

Automated estimation of total lung volume using chest radiographs and deep learning

Ecem Sogancioglu | Keelin Murphy | Ernst Th.Scholten | Luuk H. Boulogne |
Mathias Prokop | Bram van Ginneken

Department of Medical Imaging, Radboud University Medical Center, Institute for Health Sciences, Nijmegen, The Netherlands

Correspondence

Ecem Sogancioglu, Radboud University Medical Center, Institute for Health Sciences, Department of Medical Imaging, Nijmegen, The Netherlands.

Email: Ecem.Sogancioglu@radboudumc.nl

Funding information

Ministry of Economic Affairs (Perspectief programme), Grant/Award Number: P15-26; DLMedIA: Deep Learning for Medical Image Analysis

Abstract

Background: Total lung volume is an important quantitative biomarker and is used for the assessment of restrictive lung diseases.

Purpose: In this study, we investigate the performance of several deep-learning approaches for automated measurement of total lung volume from chest radiographs.

Methods: About 7621 posteroanterior and lateral view chest radiographs (CXR) were collected from patients with chest CT available. Similarly, 928 CXR studies were chosen from patients with pulmonary function test (PFT) results. The reference total lung volume was calculated from lung segmentation on CT or PFT data, respectively. This dataset was used to train deep-learning architectures to predict total lung volume from chest radiographs. The experiments were constructed in a stepwise fashion with increasing complexity to demonstrate the effect of training with CT-derived labels only and the sources of error. The optimal models were tested on 291 CXR studies with reference lung volume obtained from PFT. Mean absolute error (MAE), mean absolute percentage error (MAPE), and Pearson correlation coefficient (Pearson's r) were computed.

Results: The optimal deep-learning regression model showed an MAE of 408 ml and an MAPE of 8.1% using both frontal and lateral chest radiographs as input. The predictions were highly correlated with the reference standard (Pearson's $r = 0.92$). CT-derived labels were useful for pretraining but the optimal performance was obtained by fine-tuning the network with PFT-derived labels.

Conclusion: We demonstrate, for the first time, that state-of-the-art deep-learning solutions can accurately measure total lung volume from plain chest radiographs. The proposed model is made publicly available and can be used to obtain total lung volume from routinely acquired chest radiographs at no additional cost. This deep-learning system can be a useful tool to identify trends over time in patients referred regularly for chest X-ray.

KEYWORDS

artificial intelligence, chest radiograph, chest x-ray, deep learning, total lung volume

1 | INTRODUCTION

Chest radiography (CXR) remains the most commonly performed imaging technique and one of the most often repeated exams because of its low cost, rapid acquisi-

tion, and low radiation exposure.¹ It was estimated that 129 million chest radiographs were performed in 2006, in the United States alone.² Chest radiographs play an important role in screening, monitoring, diagnosis, and management of thoracic diseases.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Medical Physics* published by Wiley Periodicals LLC on behalf of American Association of Physicists in Medicine.

Wide availability of CXR has motivated researchers to build artificial intelligence (AI) systems that can automatically detect a variety of abnormalities^{3–5} and extract quantitative clinical measurements from them.^{6,7} AI systems have potential use for routine quantification of numerous biomarkers related to lung diseases, cardiac health, or osteoporosis. Applying such systems, whenever a chest radiograph is acquired, would be a step toward routine quantitative radiology reporting.

This work focuses on an important quantitative biomarker, total lung volume (TLV), and investigates whether it can be measured automatically from plain chest radiographs using state-of-the-art deep-learning approaches. Total lung volume is used for assessing severity, progression, and response to treatment in restrictive lung diseases.^{8,9} Specific temporal changes in TLV can be identified in patients with obstructive and restrictive lung diseases, such as emphysema, pulmonary fibrosis, or asthma. Further, TLV has been shown to correlate with mortality and health status.¹⁰

Currently, the gold standard for measurement of TLV is the pulmonary function test (PFT), using special techniques such as body plethysmography, helium, or nitrogen dilution techniques.⁹ Several studies^{11–13} demonstrated that TLV measured from CT strongly correlates to TLV obtained from PFTs. Alternatively, several studies investigated TLV estimation from CXR using predictive equations. In fact, this has been a research interest for a century, with the first relevant paper appearing in 1918¹⁴ demonstrating the correlation of external measurements to the PFT (using gas dilution technique). All such previous literature, investigating predictive equations, was either based on the use of planimetric techniques^{15–18} or made assumption of a given geometry,^{19–21} or required several manual linear measurements to estimate TLV from CXR. However, all these studies required manual measurements to estimate TLV and used small sample sizes, making it unclear whether the techniques could be generalized to other populations.

In this study, we investigate, to the best of our knowledge, for the first time, whether CXR can be used to automatically predict TLV in a fully automated fashion using large datasets and deep learning. We examine the role of TLV labels derived from thoracic CT imaging in training deep-learning systems. In order to account for variations in inspiration and dataset complexity, experiments with simulated and real chest radiographs in three different datasets were designed in a stepwise fashion. For each experiment, we optimized various state-of-the-art deep-learning regression approaches to predict TLV using only posterioranterior (PA) view, lateral view, or both views. The purpose of our study was to determine the accuracy of fully automatic measurement of TLV from CXR using deep-learning-based models.

2 | MATERIALS AND METHODS

2.1 | Data and preprocessing

The data used in this study was obtained from two sources: the COPDGene study²² and Radboud University Medical Center (RUMC). To facilitate our stepwise experimentation, demonstrating sources of error, we experimented with simulated CXR images (digitally reconstructed radiographs), which are obtained from a forward-projection using a parallel beam geometry (i.e., average intensity projections [AIPs]) on thoracic CT, as well as with true CXR images. Reference TLV labels were obtained by two means; through segmentation of the lungs in CT or PFTs.

The datasets are detailed in the following sections and in Figure 1.

2.2 | COPDGene-sim

Inspiration chest CT studies (1000) from unique patients were randomly selected from the COPDGene study,²² which is publicly available for research purposes. The images in this study are acquired from patients with chronic obstructive pulmonary disorder (COPD), varying from mild to very severe. From the 1000 randomly selected CT studies, 800 (600 training, 200 development) were used for training and validation, and 200 were retained as a held-out test set as illustrated at the top of Figure 1.

Lung segmentations were obtained by an automated algorithm and manually corrected by trained analysts with radiologist supervision.²³ Reference TLV was calculated for each CT scan by multiplying CT voxel size by the number of voxels segmented.

All CT scans were first resampled to 1 mm, 1 mm, 1 mm spacing and simulated CXRs were generated from CT by creating AIP²⁴ from coronal and sagittal planes, resulting in frontal and lateral view simulated CXR. This dataset, which we refer to as COPDGene-sim, was used to demonstrate model performance in an ideal scenario where there is no inspiration difference between the label source (CT) and the (simulated) CXR image, CT segmentations are manually corrected, and the variety of pathologies is limited.

2.3 | RUMC datasets

This data was obtained from routine clinical care in RUMC, Nijmegen, the Netherlands. This study was approved by the research ethics committee of the Radboud University Nijmegen Medical Centre. Dataset was collected and anonymized according to local guidelines and informed consent was obtained from all participants.

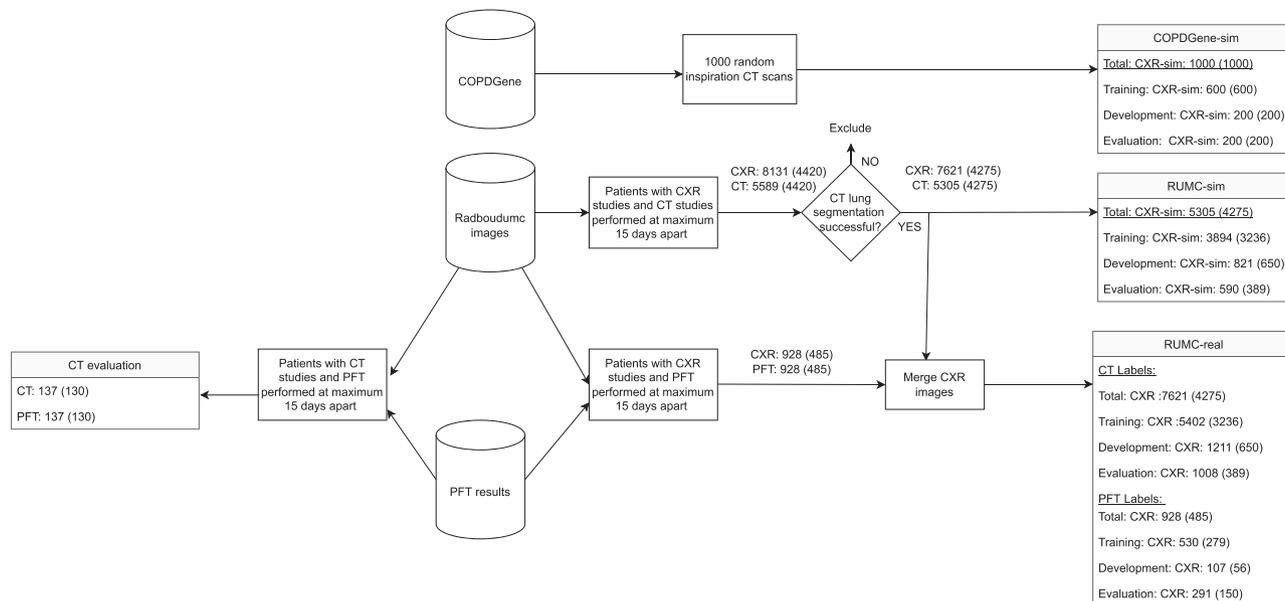


FIGURE 1 Flowchart of the data-selection procedure. Flowchart that shows the criteria to select the data to be used in the experiments. Numbers of images are shown with numbers of patients in brackets. CXR = chest radiographs, CXR-sim = simulated chest radiographs from CT, PFT = pulmonary function test

All research was performed in accordance with relevant guidelines and regulations.

We retrospectively collected CXR studies and chest CT acquired between 2003 and 2019 resulting in 321k CXR studies and 120k CT studies. Patients with both CT and CXR (with PA and lateral view), performed a maximum of 15 days apart, were selected (4420 patients). The reference standard TLV measurements were obtained by a CT lung segmentation algorithm²³ and segmentation failure cases were visually identified and excluded (284 CT). This resulted in 7621 CXR studies and 5305 CT studies from 4275 patients (Figure 1). Multiple CXR studies from a single patient could be matched to a single CT reference standard.

A group of patients being assessed for lobectomy was used to provide subjects with both PFT and CXR data acquired within 15 days of each other. This resulted in 928 CXR studies from 485 patients. Reference TLV was determined using the helium dilution technique.²⁵

From this dataset, we created two sets for experimentation. The first is referred as RUMC-sim and used simulated CXR generated from CT as described earlier. The second is RUMC-real, consisting of real CXR with CT-derived and PFT-derived labels for TLV. To investigate the relationship between CT-derived and PFT-derived labels, we created a dataset, CT-evaluation, where both CT and PFT were acquired within 15 days of each other. We made sure that there was no patient overlap between training and held-out evaluation sets for all the datasets. These datasets are detailed later and illustrated in Figures 1 and 2.

RUMC-sim: In this dataset, PA and lateral CXRs were simulated from 5305 CT studies (4275 patients). Of these, 389 patients (590 CT studies) were randomly selected and used as a held-out evaluation set, whereas the remaining 3886 were used for training. This dataset was used, with CT-derived TLV labels, to illustrate the model performance in a set of images with a large variety of abnormalities (compared to COPDGene-sim). The use of simulated CXR images removes any possibility of error related to inspiration effort, or patient position between the label source (CT) and the (simulated) CXR.

RUMC-real: This dataset consists of patients with real CXR studies (PA and lateral) and with TLV reference from two sources, CT and PFT. For CT-derived label data, the same subject partition was used as in RUMC-sim. This resulted in 7621 CXR studies with CT-derived labels. PFT-derived labels were used for 928 CXR studies as seen in Figure 1. As held-out evaluation sets, 590 patients with 1008 CXR (CT-derived labels) and 150 patients with 291 CXR (PFT-derived labels) were randomly selected. All CXR images were resized to have 1 mm by 1 mm spacing.

CT evaluation dataset: We identified patients with PFT results that were also in the RUMC-sim dataset, and selected those with PFT acquired a maximum of 15 days apart from CT (137 CT studies from 130 patients). CT lung volume was calculated by a CT lung segmentation algorithm,²³ and the results were visually inspected, identifying no obvious failed segmentations. This set was used to demonstrate the relationship between CT-derived and PFT-derived labels.

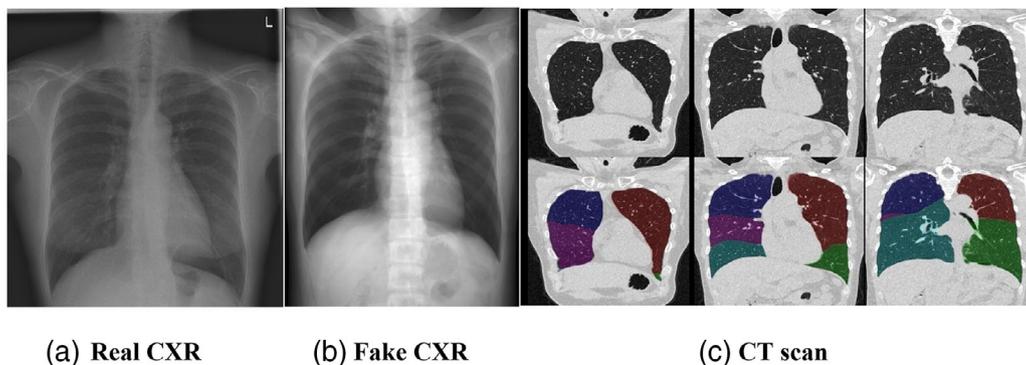


FIGURE 2 Sample data from a single subject. Real CXR (a), Simulated CXR (b) and coronal CT slices (c) from a patient in the RUMC-real dataset. Lobe segmentation results in CT are illustrated in the bottom row of (c). CT-derived TLV is calculated as the sum of the lobe volumes. Each color in the figure represents a lung lobe which was segmented by an automated algorithm. CT-derived TLV for this subject was 3.8 L, while PFT-derived TLV was 4.3 L

3 | METHODS

We experiment with 5 different deep-learning architectures, 4 of which are widely used popular classification architectures (DenseNet121,²⁶ ResNet34, ResNet50,²⁷ VGG-Net19²⁸), and one, referred as 6-layer CNN, was designed to represent a shallow architecture. The 6-layer CNN consisted of 6 CNN layers, each followed by RELU, batch normalization, and a max-pooling layer. The first CNN layer had 32 feature maps, and the number of feature maps was doubled in each layer. The final CNN layer was followed by 3 fully connected layers using linear activation function, which mapped the number of features to 512, 128, and 1, respectively.

These network architectures were trained from scratch with 3 possible inputs (PA CXR, lateral CXR or both, and methods of combining their outputs (see Figure 3)). Each network outputs a regression value representing TLV in liters. Before feeding the input to the network, all real and simulated CXR images were padded with 0 to reach 512×512 pixels. Images underwent standard normalization in the range of -1 to 1 , and the corresponding TLV measurements were normalized between 0 to 1. The dual CNN architecture, which receives both frontal and lateral radiographs as input, consists of two branches with a backbone architecture that is either VGG-Net, ResNet34 or 6-layer CNN, and concatenates the features from these branches before the first fully connected layer. Due to memory limitations, Densenet121 and ResNet50 architectures were not investigated for the dual CNN model.

For each model trained, a hyperparameter optimization was carried out to ensure the best possible result for that architecture/input combination on the validation set. A variety of aspects for training a convolutional neural network were considered as hyperparameters: They were learning rate, optimizer, oversampling technique, and data augmentation as seen in Figure 4. Data

augmentation techniques consisted of brightness and contrast, rotation, translation, and horizontal flip. If oversampling hyperparameter was chosen, the chest X-rays with low and high TLV labels, which were underrepresented in the dataset, were oversampled during training (more details regarding the hyperparameter optimization in Supporting Information). Random hyperparameter optimization was employed given a predefined range for hyperparameters for each model (frontal, lateral, and dual CNN) separately.

Each model was trained by optimizing the mean squared error loss between the predicted TLV and the reference standard TLV. The model was trained for a maximum of 300 epochs, terminating if there was no improvement in the validation set performance for 50 successive epochs. We selected the epoch that yielded the least mean squared error in the validation set.

For each of our three datasets, the optimal combination of architecture and hyperparameters was identified for each of the three possible input types on the validation set. These models were then applied to the held-out evaluation set. In addition, the average of the two outputs from the networks using single (frontal or lateral) inputs is calculated and presented as Ensemble CNN output (Figure 3).

Our TLV prediction experiments were constructed in a stepwise fashion, to identify potential sources of error as the task becomes increasingly difficult. This is given in Table 1. CT-derived volume labels are used in all experiments except the final one where the network is additionally fine-tuned with PFT-derived labels. We begin with the COPDGene-sim dataset, where errors related to patient position and inspiration effort as well as errors related to CT segmentation accuracy and diversity of underlying pathologies are eliminated. In RUMC-sim, we introduce the potential for errors from minor CT segmentation inaccuracies, and from the diverse pathology within the dataset, which is likely to increase the

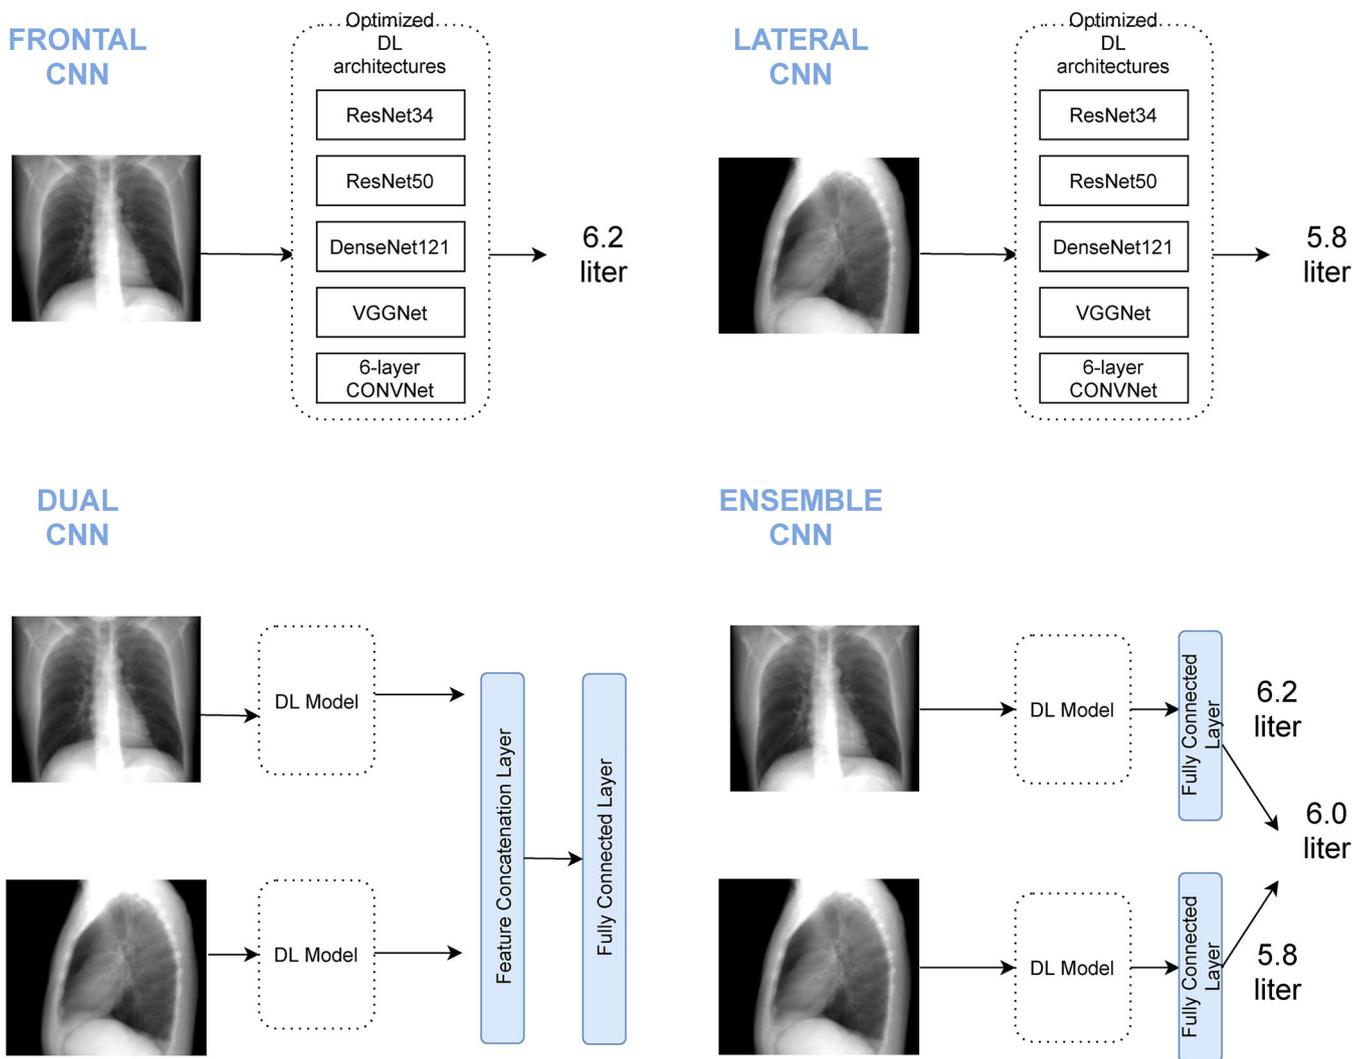


FIGURE 3 Illustration of architecture pipelines. Four different experimental designs were considered: frontal CNN, lateral CNN, dual CNN (combining frontal and lateral models by layer concatenation) and ensemble CNN (combining optimal frontal and lateral models by averaging their outputs)

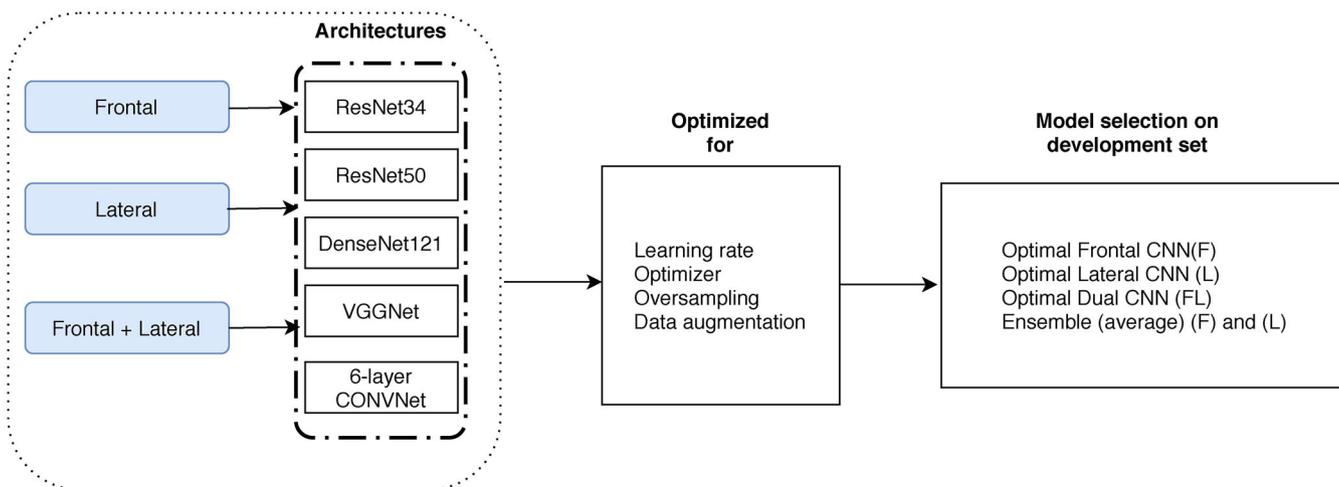


FIGURE 4 Illustration of our model selection process on validation set. Different network architectures were systematically optimized for three different inputs, namely frontal, lateral, and dual (frontal+lateral), separately. Each of them was optimized systematically for hyperparameters, and the model with the least mean absolute percentage error on the validation set was selected

TABLE 1 Datasets characteristics in stepwise experiments

		COPDGene-sim	RUMC-sim	RUMC-real (CT-labels)	RUMC-real (PFT-labels)
Reference TLV measurement	Label type	CT-derived	CT-derived	CT-derived	PFT-derived
Possible sources of label error	Patient position difference			Y	Y
	Inspiration effort difference			Y	Y
	CT segmentation				
	Inaccuracy		Y	Y	
	Diverse pathologies		Y	Y	Y

Note: RUMC-real (PFT-labels) was used to finetune the models which were pretrained on RUMC-real (CT-derived) dataset. Y indicates that the condition holds true.

variability in image appearance. Finally in RUMC-real, we first experiment with predicting CT-derived TLV from chest radiographs (RUMC-real [CT-labels]), and subsequently with PFT-derived TLV (RUMC-real [PFT labels]). The networks to predict CT-derived TLV were trained from scratch with real CXR images. In this last experiment, since there is only a small number of gold-standard PFT labels available (487 patients), the network trained with CT-labels is used as pretrained model, and fine-tuned using CXR images with associated PFT-labels.

As an additional experiment, we investigate the relationship between PFT-derived TLV and CT-derived TLV, in a scenario where they are acquired at most 15 days apart from each other, using the CT-evaluation dataset.

All the experiments in the paper were implemented in Python using PyTorch and other standard data-processing libraries such as pandas, sklearn, and imageio.

3.1 | Statistical analysis

Mean absolute error (MAE), mean absolute percentage error (MAPE), and Pearson correlation coefficient were computed to demonstrate the relationship between predicted and reference TLV values. About 95% limits of agreement were estimated by means of a nonparametric method for Bland–Altman plot since data distribution was not normal, assessed with Shapiro–Wilk test²⁹ and quantile–quantile plot.³⁰

4 | RESULTS

Model training for each model, namely frontal CNN, lateral CNN, and dual CNN, took between 8 to and 14 h on the RUMC-sim and RUMC-real datasets (depending on the network architecture), and 2–4 h on the COPDGene-

sim dataset using a variety of GPUs such as TitanX, GTX1080, GTX1080ti, GTXTitanx, and TitanV. The mean processing time per test image was 0.3 s.

Three trained models (frontal, lateral, dual) were selected for each dataset, based on optimization using the validation set, and applied to the held-out evaluation data. Additionally, the outputs of the optimized frontal and lateral models were averaged and presented as “Ensemble” model. The selected architectures and their performance on the held-out evaluation data are provided in Table 2.

In the COPDGene-sim dataset, where chest radiographs were simulated from CT and potential sources of label error were minimal, VGG-Net, 6-layer CNN, and Densenet121 architectures were selected. On the held-out evaluation set, the model with the lowest error according to all 3 metrics was the dual CNN with 6-layer CNN architecture. This model achieved a MAPE of 2.2% and MAE of only 112 ml. The scatter plot of model predictions against the reference standard from CT volumes and Bland–Altman-like plot for analyzing differences between the reference standard and predicted TLV measurements are shown in Figure 5a,b, respectively. As shown in Figure 5b, 95% of differences between predicted and reference standard TLV were from –351 to 261 ml.

On the RUMC-sim dataset, which contains more abnormal images compared to COPDGene-sim, Densenet121 and ResNet architectures were selected from the development set experiments. As in the COPDGene-sim experiments, the lateral CNN model performed better than the frontal CNN model and the best performance on the evaluation set was, once again, achieved by the dual CNN with MAPE of 2.9% and MAE of 112 ml as seen in Table 2 and plotted in Figure 5c. Limits of agreement of the differences between the predicted and reference standard TLV measurements were between -348 and 235 ml as shown in Figure 5d.

TABLE 2 Results of the selected models on the held-out evaluation sets

Evaluation Datasets (#images)	Model	Architecture	MAPE (%)	MAE (ml)	Pearson's r
COPDGene-sim (200)	Frontal CNN	DenseNet121	4.3	226	0.978
	Lateral CNN	VGG-Net	3.6	198	0.983
	Dual CNN	6-layer CNN	2.2	112	0.995
	Ensemble CNN	Densenet121&VGG-net	2.6	139	0.992
	Frontal CNN	Densenet121	5.5	220	0.978
RUMC-sim (590)	Lateral CNN	Densenet121	5.0	200	0.984
	Dual CNN	ResNet34	2.9	112	0.993
	Ensemble CNN	Densenet121&Densenet121	3.8	154	0.989
	Frontal CNN	VGG-Net	16.9	650	0.826
	Lateral CNN	Densenet121	16.8	639	0.831
RUMC-real (CT-labels) (1008)	Dual CNN	ResNet34	16.1	592	0.855
	Ensemble CNN	VGG-Net&Densenet121	15.7	597	0.851
	Frontal CNN	VGG-Net	10.3	509	0.870
	Lateral CNN	Densenet121	9.2	472	0.875
	RUMC-real (PFT-labels) (291)	Dual CNN	ResNet34	8.1	408
	Ensemble CNN	VGG-Net&Densenet121	8.5	420	0.907

Note: Mean absolute error is calculated against the reference standard for TLV measurements. MAE = mean absolute error (in milliliters), MAPE = mean absolute percentage error, Pearson's r = Pearson correlation coefficient. Bold font indicates best performance per dataset and metric.

Finally, in the RUMC-real dataset, where real chest radiographs were used, dual CNN and ensemble CNN performed very similarly, and the best result obtained (with the least MAPE) with CT-derived labels was achieved by the ensemble CNN, as shown in Table 2. This model achieved 15.7% MAPE and MAE of 597 ml. The model predictions and references for the evaluation set of 1008 CXRs are plotted in Figure 6a; and the differences between predicted TLV and reference standard are analyzed in Figure 6b. As shown in Figure 6b, the model tended to underestimate TLV where reference standard was higher than 6 L, and overestimate TLV where reference standard was lower than 4 L.

For the final experiment using PFT-derived labels, the best models trained on the RUMC-real (CT-labels) data for frontal, lateral, dual CNN were used as pretrained models and further fine-tuned on 637 CXR images with PFT-derived labels. The results achieved on 291 CXR images with PFT-derived labels are shown in Table 2 (RUMC-real [PFT-labels]). The best model on the held-out evaluation set was the dual CNN with ResNet34 architecture and achieved MAE of 408 ml and MAPE of 8.1%. The model predictions and PFT-derived reference standard were highly correlated with Pearson correlation coefficient of 0.92 as illustrated in Figure 6c; 95% of differences between predicted and reference standard TLV measurements were from -1 L to 938 ml (Figure 6d). Example cases of this model predictions are shown in Figure 7a–f.

Figure 8a,b shows the results of the comparison between CT-derived TLV and PFT-derived TLV on the CT evaluation set of 137 subjects. These two measurements were well correlated with Pearson's r of 0.78; however, considerable variations were observed between the two measurements for some patients. TLV was consistently underestimated by CT-based measurements where median differences (bias) between CT-derived and PFT-derived was -560 ml as shown in Figure 8b.

5 | DISCUSSION

This study demonstrated that state-of-the-art deep-learning solutions can measure TLV from PA and lateral CXRs, using primarily CT-derived labels and a small number of PFT-derived measures. To demonstrate the sources of error, the experiments were conducted in a stepwise fashion with increasing levels of complexity. Using simulated CXRs eliminated potential error related to the patient position or inspiration level between the CT and CXR image acquisition. Results on both simulated datasets show extremely low error (MAPE of 2.2% and 2.9%) and high correlation with the reference labels ($r = 0.99$ and $r = 0.99$). The slightly better performance on the COPDGene-sim dataset may be attributed to the fact that this dataset contains a limited range of pathologies and that the CT segmentations were manually

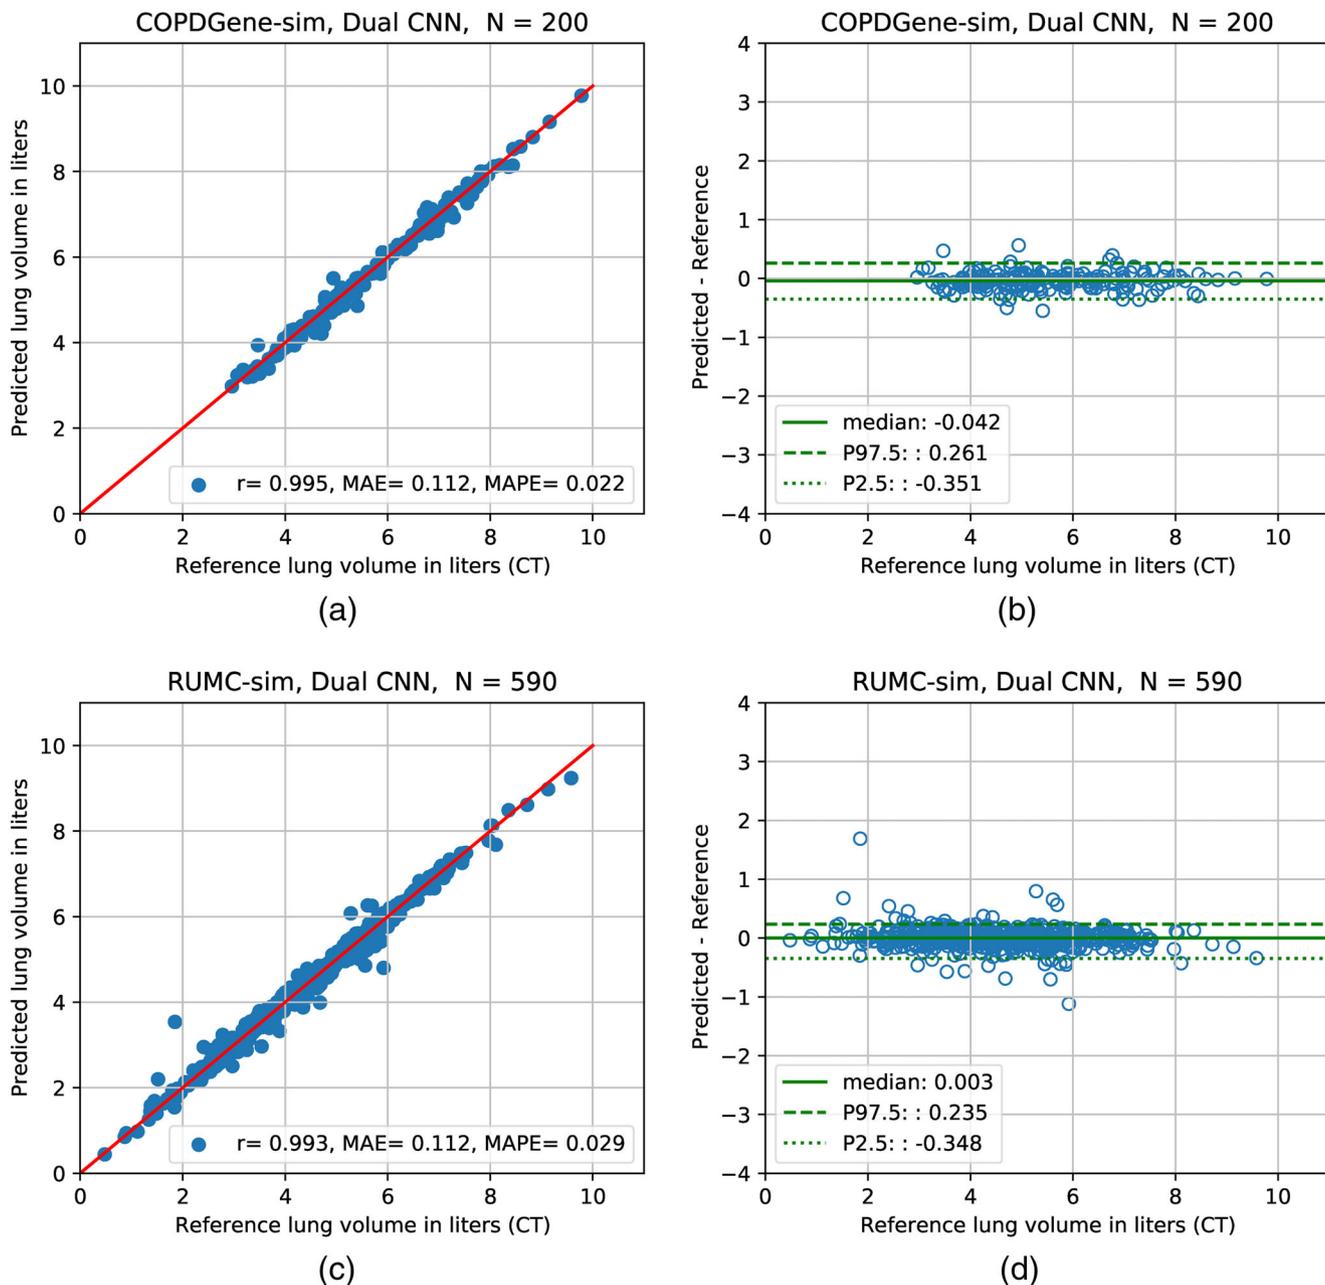


FIGURE 5 Results on simulated datasets in stepwise experiments. Left: The TLV predictions of the best model against the reference standard measurements on the held-out evaluation sets. (a) COPDGene, (c) RUMC-sim. Red line is line of identity (ideal agreement). Right: Bland–Altman-like plot to analyze the differences between predicted and reference standard TLV measurements. Nonparametric method was used to estimate 95% limits of agreement. r = Pearson correlation coefficient, MAE = mean absolute error, MAPE = mean absolute percentage error, N = number of data, P2.5 = 2.5th percentile P97.5 = 97.5th percentile

corrected, meaning that even very small inaccuracies were eliminated.

In the dataset of clinical CXR with CT-derived volumes (RUMC-real dataset), we see a substantial increase in the prediction error with MAPE of 15.7%, which we attribute largely to the difference in patient position and inspiration effort between the CT and the CXR image acquisition. It is likely that the degree of inspiration in the CXR and CT images is different, particularly given that there is known to be a high

intraindividual deviation in TLV between routine CT scans.³¹ The indication from this experiment is that CT-derived labels are useful, but not optimal, to learn the TLV from CXR. As an additional check, we investigated the relationship between CT-derived and PFT-derived volumes in 137 cases where both were available. This provides results in line with previous studies on CT-derived lung volumes^{31,32}: Although CT-derived lung volume and TLV are well correlated ($r = 0.78$), there are considerable differences in some patients.

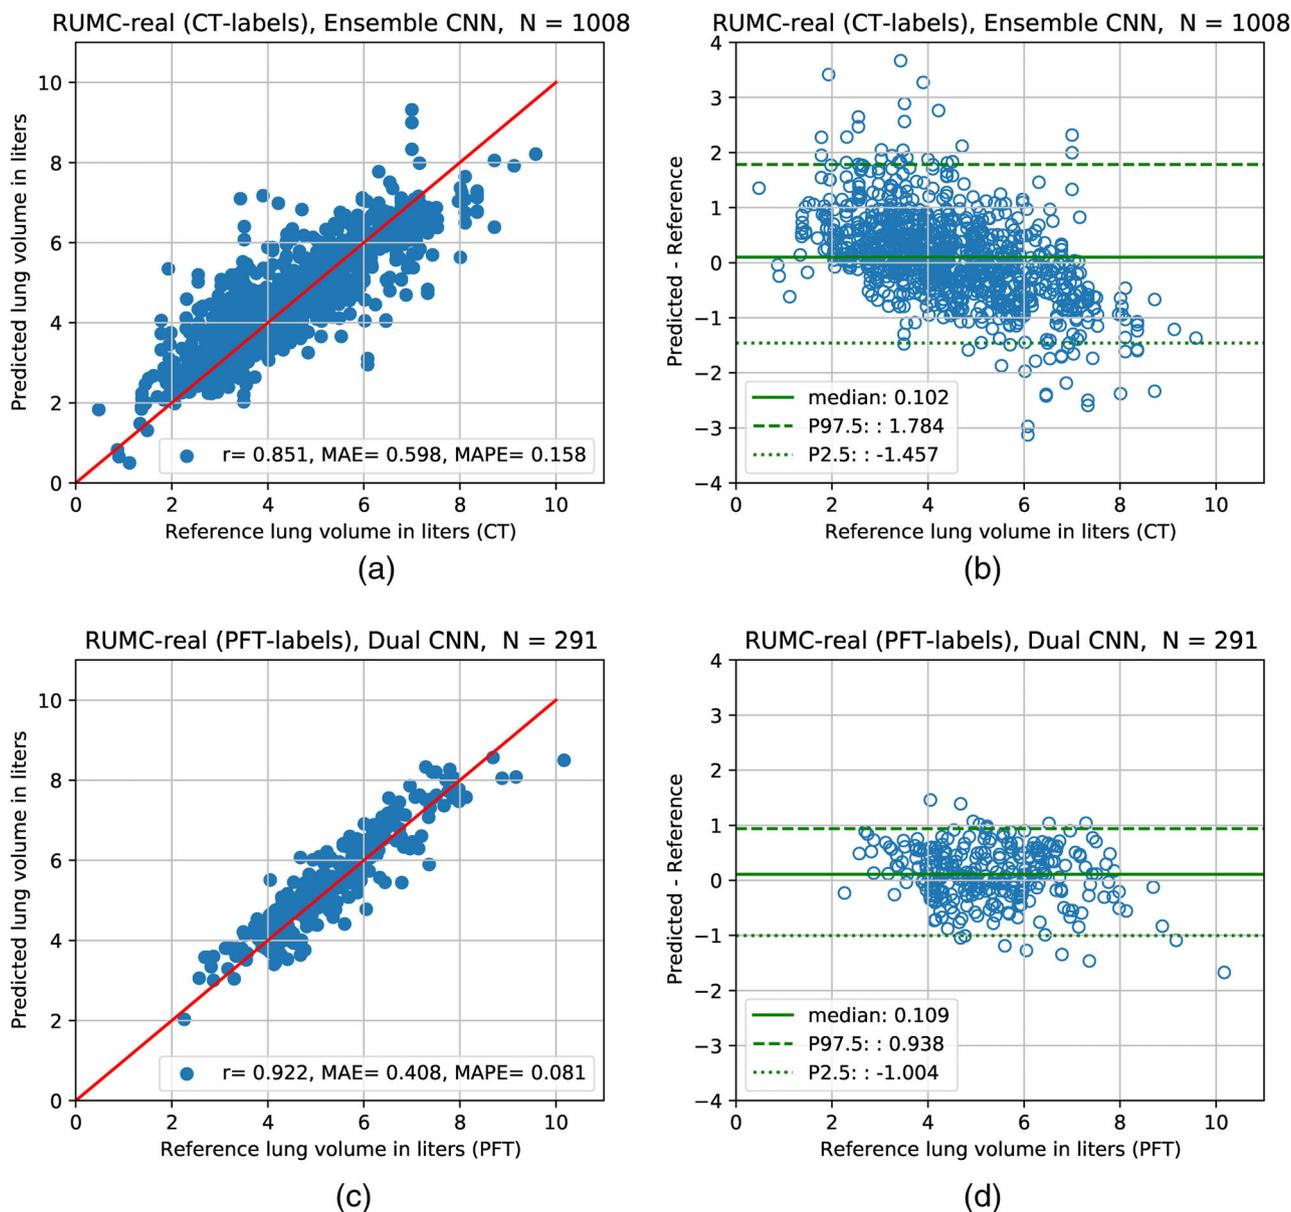


FIGURE 6 Results on real datasets in stepwise experiments. Left: The TLV predictions of the best model against the reference standard measurements on the held-out evaluation sets. (a) RUMC-real, (c) RUMC-real (PFT-labels). Red line is line of identity (ideal agreement). Right: Bland–Altman-like plot to analyze the differences between predicted and reference standard TLV measurements. Nonparametric method was used to estimate 95% limits of agreement. r = Pearson correlation coefficient, MAE = mean absolute error, MAPE = mean absolute percentage error, N = number of data, P2.5 = 2.5th percentile P97.5 = 97.5th percentile

To overcome the issues with the CT-derived labels on the RUMC-real dataset, we further fine-tuned the best networks from that experiment with PFT-derived labels. Evaluation on an independent dataset of 291 subjects that were not used for training showed that the error of the estimated TLV from CXR relative to the measured TLV from PFT is reduced considerably, achieving MAPE of 8.1% and Pearson's correlation coefficient of 0.92. This algorithm is publicly available at <https://grand-challenge.org/algorithms/cxr-total-lung-volume-measurement/>

In all experiments, the model was optimized to use the best performing architecture and input. In the experiments using simulated CXR images, it is notable that the networks using lateral images as input perform better than the networks using frontal images. This may indicate that the lateral projection image contains more information related to CT-derived TLV. However we note also that in all experiments the combination of frontal and lateral images produced the optimal results, either by use of a dual-CNN or through an ensemble.

of 2.2% and 2.9%). For studies that investigated predictive equations to estimate PFT-derived TLV from real CXR,^{16,19,20} the coefficient of correlation between predictions and reference standard (body plethysmography or helium dilution technique) generally ranged from 0.80 to 0.93 (compared to our method with 0.92). Sample sizes in these papers ranged from 21 to 100 patients. However, it should be noted that many of these studies used spirometric control to regulate the level of inspiration during CXR acquisition. In fact, one study³⁴ has shown that without spirometric control, the correlation of predicted TLV and PFT-derived reference standard was only 0.47, compared to 0.82 with spirometric control. In this work, however, we experiment with routinely taken chest radiographs (with no spirometric control), and produce TLV predictions that are highly correlated ($r = 0.92$) to PFT-derived results. Our work is the first to demonstrate automated measurement of TLV from chest radiographs and achieves a comparable or lower error range with a remarkably larger sample size compared to previous literature.

There are several limitations in this study. First, the algorithms were evaluated on an internal dataset from a single institution; validation of the models on an external dataset is an important next step to assess the algorithm robustness. Second, the datasets were constructed from routinely taken studies with the assumption that TLV would not change in 15 days, which might not hold true for extreme cases. This selection criterion also yielded an underrepresentation of healthy subjects but reflects a clinical population in which TLV measurements are of clinical interest. The PFT-derived reference standard measurements were obtained using the helium dilution technique which might underestimate TLV in patients with severe airway obstruction. Further, simulated CXRs were obtained using a parallel beam geometry; other techniques such as cone-beam geometry or advanced machine learning techniques might produce more realistic simulated CXRs. However, since the simulated radiographs were only used for pretraining, the overall results are expected not to be improved significantly. Furthermore, inspiration levels were not controlled in a similar fashion to PFT in these routine chest radiographs, which could have introduced a source of error in our predictions, but this represents regular clinical practice. One possible solution to address this issue would be to develop an automated algorithm to assess the inspiration level on CXR, for example by rib counting.³⁵ Moreover, our held-out evaluation set was constructed with patients assessed for lobectomy since their PFT results were readily available; future research should address the evaluation of the algorithm on a population with other clinically relevant pathologies, including fibrosis.

In conclusion, we demonstrated that TLV can be automatically estimated from CXR using a deep-learning approach, with an accuracy that is superior or comparable to the previous literature using semiautomated

methods. Further, we showed that the deep-learning system can be trained primarily with CT-derived labels from automatically segmented chest CT images and fine-tuned on gold-standard PFT-derived labels. This automated system could be routinely applied to clinical chest radiographs and serve as a tool for identifying temporal change in total lung volume in patients with restrictive and obstructive lung diseases.

ACKNOWLEDGMENTS

This work was supported by the Dutch Technology Foundation STW, which formed the NWO Domain Applied and Engineering Sciences and partly funded by the Ministry of Economic Affairs (Perspectief programme P15-26 “DLMedIA: Deep Learning for Medical Image Analysis”). We thank Weiyi Xie and Luuk Boulogne for their help in dataset collection, and Jesus Lago for useful discussion on statistical analysis.

CONFLICT OF INTEREST

Bram van Ginneken receives royalties and funding from Delft Imaging Systems and Mevis Medical Solutions and stock, royalties, and funding from Thirona. All other authors report no conflict of interest.

REFERENCES

1. Raoof S, Feigin D, Sung A, Raoof S, Irugulpati L, Rosenow EC. Interpretation of plain chest roentgenogram. *Chest*. 2012;141(2):545-558.
2. Mettler FA, Bhargavan M, Faulkner K, et al. Radiologic and nuclear medicine studies in the united states and worldwide: frequency, radiation dose, and comparison with other radiation sources—1950–2007. *Radiology*. 2009;253(11):520-531.
3. Annarumma M, Withey SJ, Bakewell RJ, Pesce E, Goh V, Montana G. Automated triaging of adult chest radiographs with deep artificial neural networks. *Radiology*. 2019;291(4):196-202.
4. Rajpurkar P, Irvin J, Ball RL, et al. Deep learning for chest radiograph diagnosis: a retrospective comparison of the CheXNeXt algorithm to practicing radiologists. *PLoS Med*. 2018;15(11):e1002686.
5. Murphy K, Smits H, Knoops AJG, et al. COVID-19 on chest radiographs: a multireader evaluation of an artificial intelligence system. *Radiology*. 2020;296(9):E166-E172.
6. Sogancioglu E, Murphy K, Calli E, Scholten ET, Schalekamp S, Ginneken BV. Cardiomegaly detection on chest radiographs: segmentation versus classification. *IEEE Access*. 2020;8:94631-94642.
7. Li Z, Hou Z, Chen C, et al. Automatic cardiothoracic ratio calculation with deep learning. *IEEE Access*. 2019;7:37749-37756.
8. Quanjer PH, Tammeling GJ, Cotes JE, Pedersen OF, Peslin R, Yernault J-C. Lung volumes and forced ventilatory flows. *Eur Respir J*. 1993;6(3):5-40.
9. Flesch JD, Dine CJ. Lung volumes. *Chest*. 2012;142(8):506-510.
10. Pedone C, Scarlata S, Chiurco D, Conte ME, Forastiere F, Antonelli-Incalzi R. Association of reduced total lung capacity with mortality and use of health services. *Chest*. 2012;141(4):1025-1030.
11. Tantucci C, Bottone D, Borghesi A, Guerini M, Quadri F, Pini L. Methods for measuring lung volumes: is there a better one? *Respiration*. 2016;91:273-280.
12. Coxson HO, Fauerbach PVN, Storness-Bliss C, et al. Computed tomography assessment of lung volume changes after bronchial valve treatment. *Eur Respir J*. 2008;32(12):1443-1450.

13. Iwano S, Okada T, Satake H, Naganawa S. 3D-CT volumetry of the lung using multidetector row CT. *Acad Radiol.* 2009;16(3):250-256.
14. Lundsgaard C, Slyke DDV. Studies of lung volume. I. *J Exp Med.* 1918;27(1):65-86.
15. Ries AL. Measurement of lung volumes. *Clinics in Chest Medicine.* 1989;10(6):177-186.
16. Harris TR, Pratt PC, Kilburn KH. Total lung capacity measured by roentgenograms. *Am J Med.* 1971;50(6):756-763.
17. Cobb S, Blodgett DJ, Olson KB, Stranahan A. Determination of total lung capacity in disease from routine chest roentgenograms. *Am J Med.* 1954;16(1):39-54.
18. Schlesinger AE, White DK, Mallory GB, Hildeboldt CF, Huddleston CB. Estimation of total lung capacity from chest radiography and chest CT in children: comparison with body plethysmography. *Am J Roentgenol.* 1995;165(7):151-154.
19. Barnhard HJ, Pierce JA, Joyce JW, Bates JH. Roentgenographic determination of total lung capacity. *Am J Med.* 1960;28(1):51-60.
20. Pierce RJ, Brown DJ, Holmes M, Cumming G, Denison DM. Estimation of lung volumes from chest radiographs using shape information. *Thorax.* 1979;34(12):726-734.
21. Loyd HM, String ST, DuBois AB. Radiographie and plethysmographie determination of total lung capacity. *Radiology.* 1966;86(1):7-30.
22. Regan EA, Hokanson JE, Murphy JR, et al. Genetic epidemiology of COPD (COPDGene) study design. *COPD: Journal of Chronic Obstructive Pulmonary Disease.* 2010;7(3):32-43.
23. Xie W, Jacobs C, Charbonnier J-P, Ginneken B. Relational modeling for robust and efficient pulmonary lobe segmentation in CT scans. *IEEE Trans Med Imaging.* 2020;39(8):2664-2675.
24. Campo MI, Pascau J, Estepar RSJ, Emphysema quantification on simulated X-rays through deep learning techniques. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, 2018.
25. Wanger J, Clausen JL, Coates A, et al. Standardisation of the measurement of lung volumes. *Eur Respir J.* 2005;26(9):511-522.
26. Huang G, Liu Z, Maaten LVD, Weinberger KQ, Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
27. He K, Zhang X, Ren S, Sun J, Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
28. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
29. Shapiro SS, Wilk MB. An analysis of variance test for normality (complete samples). *Biometrika.* 1965;52(12):591.
30. Wilk MB, Gnanadesikan R. Probability plotting methods for the analysis of data. *Biometrika.* 1968;55(3):1.
31. Haas M, Hamm B, Niehues SM. Automated lung volumetry from routine thoracic CT scans. *Acad Radiol.* 2014;21(5):633-638.
32. Daghighi A, Tropp H. Computed tomography lung volume estimation and its relation to lung capacities and spine deformation. *Journal of Spine Surgery.* 2019;5(3):132-141.
33. Park CH, Haam SJ, Lee S, Han KH, Kim TH. Prediction of anatomical lung volume using planimetric measurements on chest radiographs. *Acta Radiol.* 2016;57(7):1066-1071.
34. Wade OL, Gilson JC. The effect of posture on diaphragmatic movement and vital capacity in normal subjects. *Thorax.* 1951;6(6):103-126.
35. Mader AO, Berg J, Fabritz A, Lorenz C, Meyer C. Localization and labeling of posterior ribs in chest radiographs using a CRF-regularized FCN with local refinement. *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018*. Springer International Publishing; 2018:562-570.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Sogancioglu E, Murphy K, Th.Scholten E, Boulogne LH, Prokop M, van Ginneken B. Automated estimation of total lung volume using chest radiographs and deep learning. *Med Phys.* 2022;49:4466-4477. <https://doi.org/10.1002/mp.15655>