-Art Methods

We compared the segmentation performance obtained through our approach to the results reported by the state-of-the-art methods using three widely used metrics (overlap score, average contour distance (ACD), and DSC) in Table III. The table reports the performance on both

lungs. The results reported here for our method are obtained on the original 2048×2048 images and not the down-sampled version. None of the other methods includes the retro-cardiac region within the segmentation. In addition, we also compared the segmentation performance with the U-Net based architecture proposed by Wang et al. [10] specifically for lung field segmentation task: the current state-of-the-art convolutional neural network for biomedical image segmentation. The U-Net architecture and its derivatives have been extensively used for segmentation in radiological and histological images, providing some of the most accurate and satisfactory performances [40], [41]. The U-net architecture is a fully convolutional network, which includes shortcut connections between a contracting encoder and a successive expanding decoder. The quantitative overall segmentation performance as well as on individual datasets (JSRT, BTP, CNHS) using Wang et al.'s approach is also reported in Table III for segmentation labels with and without retro-cardiac space. Although, a range of hyper-parameters were tested; however, ones proposed by Wang et al. were found to be optimal for the task. Exact same architecture and hyper parameters as reported in [10] were used. Furthermore, since Wang et al. used level-set based post-processing step, the results in Table III are reported using both without and with post-processing step [?] for a fair comparison since our proposed approach does not use any post-processing. In our experiments, we found that while post-processing step helps in smoothing the contours, removing some holes, and false positives; however, by using diverse training data and optimal hyper-parameters suggested by Wang et al. [10], most of these issues have been taken care of by the original architecture and the post-processing step offers negligible improvement. This as can be observed through quantitative results presented in Table III.

D. Qualitative Comparison with State-of-the-Art Methods

Fig. 9 presents the qualitative results of performing the lung segmentation using the proposed pipeline (ESL+MaShDL). The figure provides a visual insight on how inclusion of retro-cardiac region results in the segmentation label that is independent to the shape and structural changes in the close-by anatomical structures such as heart. Fig. 10 shows the qualitative results of performing the lung segmentation using Wang et al.'s method [10]. As discussed earlier, the shape specificity is not preserved for the lung field labels obtained using [10]. This is further evident through the results presented in Table III. Moreover, unlike the proposed method, Wang et al.'s approach uses an overlapping-based objective function (e.g., cross-entropy) which provides satisfactory results in cases with reduced shape variability. However, in the particular case of thoracic radiographs, the lung field labels without retro-cardiac space present higher shape variability than those observed when including this region. This could be a possible explanation of a slightly better overlapping-based performance (i.e., Overlap and DSC) by [10] when including the retro-cardiac space than without including it.

V. Discussion and Conclusion

This work introduced a generic representation learning framework for the deformable object segmentation via space (translation, orientation, anisotropic scaling) and shape parameter estimation. The boundary detectors in the conventional statistical shape models (SSM) do not work consistently well on the data with complex patterns or with poor contrast and edge

information. Furthermore, since the SSMs are known to be sensitive to initial shape estimation, an efficient learning-based mechanism to estimate the space parameters (translation, scale, and orientation) was also presented in this work to initialize the mean shape. On the contrary, state-of-the-art learning-based approaches such as U-Net fail to accommodate the shape of the object-of-interest as complementary extra information and therefore perform sub-optimally as has been demonstrated through our results.

Our solution to space parameter learning, ensemble space learning (ESL), was significantly more accurate than current state-of-the-art marginal space learning (MSL) [28] and marginal space deep-learning (MSDL) [29] approaches as demonstrated through rigorous experiments. Although ESL has the potential to be generically applicable for the localization of objects of interest in 2D/3D images; however, the algorithm, in its current form, assumes symmetry (such as lung field) of the object of interest as well as the neighborhood context information for efficiency purposes. Therefore, while ESL is envisioned to demonstrate best performance for the localization of objects in medical images where organ symmetry as well as the contextual information can be somewhat guaranteed; for the general computer vision tasks, the ESL may need to be modified for optimal results. Furthermore, due to the existence of clinical acquisition protocols, large rotational variation is not expected in various CXR images; therefore, rotation estimation is still performed sequentially after translation and scaling rather than independently. The method can be easily modified for tasks where large rotational variation in the training data is expected such as using multiple passes to obtain the minimum area bounding box.

A. Training Considerations within the Proposed Framework

Although convolutional neural network (CNN) based DL architectures can be used in our framework. The principal difference between the fully connected networks and the CNN is that the convolutional filters in CNNs limit the scope of connections to a local neighborhood. The idea was motivated by the fact that the pixel information is dependent upon a small neighborhood around the pixel and is independent of pixels outside that neighborhood; therefore, CNNs were introduced as a method for reducing the number of trainable parameters within the network. Since in our approach we are extracting the image patches prior to the deep neural network, the spatial context is already established and the number of trainable parameters have been limited already therefore there is no additional benefit of using CNN in this scenario. Furthermore, although we do find any significant difference in performance with and without pre-training. However, the idea of pre-training is predicted make significant difference both in terms of accuracy and convergence when new training data from another site with substantially different protocol will be included.

Furthermore, certain empirically determined heuristics were applied during the training process (eqns. (11), (12), and (15)); however, there is a logical explanation to that can help the reader extend the idea to other datasets as well as other applications. The clean separation between positive (eq. (11)) and negative (eq. (12)) hypotheses is created for effective training of the classifier. Furthermore, the lower (0.25) limit in eq. (12) is based on the fact that there has to be a reasonable difference between positive and negative hypothesis of a mode in order to effectively train a classifier. The lower limit may vary depending upon

the application, acquisition protocol, and the resolution of images. The upper limit (3.0) is picked up from the literature on statistical shape modeling (Cootes et al. [37]).

B. Limitations of the Proposed Framework

For a given performance accuracy, our formulation for marginal shape deep learning (MaShDL) estimated 2D/3D deformable shape parameters significantly faster than the conventional SSM-based methods. As has been stressed throughout in the manuscript, MaShDL extends the ASM into deep learning realm; which results in better overall accuracy as demonstrated by rigorous set of experiments. However, since the mathematical framework behind MaShDL is still similar to the ASM, some of the limitations of the traditional ASM that are part of the mathematical framework exists in the MaShDL framework as well, that includes: (1) the tedious task of labeling training images that becomes unacceptable especially with the large training set. As pointed out previously in the manuscript that although the current scheme of six manually defined landmarks was found sufficient for accurate lung field segmentation (Dice score between the manual ground truth label and the label obtained using the interpolated landmarks = 0.994 ± 0.001), different number of manually annotated landmarks can be tested based on the application and the object of interest. (2) Similar to the traditional ASM, parameters such as the number of modes of variation still need to be specified. Furthermore, as the difference between the positive and negative hypothesis becomes more subtle at higher modes; in order to learn modes beyond a certain limit, approaches such as the use of deeper networks and mode-dependent thresholds for hypotheses testing need to be investigated. (3) The use of global statistical shape models by approaches like ASM is one of the most successful method to impose shape and anatomical constraints in medical image segmentation. However, while providing robust and anatomically accurate constraints, it also limits the flexibility of the method to deal with small localized shape details, such as the region around the diaphragm in chest radiographs. The reason for subtle boundary imperfections (under-segmentation or over-segmentation) by any statistical shape modeling method including MaShDL are the localized intensity changes and contrast variation amongst the pixels that are supposed to be on the one side of the boundary due to noise in the acquisition process. To overcome this limitation, we intend to extend our previous work on partitioned shape modeling [42] to MaShDL in the future. Although we exemplified an application of our framework through the segmentation of the lung field from CXR using diversified populations (i.e., age, pathology, and source); however, even with this diversity, as long as the standard acquisition protocols of routine clinical environment were followed, the algorithm was designed to robustly handle variation in shape. Our framework is applicable to general deformable object segmentation in both 2D and 3D image data, as a faster and potentially more accurate alternative to statistical appearance and shape model.

C. Insight Analysis of Various Segmentation Approaches

As described in Table I, the automated segmentation techniques for lung field from CXRs can be divided to four major categories, namely, rule-based methods, feature classification-based methods, deformable shape model-based methods, hybrid methods. In Table III, we summarized the performance of methods belonging to each category. As expected, the hybrid methods have demonstrated the best overall performance (min. mean overlap

score=0.907), followed by not so distant feature classification-based methods (min. mean overlap score=0.897), and finally the solo deformable shape model-based methods (min. mean overlap score=0.851). The inferior performance of deformable shape model-based methods can be attributed to the lack of sufficient training data and the inflexibility of these models to accommodate subtle changes across the population especially changes not accommodated within the training data. However, it is important to emphasize here that the performance of any hybrid method depends strongly on the individual algorithmic components used to produce the hybrid. The selection of these individual components and their task-specific hyper-parameter optimization is pivotal in the performance of overall pipeline. We believe that the future research in this direction will involve optimal selection of individual algorithmic components as well as hyper-parameters specific to the task and the data.

Another important observation made during the study was the diminishing need for post-processing steps with the state-of-the-art segmentation techniques. Our proposed method does not use any post-processing while the effect of post-processing on the performance of the method proposed in [10] was found to be negligible. We believe that with the advent of more sophisticated methods and the availability of larger training data, the need for post-processing will reduce even further.

Acknowledgments

This project was funded by NIH grants UL1TR000075/ KL2TR000076 and R41HL145669.

References

- [1]. Candemir S, Antani S, Jaeger S, Browning R, and Thoma GR, "Lung boundary detection in pediatric chest X-rays," in SPIE Medical Imaging. International Society for Optics and Photonics, 2015, pp. 94 180Q–94 180Q.
- [2]. Smeets M, Brunekreef B, Dijkstra L, and Houthuijs D, "Lung growth of pre-adolescent children," *European Respiratory Journal*, vol. 3, no. 1, pp. 91–96, 1990. [PubMed: 2311736]
- [3]. Brown MS, Wilson LS, Doust BD, Gill RW, and Sun C, "Knowledge-based method for segmentation and analysis of lung boundaries in chest X-ray images," *Computerized Medical Imaging and Graphics*, vol. 22, no. 6, pp. 463–477, 1998. [PubMed: 10098894]
- [4]. Duryea J and Boone JM, "A fully automated algorithm for the segmentation of lung fields on digital chest radiographic images," *Medical Physics*, vol. 22, no. 2, pp. 183–191, 1995. [PubMed: 7565349]
- [5]. Armato SG, Giger ML, and MacMahon H, "Automated lung segmentation in digitized posteroanterior chest radiographs," *Academic Radiology*, vol. 5, no. 4, pp. 245–255, 1998. [PubMed: 9561257]
- [6]. Li L, Zheng Y, Kallergi M, and Clark RA, "Improved method for automatic identification of lung regions on chest radiographs," *Academic Radiology*, vol. 8, no. 7, pp. 629–638, 2001. [PubMed: 11450964]
- [7]. Van Ginneken B and ter Haar Romeny BM, "Automatic segmentation of lung fields in chest radiographs," *Medical Physics*, vol. 27, no. 10, pp. 2445–2455, 2000. [PubMed: 11099215]
- [8]. McNitt-Gray MF, Huang H, and Sayre JW, "Feature selection in the pattern classification problem of digital chest radiograph segmentation," *Medical Imaging, IEEE Transactions on*, vol. 14, no. 3, pp. 537–547, 1995.

- [9]. Dai W, Doyle J, Liang X, Zhang H, Dong N, Li Y, and Xing EP, "Scan: Structure correcting adversarial network for chest x-rays organ segmentation," arXiv preprint arXiv:1703.08770, 2017.
- [10]. Wang C, "Segmentation of multiple structures in chest radiographs using multi-task fully convolutional networks," in Scandinavian Conference on Image Analysis. Springer, 2017, pp. 282–289.
- [11]. Dawoud A, "Fusing shape information in lung segmentation in chest radiographs," in Image Analysis and Recognition. Springer, 2010, pp. 70–78.
- [12]. Annangi P, Thiruvankadam S, Raja A, Xu H, Sun X, and Mao L, "A region based active contour method for X-ray lung segmentation using prior shape and low level features," in Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on. IEEE, 2010, pp. 892–895.
- [13]. Sohn K, "Segmentation of lung fields using chan-vey active contour model in chest radiographs," in SPIE Medical Imaging. International Society for Optics and Photonics, 2011, pp. 796332–796332.
- [14]. Chan TF and Vese LA, "Active contours without edges," Image processing, IEEE transactions on, vol. 10, no. 2, pp. 266–277, 2001.
- [15]. Shi Y, Qi F, Xue Z, Chen L, Ito K, Matsuo H, and Shen D, "Segmenting lung fields in serial chest radiographs using both population-based and patient-specific shape statistics," Medical Imaging, IEEE Transactions on, vol. 27, no. 4, pp. 481–494, 2008.
- [16]. Xu T, Mandal M, Long R, Cheng I, and Basu A, "An edge-region force guided active shape approach for automatic lung field detection in chest radiographs," Computerized Medical Imaging and Graphics, vol. 36, no. 6, pp. 452–463, 2012. [PubMed: 22608158]
- [17]. Shao Y, Gao Y, Guo Y, Shi Y, Yang X, and Shen D, "Hierarchical lung field segmentation with joint shape and appearance sparse learning," Medical Imaging, IEEE Transactions on, vol. 33, no. 9, pp. 1761–1780, 2014.
- [18]. Candemir S, Jaeger S, Palaniappan K, Musco JP, Singh RK, Xue Z, Karargyris A, Antani S, Thoma G, and McDonald CJ, "Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration," Medical Imaging, IEEE Transactions on, vol. 33, no. 2, pp. 577–590, 2014.
- [19]. Ibragimov B, Likar B, Pernuš F, and Vrtovec T, "Accurate landmark-based segmentation by incorporating landmark misdetections," in Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on. IEEE, 2016, pp. 1072–1075.
- [20]. Cosío FA, "Automatic initialization of an active shape model of the prostate," Medical Image Analysis, vol. 12, no. 4, pp. 469–483, 2008. [PubMed: 18346931]
- [21]. Zhang S, Zhan Y, Dewan M, Huang J, Metaxas DN, and Zhou XS, "Towards robust and effective shape modeling: Sparse shape composition," Medical Image Analysis, vol. 16, no. 1, pp. 265–277, 2012. [PubMed: 21963296]
- [22]. Cootes TF, Taylor CJ, Cooper DH, and Graham J, "Active shape models-their training and application," Computer Vision and Image Understanding, vol. 61, no. 1, pp. 38–59, 1995.
- [23]. Zheng Y, Barbu A, Georgescu B, Scheuering M, and Comaniciu D, "Four-chamber heart modeling and automatic segmentation for 3D cardiac CT volumes using marginal space learning and steerable features," Medical Imaging, IEEE Transactions on, vol. 27, no. 11, pp. 1668–1681, 2008.
- [24]. Ghesu FC, Georgescu B, Zheng Y, Hornegger J, and Comaniciu D, "Marginal space deep learning: Efficient architecture for detection in volumetric image data," in Medical Image Computing and Computer- Assisted Intervention–MICCAI 2015. Springer, 2015, pp. 710–718.
- [25]. Anthimopoulos M, Christodoulidis S, Ebner L, Christe A, and Mougiakakou S, "Lung pattern classification for interstitial lung diseases using a deep convolutional neural network," IEEE Transactions on Medical Imaging, vol. PP, no. 99, pp. 1–1, 2016.
- [26]. Shin H-C, Orton MR, Collins DJ, Doran SJ, and Leach MO, "Stacked autoencoders for unsupervised feature learning and multiple organ detection in a pilot study using 4D patient data," Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 35, no. 8, pp. 1930–1943, 2013.

- [27]. Ronneberger O, Fischer P, and Brox T, “U-net: Convolutional networks for biomedical image segmentation,” in International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, 2015, pp. 234–241.
- [28]. Zheng Y, Barbu A, Georgescu B, Scheuering M, and Comaniciu D, “Four-chamber heart modeling and automatic segmentation for 3D cardiac CT volumes using marginal space learning and steerable features,” *Medical Imaging, IEEE Transactions on*, vol. 27, no. 11, pp. 1668–1681, 2008.
- [29]. Ghesu FC, Krubasik E, Georgescu B, Singh V, Zheng Y, Hornegger J, and Comaniciu D, “Marginal space deep learning: efficient architecture for volumetric image parsing,” *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1217–1228, 2016. [PubMed: 27046846]
- [30]. Davies R, Taylor C et al., *Statistical models of shape: Optimisation and evaluation*. Springer Science & Business Media, 2008.
- [31]. Zheng Y, Lu X, Georgescu B, Littmann A, Mueller E, and Comaniciu D, “Robust object detection using marginal space learning and ranking-based multi-detector aggregation: Application to left ventricle detection in 2d mri images,” in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on IEEE*, 2009, pp. 1343–1350.
- [32]. Schwing A, Zheng Y, Harder M, and Comaniciu D, “Method and system for anatomic landmark detection using constrained marginal space learning and geometric inference,” 1 29 2013, uS Patent 8,363,918.
- [33]. Lu X, Georgescu B, Zheng Y, Otsuki J, and Comaniciu D, “Autompr: Automatic detection of standard planes in 3d echocardiography,” in *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on IEEE*, 2008, pp. 1279–1282.
- [34]. Erhan D, Bengio Y, Courville A, Manzagol P-A, Vincent P, and Bengio S, “Why does unsupervised pre-training help deep learning?” *Journal of Machine Learning Research*, vol. 11, no. Feb, pp. 625–660, 2010.
- [35]. Vincent P, Larochelle H, Bengio Y, and Manzagol P-A, “Extracting and composing robust features with denoising autoencoders,” in *Proceedings of the 25th international conference on Machine learning. ACM*, 2008, pp. 1096–1103.
- [36]. Zheng Y and Comaniciu D, *Marginal Space Learning for Medical Image Analysis*. Springer, 2014.
- [37]. Cootes TE and Lanitis A, “Active shape models: Evaluation of a multi-resolution method for improving image search,” in *Proc. British Machine Vision Conference*, 1994, pp. 327–338.
- [38]. Krizhevsky A, Sutskever I, and Hinton GE, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [39]. Van Ginneken B, Stegmann MB, and Loog M, “Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database,” *Medical image analysis*, vol. 10, no. 1, pp. 19–40, 2006. [PubMed: 15919232]
- [40]. Kamnitsas K, Ledig C, Newcombe VF, Simpson JP, Kane AD, Menon DK, Rueckert D, and Glocker B, “Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation,” *Medical image analysis*, vol. 36, pp. 61–78, 2017. [PubMed: 27865153]
- [41]. Chen H, Qi X, Yu L, and Heng P-A, “Dcan: Deep contour-aware networks for accurate gland segmentation,” in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016, pp. 2487–2496.
- [42]. Mansoor A, Cerrolaza JJ, Idrees R, Biggs E, Alsharid MA, Avery RA, and Linguraru MG, “Deep learning guided partitioned shape model for anterior visual pathway segmentation,” *IEEE transactions on medical imaging*, vol. 35, no. 8, pp. 1856–1865, 2016. [PubMed: 26930677]

Figure 1:

Illustration of age-related anatomical differences captured within CXRs. CXR obtained from: (a) 2-month old subject, (b) 4-year old subject, (c) 44-year old subject. (d) Structural differences in the lung field between the adults and pediatrics based on the aspect ratio, (e) Structural differences in the lung field between the adults and pediatrics based on the two largest modes of principal component analysis. Intuitively, mode values define variation in shape within the data. Largest modes define the highest variation.

Figure 2:

A chest radiograph with lung field delineation overlay: (a) without retro-cardiac region, (b) with retro-cardiac region. (c) Correlation between the lung volume estimated from computed tomography (CT) scan with segmented lung field area from CXR (without retro-cardiac region, $R=0.80$, including retro-cardiac region, $R=0.86$). Red and blue boundaries indicate the lung field with and without the retro-cardiac region respectively.

Figure 3:

The overview flow diagram of the proposed method for generic space and shape learning.

Figure 4:

Illustration of the differences in approach between the marginal space learning (MSL) and the ensemble space learning (ESL) to estimate Ω_{Space} . (a) MSL, (b) ESL. To estimate Ω_{Space} , MSL uses minimum area bounding box while ESL employs linearly independent bounding lines. The dashed green patch of size $(2r + 1, Ny)$ shows a positive hypothesis satisfying eq. (11) for line l_1 .

Figure 5:

Architectures of the proposed deep learning framework for (a) line estimation, (b) orientation estimation, the pendulum indicates the orientation of the bounding box and (c) modes of shape variation. The squares indicate the patches (size = $q \times q$) extracted around landmarks. Yellow and orange patches indicate the first patch for right and left lungs respectively. The clock-wise and counter clock-wise white arrows indicate the direction of concatenation of patches to obtain the feature set.

Figure 6:

Four highest modes of variation in training data are shown from left to right. For each mode, we show the superimposed synthesized shapes (step size= 0.05): The landmarks forming the positive shape hypotheses of the mode is shown in green while the landmarks forming the negative shape hypotheses (eq. (15)) are shown in red.

Figure 7: Boxplots of space parameters estimation ($\hat{\Omega}_{\text{Space}}$) error using MSL [28] and ESL (proposed). (a) Translation (**T**) and scale (**S**), (b) Orientation θ , and (c) Performance comparison of ESL and MSL in aligning the mean shape ($\bar{x}$).

Figure 8:

(a) Boxplots of shape parameters estimation ($\hat{\Omega}_{\text{Shape}}$) error using ASM [22] and MaShDL (proposed) using $K = 15$ largest modes of shape variation. (b) Average segmentation accuracy, measured using DSC, obtained as a function of the number of shape parameters using ASM [22] and MaShDL (proposed). Mode 0 denotes the aligned mean shape ($\bar{\mathbf{x}}$), the mean shape for both ASM and MaShDL are aligned using ESL.

Figure 9:

Qualitative lung field segmentation results using the proposed framework (ESL+MaShDL). The cases shown were randomly chosen from the dataset. The segmentation labels obtained are overlaid in over the input CXR. The blue region denotes the overlap area between the ground expert segmented manual ground truth and the segmentation produced using the proposed framework, the green region denotes ground truth area, and the red region denotes the segmentation obtained using the proposed framework. The heart is also visible underneath the segmentation label.

Figure 10:

Qualitative lung field segmentation results obtained using the state-of-the-art U-Net based architecture proposed by Wang et al. [10] for the segmentation of structures in CXR. The cases shown were randomly chosen from the dataset. The segmentation labels obtained, without using the post-processing step [?], are overlaid in over the input CXR. The blue region denotes the overlap area between the ground expert segmented manual ground truth and the segmentation produced using [10], the green region denotes ground truth area, and the red region denotes the segmentation obtained using [10]. As can be seen from the overlaid labels that shape specificity is not preserved by the U-Net architectures.

Table I:**BRIEF DESCRIPTION OF THE STATE-OF-THE-ART LUNG FIELD SEGMENTATION TECHNIQUES.**

Rule-Based Methods	
Brown <i>et al.</i> [3]	Matches the anatomical model of lung to extracted edges from image.
Dureya <i>et al.</i> [4]	Extracts diaphragm for lung field extraction.
Armato <i>et al.</i> [5]	Uses global and local intensities.
Li <i>et al.</i> [6]	Combines edge-based feature classification with iterative contour smoothing.
Feature Classification-Based Methods	
van Ginneken <i>et al.</i> [7]	Uses k-NN classifier with Gaussian derivative filters for multiscale pixel classification.
McNitt-Gray <i>et al.</i> [8]	Employs linear discriminator and neural networks with features selected.
Dai <i>et al.</i> [9]	Uses adversarial network that jointly train a segmentation network and a critic network.
Wang <i>et al.</i> [10]	Uses fully convolutional network (FCN) to simultaneously segment multiple structures including the lung field within chest radiographs.
Candemir <i>et al.</i> [1]	Use multiple shape models for pediatrics and adults based on age and shape information.
Deformable Shape Model-Based Methods	
Dawoud <i>et al.</i> [11]	Fuses shape prior with intensity threshold.
Annangi <i>et al.</i> [12]	Integrates lung edge and castrophenic angle into level set information.
Sohn <i>et al.</i> [13]	Uses active contour model [14] for lung field segmentation.
Shi <i>et al.</i> [15]	Combines cohort-specific statistics for constraining the deformable contour.
Xu <i>et al.</i> [16]	Combines edge and region forces for shape model deformation.
Hybrid Methods	
Shao <i>et al.</i> [17]	Uses local shape and appearance sparse learning in hierarchical deformation framework.
Candemir <i>et al.</i> [18]	Uses multiple atlases with non-rigid registration.
Ibragimov <i>et al.</i> [19]	Employs Haar-like features with random forest classifier to model the appearance of the landmarks and shape-based Gaussian distribution to model the spatial relationship amongst those landmarks.

Note: None of the methods include the retro-cardiac region as part of the lung field label.

Table II:

PARAMETER DETAILS OF THE DEEP-LEARNING ARCHITECTURES.

Architecture	Number of Parameters	
	Pre-Training	Fine-Tuning
Line Estimation (a)	321, 201 (801×401)	65, 561, 401 ($801 \times 401 \times 201$)
Orientation	321, 201 (801×401)	65, 561, 401
Estimation (b)		($801 \times 401 \times 201$)
Shape Mode Estimation (c)	1,282,401 (1601×801)	519, 385, 229 ($1601 \times 801 \times 401 \times 101$)

QUANTITATIVE COMPARISON OF THE PROPOSED APPROACH WITH THE CURRENT STATE-OF-THE-ART METHODS IN LUNG FIELD SEGMENTATION. HIGHEST PERFORMING METHODS ARE HIGHLIGHTED IN BOLD.

Table III:

Method	Overlap Score/ Jaccard Index	ACD (mm)	Mean±Standard Deviation (Min/Max)	DSC	Elapsed Time (sec)
Deformable Shape Model-Based Methods					
Arrangi et al. [12]	-	-	-	0.880±0.070	-
Dawoud et al. [11] [†]	0.940±0.005	2.460±2.060	-	-	-
Sohn et al. [13] [†]	0.851±0.046	-	-	0.952±0.016	-
Shi et al. [15] [‡]	0.920±0.031	2.492±1.092	-	-	-
Shao et al. [17] [†]	0.946±0.019	1.669±0.762	-	0.972±0.010	-
Van Ginneken et al. [7]	-	-	-	-	-
AAM Whiskers	0.913±0.032	2.700±1.100	-	-	-
ASM tuned	0.927±0.032	2.300±1.030	-	-	-
Feature Classification-Based Methods					
Dai et al. [9] [†]	0.929±0.500	-	-	0.963±0.3	-
Candemir et al. [1] ^{††} Van Ginneken et al. [7]	0.897	0.945	-	1.110	-
Pixel Class. (PC)	0.938±0.027	3.250±2.650	-	-	-
Hybrid Methods Van Ginneken et al. [7]					
Hybrid AAM+PC	0.933±0.025	2.060±0.840	-	-	-
Hybrid ASM+PC	0.934±0.037	2.080±1.400	-	-	-
PC+Post-processing	0.945±0.022	1.610±0.800	-	-	-
Hybrid Voting	0.949±0.020	1.620±0.660	-	-	-
Candemir et al. [18] [‡]	0.954±0.015	1.321±0.316	-	-	-
Ibragimov et al. [19] [†]	0.953±0.2	1.430±0.850	-	-	-
Inter-Observer Agreement					
With Retro-Cardiac Region	0.932±0.033	1.726±1.252	-	0.951±0.029	-
Without Retro-Cardiac Region [39]	0.946 ±0.018	1.640±0.690	-	-	-

U-Net Implementation by Wang et al. [10] (With Retro-Cardiac Region)

Complete Pipeline including Level Set Post-Processing ^{†††}; Hybrid Method

Method	Overlap Score/ Jaccard Index	ACD (mm)	Mean±Standard Deviation (Min/Max)	DSC	Elapsed Time (sec)
overall	0.907±0.062(0.604/0.973)	1.627±0.839(0.540/4.869)	0.947±0.074(0.703/0.986)		0.0072±0.004
JSRT	0.909±0.036(0.719/0.955)	3.018±0.997(1.297/4.859)	0.952±0.021(0.836/0.977)		0.0074±0.004
BTP	0.943±0.042(0.674/0.973)	2.893±0.443(2.01/3.496)	0.968±0.027(0.873/0.986)		0.0071±0.006
CNHS	0.894±0.064(0.604/0.960)	1.295±0.703(0.471/4.869)	0.929±0.083(0.703/0.980)		0.0072±0.004
Without Level Set Post-Processing: Feature Classification-Based Method					
overall	0.91±0.059(0.603/0.971)	1.622±0.826(0.540/4.872)	0.941±0.065(0.611/0.985)		0.0072±0.004
JSRT	0.91±0.039(0.678/0.960)	2.974±0.819(1.361/4.959)	0.952±0.023(0.808/0.980)		0.0074±0.004
BTP	0.935±0.031(0.837/0.973)	2.689±0.563(2.126/3.251)	0.966±0.017(0.912/0.986)		0.0071±0.006
CNHS	0.897±0.068 (0.607/0.967)	1.283±0.603(0.459/4.981)	0.927±0.078(0.714/0.983)		0.0072±0.004
U-Net Implementation by Wang et al. [10] (Without Retro-Cardiac Region)					
Complete Pipeline including Level Set Post-Processing ^{†††}; Hybrid Method					
overall	0.906±0.06(0.638/0.98)	1.265±0.676(0.551/9.821)	0.938±0.035(0.789/0.987)		0.007±0.004
JSRT	0.959±0.019(0.807/0.98)	1.132±0.859(0.618/9.821)	0.978±0.01(0.893/0.987)		0.006±0.005
BTP	0.842±0.075(0.638/0.936)	2.323±0.59(1.543/3.123)	0.922±0.047(0.789/0.967)		0.007±0.004
CNHS	0.887±0.045(0.670/0.954)	1.221±0.5(0.551/7.591)	0.933±0.026(0.802/0.977)		0.006±0.004
Without Level Set Post-Processing: Feature Classification-Based Method					
overall	0.904±0.063(0.631/0.980)	1.237±0.775(0.454/3.836)	0.948±0.037(0.774/0.990)		0.007±0.004
JSRT	0.949±0.029(0.759/0.976)	1.775±1.252(0.723/3.921)	0.973±0.016(0.863/0.988)		0.006±0.005
BTP	0.842±0.075(0.643/0.937)	2.218±0.426(1.891/3.156)	0.912±0.046(0.783/0.967)		0.007±0.004
CNHS	0.877±0.055(0.600/0.965)	1.343±0.747(0.514/7.905)	0.933±0.034(0.747/0.976)		0.006±0.004
Proposed Method (With Retro-Cardiac Region)					
overall	0.949±0.061(0.798/0.978)	1.510±0.484(0.412/2.327)	0.968±0.025(0.794/0.989)		0.0055±0.003
JSRT	0.954±0.044(0.833/0.980)	1.671±0.715(0.414/2.711)	0.958±0.009(0.855/0.989)		0.0056±0.003
BTP	0.945±0.063(0.813/0.974)	1.466±0.381(0.717/1.917)	0.968±0.021(0.833/0.987)		0.0055±0.004
CNHS	0.948±0.132(0.818/0.960)	1.466±0.713(0.718/2.299)	0.954±0.033(0.789/0.982)		0.0055±0.003
Proposed Method (Without Retro-Cardiac Region)					
overall	0.961±0.091(0.784/0.979)	1.491±0.757(0.682/2.786)	0.971±0.041(0.772/0.988)		0.0056±0.004
JSRT	0.961±0.067(0.852/0.974)	1.640±0.877(0.622/2.548)	0.969±0.011(0.871/0.988)		0.0056±0.004
BTP	0.962±0.092(0.787/0.979)	1.585±0.402(0.814/2.784)	0.975±0.037(0.778/0.988)		0.0055±0.003
CNHS	0.953±0.101(0.848/0.967)	1.223±0.698(0.479/2.341)	0.969±0.052(0.821/0.978)		0.0055±0.004

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

$$DSC = 2 \frac{|GT \cap SEG|}{|GT| + |SEG|} \quad \text{Overlap (Jaccard index)} = \frac{|GT \cap SEG|}{|GT \cup SEG|} \quad ACD = \frac{1}{2} \left(\frac{|GT|}{\sum_{p \in SEG} d(p, GT)} + \frac{|SEG|}{\sum_{p \in GT} d(p, SEG)} \right)$$

GT= binary labels of manual ground truth. SEG= binary labels produced by the proposed method. The operator $||$ denotes cardinality.

^{*} Method tested on the JSRT dataset.

^{*} Method tested on the JSRT dataset among others.

^{**} Method tested on the pediatric dataset.

^{***} Level set implementation described in [?] was used.