

RESEARCH ARTICLE

Automatic emphysema detection using weakly labeled HRCT lung images

Isabel Pino Peña¹*, Veronika Cheplygina^{2,3}*, Sofia Paschaloudi⁴‡, Morten Vuust⁴‡, Jesper Carl⁵‡, Ulla Møller Weinreich^{5,6}‡, Lasse Riis Østergaard¹‡, Marleen de Bruijne^{3,7}‡

1 Department of Health Science and Technology, Aalborg University, Aalborg, Denmark, **2** Department of Biomedical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands, **3** Biomedical Imaging Group Rotterdam, Erasmus Medical Center, Rotterdam, The Netherlands, **4** Department of Diagnostic Imaging, Vendsyssel Hospital, Fredrikshavn, Denmark, **5** Department of Clinical Medicine, Aalborg University Hospital, Aalborg, Denmark, **6** Department of Pulmonary Medicine, Aalborg University Hospital, Aalborg, Denmark, **7** Department of Computer Science, University of Copenhagen, Copenhagen, Denmark

* These authors contributed equally to this work.

‡ These authors also contributed equally to this work.

* ipino@hst.aau.dk (IPP); v.cheplygina@tue.nl (VC)



OPEN ACCESS

Citation: Pino Peña I, Cheplygina V, Paschaloudi S, Vuust M, Carl J, Weinreich UM, et al. (2018) Automatic emphysema detection using weakly labeled HRCT lung images. PLoS ONE 13(10): e0205397. <https://doi.org/10.1371/journal.pone.0205397>

Editor: Mathieu Hatt, INSERM, FRANCE

Received: February 3, 2018

Accepted: September 25, 2018

Published: October 15, 2018

Copyright: © 2018 Pino Peña et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: We have made derived data needed to replicate the study findings in this paper available on Figshare. These data include the histograms extracted from the CT images (<https://figshare.com/s/2d84a5aa0e809747e406>), and the code for the classifier results (<https://figshare.com/s/1ec45b700f685b91a1a6>). Representative, de-identified CT slices have been made available as Figures. In order to protect privacy and confidentiality of the study participants, complete CT data is not publicly available. Due to legal restrictions applied by Aalborg University Hospital

Abstract

Purpose

A method for automatically quantifying emphysema regions using High-Resolution Computed Tomography (HRCT) scans of patients with chronic obstructive pulmonary disease (COPD) that does not require manually annotated scans for training is presented.

Methods

HRCT scans of controls and of COPD patients with diverse disease severity are acquired at two different centers. Textural features from co-occurrence matrices and Gaussian filter banks are used to characterize the lung parenchyma in the scans. Two robust versions of multiple instance learning (MIL) classifiers that can handle weakly labeled data, miSVM and MILES, are investigated. Weak labels give information relative to the emphysema without indicating the location of the lesions. The classifiers are trained with the weak labels extracted from the forced expiratory volume in one minute (FEV₁) and diffusing capacity of the lungs for carbon monoxide (DLCO). At test time, the classifiers output a patient label indicating overall COPD diagnosis and local labels indicating the presence of emphysema. The classifier performance is compared with manual annotations made by two radiologists, a classical density based method, and pulmonary function tests (PFTs).

Results

The miSVM classifier performed better than MILES on both patient and emphysema classification. The classifier has a stronger correlation with PFT than the density based method, the percentage of emphysema in the intersection of annotations from both radiologists,

following the recommendations of the Danish data protection authorities (<http://www.datatilsynet.dk>), original scans can only be available to researchers who meet the criteria for accessing confidential data. Therefore, researchers who would like to have access to the CT data sets can contact MD. Jens Brøndum Frøkjær (jebf@dcm.aau.dk), chief Physician at Department of Diagnostic Radiology and co-author MD. Ulla M. Weinreich (ulw@rn.dk), chief physician at Department of Respiratory Diseases, Aalborg Hospital who are responsible for data access.

Funding: M.B and V.C. received financial support from the Netherlands Organization for Scientific Research (NWO), grant no. 639.022.010, www.nwo.nl/en.

Competing interests: The authors have declared that no competing interests exist.

and the percentage of emphysema annotated by one of the radiologists. The correlation between the classifier and the PFT is only outperformed by the second radiologist.

Conclusions

The presented method uses MIL classifiers to automatically identify emphysema regions in HRCT scans. Furthermore, this approach has been demonstrated to correlate better with DLCO than a classical density based method or a radiologist, which is known to be affected in emphysema. Therefore, it is relevant to facilitate assessment of emphysema and to reduce inter-observer variability.

1 Introduction

Chronic obstructive pulmonary disease (COPD) is the most important respiratory disease worldwide and one of the most important causes of death in high and middle-income countries [1, 2]. COPD is described as a progressive and irreversible airflow limitation. Emphysema is one of the most common disease manifestations that causes this limitation due to the destruction of alveolar walls and loss of elasticity [3]. Emphysema can be identified visually in computed tomography (CT) scans as low attenuation areas (LAA). However, to enable the detection of lesions smaller than 5 mm, thin slice reconstructions, such as high-resolution computed tomography (HRCT) scans, are preferred.

The automatic identification and quantification of emphysema provides objectivity and more reliability to the clinical routine in the assessment of COPD. Currently, emphysema is assessed visually, which is time consuming, subjective and suffers from inter- and intra-observer variability [4]. Over the years, the most used methods for automatically quantifying emphysema have been density based [5–7]. These methods use a threshold based on percentile density or LAA, generally lower than -950 Hounsfield units (HU). However, these methods are very dependent on, among others, the inspiration level, scanner reconstruction kernel, exposure dose and scanners. Therefore, there is no consensus on the best threshold for quantifying emphysema [8, 9]. Other quantification methods have been reported based on texture features, which collect information about the spatial relationship of the intensity values in the scan [10–12]. Machine learning methods based on texture analysis extract information to learn normal and abnormal lung tissues, which facilitates the recognition of disease patterns and can therefore lead to a more reliable diagnosis [13]. In general, machine learning methods use supervised classifiers that require annotated regions of interest (ROIs) or labeled patches based on manual annotations of emphysema performed by clinical experts [14–17]. Manual annotations are even more time consuming than visual assessment of emphysema and also suffer from inter-observer variation [18].

Learning from weak labels, which assign a label to the entire image, is proposed in the literature as the less time-consuming alternative to the manual annotation of patches, and it is being increasingly used in different medical image analysis applications [19–21]. Weak labels are easier to acquire than manual annotations because they can be obtained from basic quantification methods or complimentary data of the patient, such as pulmonary function tests (PFTs) or bio-markers. Classifiers which learn from weak labels are referred to as multiple instance learning (MIL) classifiers. All MIL classifiers can learn to label entire scans. For example, Sørensen et al. [22] and Cheplygina et al. [19] used spirometry results, which is the most common PFT to clinically assess COPD, to assign labels to scans from the Danish Lung Cancer

Screening Trial, and trained different types of MIL classifiers to detect COPD in previously unseen scans from the same trial. However, a subset of MIL classifiers can also learn to classify individual patches, thus identifying regions with signs of COPD, including emphysema. Neither [22] nor [19] evaluated MIL classifiers for this purpose. For example, more than half of the classifiers studied in [19] including the best performing classifier, could not provide individual patch labels.

In contrast with previous studies, this study aims to automatically identify emphysema regions in patients with COPD using HRCT scans without local annotations. Different texture-based methods and MIL classifiers are investigated. Furthermore, in this study, more robust versions of two MIL classifiers are proposed. The results from the classifiers are evaluated with manual annotations made by two radiologists.

2 Materials and methods

This study focuses on automatically distinguishing emphysema without using manual annotations to train the classifiers. For this purpose, different types of texture features are extracted to characterize emphysema, and two variations of MIL classifiers are investigated. Fig 1 presents an overview of the method used.

2.1 Features

Two different types of texture features are computed: features from co-occurrence matrices and Gaussian derivative features. The co-occurrence matrix algorithm is used in 3D, and it aims to capture the spatial dependence of gray-level intensities through multiple slices. The co-occurrence of voxel pairs is evaluated in 13 directions and at five different distances. After obtaining the co-occurrence matrices, the spatial dependencies of gray-level values are described by 12 Haralick textural features: energy, entropy, correlation, contrast, homogeneity,

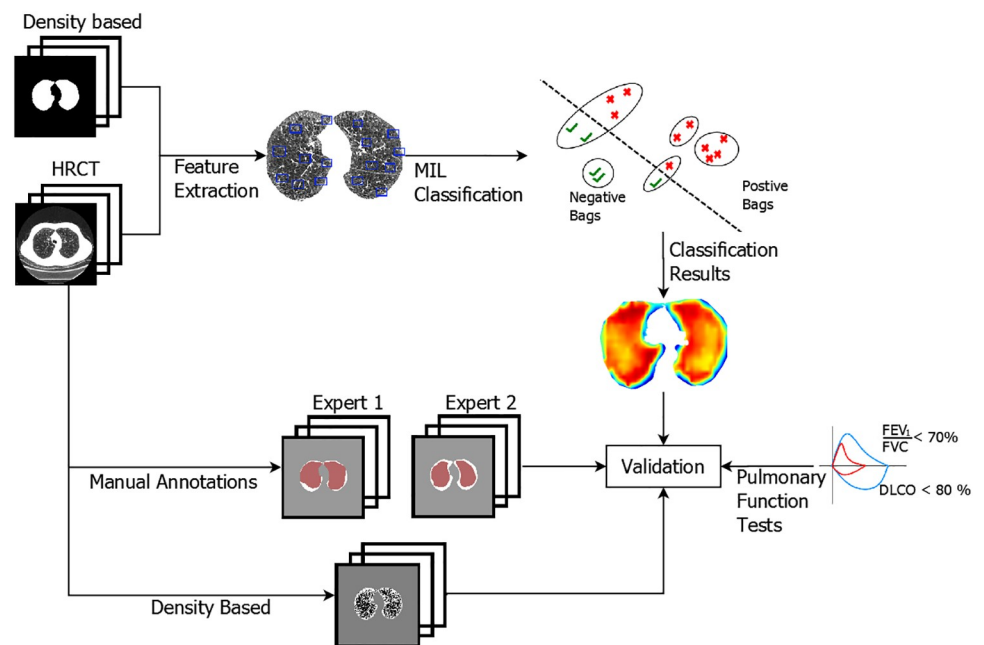


Fig 1. Summary of the methodology. Texture features are extracted from the lung parenchyma. Two different MIL classifiers are trained and are tested on previously unseen scans. The results are evaluated against manual annotations performed by two radiologists, a density based analysis, and pulmonary function tests.

<https://doi.org/10.1371/journal.pone.0205397.g001>

variance, sum mean, inverse difference moment, inertia, cluster shade, cluster tendency and max probability [23].

Gaussian derivative features aim to capture the presence of structures such as edges and blobs. Each image is first convolved (using normalized convolution) with a Gaussian function: $G(\mathbf{v}, \sigma) = \frac{1}{((2\pi)^{1/2}\sigma)^3} \exp\left(-\frac{\|\mathbf{v}\|_2^2}{2\sigma^2}\right)$, where σ represents the standard deviation of the Gaussian, or the scale at which the texture is examined, and $\mathbf{v} = [x, y, z]^T$ is a voxel. Similar to [22], eight filters are computed: smoothed image, gradient magnitude, Laplacian of Gaussian, three eigenvalues of the Hessian, Gaussian curvature, and eigen magnitude. The filters are computed at four different scales: 0.6mm, 1.2mm, 2.4mm, and 4.8mm. The filtered outputs are summarized using histograms with ten bins, where the bin sizes are determined by adaptive binning [24] on an independent dataset [22] prior to this study.

2.2 Classifiers

MIL is originally a binary classification problem, although multi-class extensions also exist [25]. MIL classifiers are trained on labeled bags $\{(B_i, y_i) | i = 1, \dots, N\}$, where i indicates the i -th out of total N subjects, and y_i is the label ($y_i = +1$ for COPD, or $y_i = -1$ for non-COPD) of the i -th subject. The bags are also referred to as positive or negative. Each bag $B_i = \{\mathbf{x}_{ij} | j = 1, \dots, n_i\} \subset \mathbb{R}^d$, is a set of n_i texture feature vectors or instances, where \mathbf{x}_{ij} describes the j -th patch of the i -th subject.

In this study, the bags represent the entire scan of an individual subject, whereas the instances are randomly selected 3D patches from inside the lungs. The bags are related to the weak labels extracted from the pulmonary function tests, and the instance labels classify the lung parenchyma into emphysematous or healthy lung tissue. Typically, MIL classifiers assume that there are instance labels y_{ij} that relate to the bag labels as follows: a bag is positive if and only if it contains at least one positive or *concept* instance: $y_i = \max_j y_{ij}$. Thus, a bag is classified as positive if at least one instance contains emphysematous tissue. In this study a less strict assumption is used, as described in Section 2.3.

There are two main strategies that MIL classifiers follow [19, 26]. The instance-level strategy is to use the bag labels to infer an instance classifier. To classify a previously unseen test bag, such classifiers classify its instances and then combine the instance labels into a bag label. The bag-level strategy is to represent the bags by some global characteristics and use supervised classifiers to classify the bags directly. Inferring the instance labels from the bag labels is not always possible in this case. In this study, the posterior probability that a classifier outputs for a bag is denoted as $f(B_i)$ and the posterior probability that a classifier outputs for an instance as $f(\mathbf{x}_{ij})$.

Two popular and effective MIL classifiers are miSVM [27] and MILES [28]. The instance-level miSVM classifier extends the traditional support vector machine (SVM) by searching for an instance classifier that separates the instances as well as possible but such that the $\max_j \{y_{ij}\} = y_i$ condition still holds. In other words, the most positive (according to the classifier) instance in each bag is positive if the bag is positive and negative if the bag is negative. Similar to a supervised SVM, the miSVM can operate on polynomial or radial basis kernels between instances. A regularization parameter C controls the trade-off between the margin, i.e. how well the instances are separated, and how many training bags are incorrectly classified with this margin. For a test bag, its instances are classified, and the most positive instance determines the label of the bag.

MILES is a bag-level approach that is able to infer instance labels. It assumes that positive and negative bags contain discriminative prototype instances. It represents each bag by a

feature vector \mathbf{s}_i that contains its similarities to all instances in the training set, where the similarity is defined as $s(B_i, \mathbf{x}) = \max_j k(\mathbf{x}_{ij}, \mathbf{x})$, in which k is a similarity function between instances, i.e., a polynomial or radial basis kernel. The maximum operator implies that the bag's similarity to an instance is high if it contains a single similar instance. The MILES classifier then selects discriminative similarities, which correspond to discriminative prototype instances. A regularization parameter C controls the trade-off between how many bags are incorrectly classified and how many discriminative prototypes are selected. For a test bag, the similarity to the discriminative prototypes determine whether the bag is positive (if it has instances sufficiently similar prototypes from positive bags).

2.3 Avoiding false positives

Because a single positive instance is sufficient to classify whether a bag is positive, miSVM and MILES may suffer from false positives. In this study, more robust formulations of miSVM and MILES are proposed, which we refer to as miSVM-Q and MILES-Q, which use the quantile rather than the maximum operator to define the label of the bag. In miSVM-Q, $\max_j \{y_{ij}\} = y_i$ is replaced by $\text{quantile}_j(\{y_{ij}\}, q) = y_i$, where q is the desired quantile. For example, if $q = 0.5$, half of the instances must be positive for a bag to be positive.

In MILES-Q, the similarity function to $s(B_i, \mathbf{x}) = \text{quantile}_j(\{k(\mathbf{x}_{ij}, \mathbf{x})\}, q)$ is adapted. This means that the bag must contain more similar instances to the prototype \mathbf{x} to be considered similar to it. For both miSVM-Q and MILES-Q, these adaptations mean that healthy subjects can still be considered healthy if they have a few emphysematous patches.

Furthermore, this study proposes an additional measure to evaluate a MIL classifier f , which is called Separability or S :

$$S = \frac{1}{\sum_{y_i=+1} n_i} \sum_{y_i=+1} f(\mathbf{x}_{ij}) - \frac{1}{\sum_{y_i=-1} n_i} \sum_{y_i=-1} f(\mathbf{x}_{ij}). \tag{1}$$

In other words, the Separability describes the difference between the average posterior probabilities of instances in true positive training bags, and the average posterior probabilities of instances in true negative training bags. The intuition behind this is that true positive bags should have a larger proportion of positive instances, and therefore the average instance posterior probability should be higher than in negative bags. This allows reasoning about the classifier's performance on instance-level, without having access to instance labels.

Consider two classifiers f_1 and f_2 which classify a positive and a negative bag, each with two instances. For the positive bag, f_1 outputs posteriors 0.51 and 0.49, and f_2 outputs posteriors 0.9 and 0.1. For the negative bag, f_1 outputs 0.49 and 0.49, while f_2 outputs 0.1 and 0.1. While both f_1 and f_2 correctly classify the bags, f_2 would be the preferred classifier on instance-level. The Separability of f_1 and f_2 are respectively 0.01 and 0.4, which reflects our preference for f_2 . Since Separability is a difference of two averages, each of which is between 0 and 1, Separability could theoretically range between -1 and 1. In our experiments we observed that most values fall between 0.05 and 0.75.

3 Experimental

3.1 Data

Two datasets are used in this study. Both datasets are named after the hospital where the HRCT scans were performed: Frederikshavn (Fre) and Aalborg (Aal). For both datasets, volumetric (helical) HRCT scans and pulmonary function tests (PFTs) are performed.

The PFTs are spirometry and diffusing capacity of the lungs for carbon monoxide (DLCO) and are acquired for each patient. PFTs and HRCT scans are performed with the patients in a steady state, i.e., no exacerbation within six weeks prior to the test, and HRCT scans are acquired with the patients in the supine position and with breath held. No contrast agents are used.

Volumetric (helical) HRCT scans from both datasets are acquired with the patients in supine position and with breath hold. No contrast agents are used. In the Frederikshavn dataset, the scans are performed on a Siemens SOMATOM Definition Flash CT scanner with scan parameters as follows: 0.6 slice thickness, 95 mAs, 120 kvp, rotation time 0.5 seconds, CTDIvol 7.96 mGy, pitch 1.2 with a image voxel resolution of 0.58×0.588×0.6 mm. In the Aalborg dataset, the scans are achieved using a GE 160 Discovery CT750 HD scanner with with scan parameters as follows: 0.625 slice thickness, 120 kvp, rotation time 0.5 seconds, CTDIvol 5.12 mGy, pitch 0.984, automatically calculated mA by GE’s Smart mA system (max 300 mA and Noise Index 40) and image voxel resolution of 0.788×0.788×0.6 mm.

Note that the dose is higher for the Siemens scan than for the GE scan, however the scans were visually inspected by the two radiologists who performed the clinical validation and they concluded that the visual quality of the scans is similar.

Table 1 presents the clinical characteristics of both datasets. The Fre dataset contains COPD subjects and non-COPD subjects. The non-COPD subjects are referred from the outpatient clinic to have a HRCT scan due to different respiratory problems. Aal contains only subjects with COPD.

3.2 Experimental Setup

Two sets of experiments are performed, which differ in the ways that the positive and negative bags are defined. There are two ways to define the bag labels: by thresholding the COPD stratification (A-D = positive, otherwise = negative), and by thresholding the DLCO predicted value (<60% = low (positive), >60% = high (negative)). The value of 60%, rather than the traditional value of 80% is chosen due to the small sample of subjects with DLCO >80%. Thus, patients with lower DLCO (which could be indicative of COPD, but does not define the COPD diagnosis) are included in the high DLCO class.

To train and evaluate the classifiers, 50 patches with a size of 41 × 41 × 41 pixels are randomly extracted from each HRCT scan. Previous studies [19, 29] have demonstrated that 50 patches is sufficient to classify an entire scan. The patches are selected inside the lung parenchyma using the lung masks.

In each patch, textural features extracted from co-occurrence matrices and Gaussian filter banks are computed. A total of 780 features are computed from co-occurrence matrices, and 320 features are obtained from Gaussian filter banks. High-dimensional feature representations are chosen because previous studies with MIL classifiers and similar feature representations [19, 22] showed good results in terms of bag-level performance. We will also briefly discuss our experiences with lower-dimensional features in the Results section.

Table 1. Clinical characteristics of subjects belonging to both datasets. GOLD stratification reflects the classification of the COPD patients according to the *GOLD combined risk stratification assessment* [3].

Dataset		Gender (M/F)	Age	Smoking			GOLD Stratification				FEV ₁ (%)	DLCO (%)
				current	former	never	A	B	C	D		
Fre	COPD	7/1	66 [48-77]	1	7	0	1	3	3	1	58 [91]	55 [32-90]
	non-COPD	3/5	56 [73]	1	2	5					96 [63-137]	74 [62-83]
Aal		34/38	66 [32-83]	23	48	1	24	12	18	18	62 [18-111]	55 [15-108]

<https://doi.org/10.1371/journal.pone.0205397.t001>

3.2.1 Cross-validation. For each set of experiments, a 4-fold stratified cross-validation is performed, such that each fold contains a similar distribution of subjects. Thus, Fre and Aal datasets are combined and each fold contains non-COPD and COPD subjects with varying degrees of COPD severity, as well as subjects with low and high DLCO values.

The 4-fold cross validation uses 3 folds for training and the fourth fold for evaluation. During training on the 3 folds, an internal 3-fold cross-validation is done to optimize the parameters. These parameters, which are selected only using the 3 folds of the training set, are then used to train a classifier on all the 3 training set folds. The classifier is evaluated on the fourth fold. This is repeated 4 times, so each of the folds is used once for evaluation.

MILES-Q and miSVM-Q classifiers are investigated, and the best parameters for each classifier are selected using the training set. The parameter ranges for both classifiers are as follows: polynomial kernel $p \in \{1, 2\}$, radial basis kernel $rbf \in \{8, 10, 12, 14, 16, 20\}$, regularization parameter $C \in \{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1\}$, and quantile parameter $q \in \{0.25, 0.5, 0.75, 0.9, 1\}$.

3.3 Evaluation

3.3.1 Classifier evaluation. During the 4-fold cross-validation, the best combination of parameters for each classifier is extracted on the three training folds. For the test results, the bag AUC (area under the receiver-operating curve) and Separability (Eq 1) are examined. The bag AUC expresses the ability of the classifier to rank a randomly drawn positive HRCT scan higher than randomly drawn negative HRCT scan (i.e., assign a higher posterior probability to a positive scan). Therefore, AUC is not sensitive to class imbalance: if a classifier assigns all cases to the majority class, the accuracy would be high, but the AUC would be equal to 0.5. The Separability reflects the classifier's ability to distinguish patches with signs of COPD and healthy lung tissue, without having access to such labels. Performance is considered good when the bag AUC is as high as possible and the Separability is as large as possible.

3.3.2 Clinical validation. In the clinical validation, the set of features and the classifier with the best performance in terms of bag AUC and Separability on the training sets are chosen for each of the test folds. The classifier is tested on 10 slices per HRCT scan. This number of slices is chosen to keep manual annotations by the radiologists feasible. The slices are spaced 25 slices apart in each HRCT scan, avoiding the slices belonging to the top and bottom parts of the lungs. In selected slices, the classifier classifies every 10th voxel in both directions in the slice, that is, inside the lung mask for that slice.

Two radiologists, expert 1 and expert 2, with 40 and ten years of experience, respectively, working with HRCT scans on a daily basis annotated all emphysema lesions in the same ten slices per scan in which the classifier is tested. The manual annotations are performed using OsiriX imaging software (www.osirix-viewer.com) using a medical display (BARCO E-2621). The annotation process is blinded. Thus, the experts do not know the outlines of the other expert or the classification results. The amount of emphysema annotated by each expert, in percentage, is computed, as is the percentage of emphysema on which both experts agreed.

For local emphysema detection, the default threshold of 0.5 is used to transform the posterior probabilities into emphysema or healthy category labels.

Spearman correlation analysis is performed, in which the emphysema percentages of the classifier are compared with the manual annotations, results from spirometry and DLCO, and a simple method based on the threshold of LAA. The threshold is set to -950 HU, which has been demonstrated to be an acceptable threshold for density based emphysema quantification [30]. A comparison of correlations from independent samples using the Fisher r-to-z

transformation is computed to assess the significance of the differences between the results from the Spearman correlation.

The inter-observer variability between experts is also investigated using the Dice similarity coefficient (DSC), which is a measurement of similarity. The values of DSC range from 0 if there is no agreement to 1 if there is a perfect match.

4 Results

4.1 Classifier performance

As shown in Table 2, miSVM-Q has higher performance than MILES-Q on both bag AUC and Separability. For miSVM-Q, Gaussian features provide larger Separability and generally better bag classification than co-occurrence features. The combination of co-occurrence matrices and Gaussian features does not improve the results obtained with these features alone. In general, both classifiers can better distinguish obstructions given by low FEV₁ (ClassCOPD) than by low DLCO (ClassDLCO).

Additionally, lower-dimensional feature combinations are briefly investigated, such as orientation-invariant co-occurrence features (60 features) and using only a histogram of intensities (10 features); however, the results were worse than those obtained using the full co-occurrence or Gaussian features.

This suggests that, similar to earlier studies [19, 22] the relatively high feature dimensionality is not detrimental to these classifiers.

To support the classifier evaluation, we additionally visualize one of the high-dimensional feature spaces (Gaussian texture features), which consists of patches from COPD subjects and non-COPD subjects in Fig 2. Note that, due to the MIL representation, there is no straightforward way to cluster the subjects themselves. This dimensionality reduction and visualization is performed using t-SNE [31], a popular method which aims to represent the local structure of the patches in the high-dimensional space as well as possible, without using the labels of the patches. The labels are only added to the plot for visualization. Note that the embedding is the same after rotation, there are therefore no meaningful names that can be assigned to the dimensions, similar to principal component analysis.

We use all the patches from non-COPD subjects, and a random sample (due to class imbalance) of patches from the COPD subjects, for an uncluttered visualization. The visualization shows a large overlapping area with patches that both types of subjects have. However, on the periphery (bottom left, and the right side of the plot) there are regions which only contain patches from COPD subjects. A subject having patches in these areas could therefore be classified as belonging to the COPD class.

4.2 Association with PFTs

Based on the classifier evaluation results, miSVM-Q and Gaussian features are selected for the clinical validation. The parameters of miSVM-Q for the different test folds are selected during cross-validation as before.

Table 2. miSVM-Q and MILES-Q results using COPD (ClassCOPD) and DLCO (ClassDLCO) labels. S: separability (×100); AUC: bag AUC (×100).

Feature	miSVM-Q				MILES-Q			
	DLCO		COPD		DLCO		COPD	
	AUC	S	AUC	S	AUC	S	AUC	S
Cooc	70.9 ± 6.3	4.1	100.0 ± 0.0	44.2	53.0 ± 10.3	0.7	93.8 ± 6.2	17.1
Gauss	81.6 ± 10.2	21.7	100.0 ± 0.0	61.1	69.1 ± 6.6	7.0	89.4 ± 6.4	27.8
Both	59.5 ± 5.4	2.9	95.0 ± 3.5	19.1	50.8 ± 11.3	-0.1	78.8 ± 18.2	17.5

<https://doi.org/10.1371/journal.pone.0205397.t002>

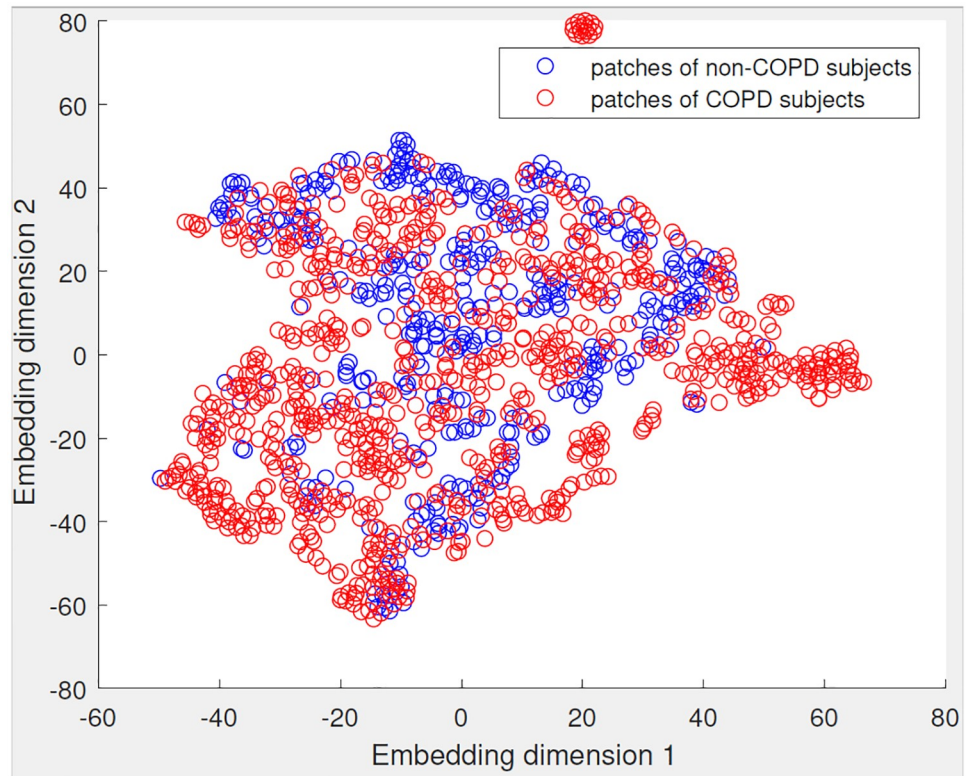


Fig 2. 2D visualization of patches from COPD and non-COPD subjects using the Gaussian feature representation. The patches have different distributions, which helps the MIL classifier to classify a subject correctly.

<https://doi.org/10.1371/journal.pone.0205397.g002>

Table 3 presents the percentage of emphysema detected by the different methods and their correlation with DLCO and FEV₁. The correlations are considered significant at the 0.05 level. Moreover, analysis using the Fisher r-to-z transformation is computed which shows that there is not a significant difference between the correlation coefficients from the Spearman analysis.

4.3 Association with manual annotations

The agreements of the annotations between the two radiologists and between the classifier and the radiologists are investigated. The corresponding scatter plots between the percentage of emphysema calculated by the classifiers and the average percentage of emphysema annotated by the two radiologists are shown in Fig 3.

Furthermore, Spearman correlation analysis are calculated between the percentage of emphysema computed from the manual annotations of the two experts ($\rho = 0.756, p = 1.71$

Table 3. Spearman correlation results with data from pulmonary tests. ClassCOPD: results from classifier with COPD label; ClassDLCO: results from classifier with DLCO label; Thr LAA: Threshold scan based on low attenuation areas; Agree Exp: area of agreement between the manual annotations of both experts; rho: correlation coefficient.

		ClassCOPD	ClassDLCO	Thr LAA	Agree Exp	Expert1	Expert2
DLCO val	rho	-0.477	-0.571	-0.513	-0.478	-0.472	-0.596
	p Value	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001	<.0001
FEV ₁	rho	-0.283	-0.383	-0.461	-0.298	-0.316	-0.314
	p Value	0.016	<0.0001	<0.0001	0.011	0.007	0.007

<https://doi.org/10.1371/journal.pone.0205397.t003>

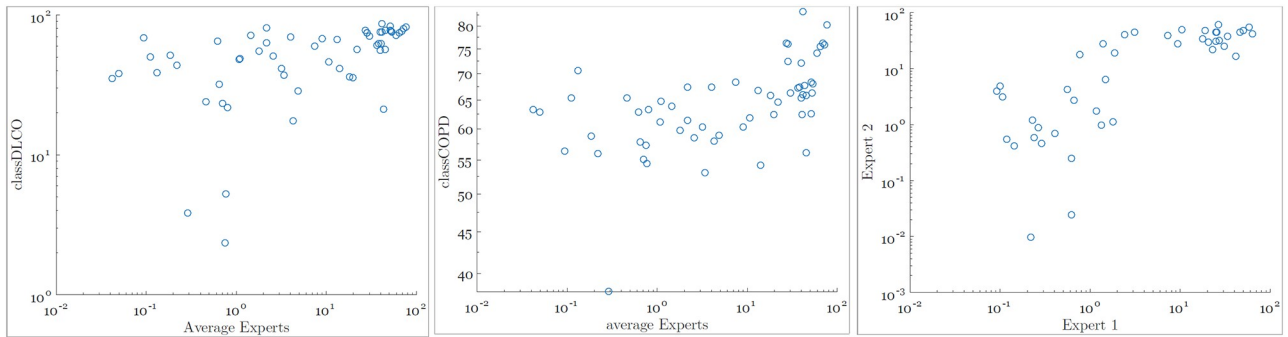


Fig 3. Percentage of emphysema (log scale for visibility) per subject annotated by the experts and computed by the classifiers.

<https://doi.org/10.1371/journal.pone.0205397.g003>

$\times 10^{-14}$), and the percentage of emphysema computed from the classifiers and the average percentage of emphysema from the manual annotations (ClassDLCO: $\rho = 0.561$, $p = 3.01 \times 10^{-7}$; ClassCOPD: $\rho = 0.515$, $p = 4 \times 10^{-6}$).

An example of the results from the classifiers, the manual annotations from the experts, and the threshold using LAA is presented in Fig 4.

Although at instance level, the agreement between the classifier and the experts is not perfect, the emphysema quantification is consistent when using ClassDLCO. In contrast to ClassCOPD, ClassDLCO identifies small emphysema areas in the same patients in which the experts do not make annotations or the annotations are small, and it identifies larger emphysema areas where experts annotate large emphysema lesions.

The agreement between the annotations from both experts is also calculated using DSC which results in a value of 0.34, indicating a weak agreement between experts.

5 Discussion

5.1 Classifier performance

In contrast with previous studies, this study uses an MIL approach to automatically identify emphysema regions in COPD patients without requiring manually annotated HRCT scans for training. Two robust versions of the MILES and miSVM classifiers are presented. Because good bag-level (patient classification) performance does not correspond to good instance-level (patch classification) performance and vice versa [32–34], both bag-level AUC and a measure of instance-level performance, called Separability, are taken into account. The best performing classifier is miSVM-Q. Other studies [19, 21] have shown that instance-level classifiers, such as miSVM, tend to have lower bag-level performance. In the present study, the bag-level performance of miSVM-Q is improved compared to the original miSVM by relaxing the condition that a bag should be classified as positive as soon as a single positive instance is detected.

The performance achieved at the bag-level is very high, with an AUC equal to 100% for the COPD class label. This result could be explained by the fact that half of the subjects in both datasets are in the severe and very severe stage of the GOLD stratification and, therefore, these stages are easier to identify by the classifier. In [19, 22] the same type of labels were used, however an AUC close to 75% was achieved. However, the dataset used was from a screening trial, and thus contained a much higher fraction of mild COPD subjects, which were difficult to classify correctly.

We used several parameters which were determined in previous studies, such as the patch size of 41x41x41 and using 50 patches to represent an entire scan. The results in the current study demonstrate that for these data, these are also reasonable choices. However, this would

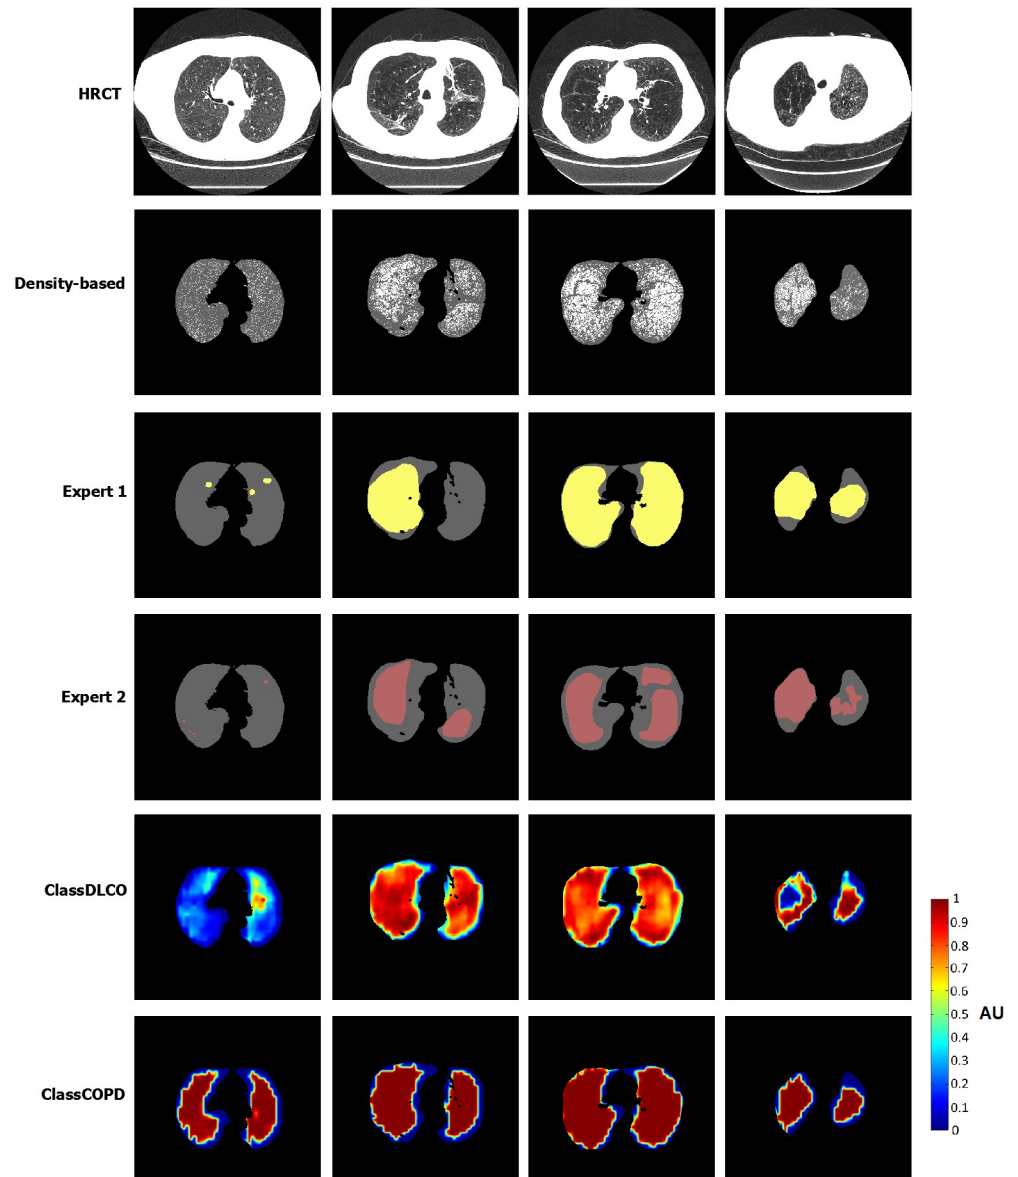


Fig 4. Example of results in randomly selected slices for the density based method, manual annotations from the experts, and classifier results using miSVM-Q and Gaussian features. From left to right: patients with mild, moderate, severe and very severe COPD.

<https://doi.org/10.1371/journal.pone.0205397.g004>

in general depend on a number of factors. For example, we would expect that if the cases in the dataset are all mild, a larger number of patches would be needed to capture some areas with emphysematous tissue. Another factor is the image resolution and therefore the physical size of the patch—for a smaller physical size, we expect that more patches would be needed. There are limits to this trade-off, as a much smaller physical size would not be able to capture the appearance of emphysema.

Earlier studies [19, 22] did not investigate the agreement of the classifier with manual instance-level annotations. Therefore, when making choices such as patch size, number of patches per scan, and so forth, the instance-level performance was not considered.

Consequently, it would be worth investigating how the patch size and number of patches (which in this study are set to the same values as in [19, 22]) would affect the instance-level performance.

Features derived from Gaussian filters, using both classifiers and both labels, provide larger Separability than co-occurrence features or their combinations. All of these feature sets are high-dimensional compared to the size of the data. However, we observed that using lower-dimensional versions of these features reduced the performance, and high-dimensional features have been used with success in previous studies [19, 22]. These observations are consistent with studies of MIL classifiers on non-medical datasets, such as [35], where for example the Web datasets have 2K instances in 5K dimensions, but where miSVM is the second best classifier, and the best classifier that can provide instance-labels. The miSVM classifier is therefore effective at dealing with high dimensional data, by using regularization to reduce overfitting.

An interesting difference with respect to the relative dimensions is between the miSVM-Q and MILES-Q. Because miSVM-Q is an instance-level approach, its effective sample size is the total number of patches used, while the effective sample size of MILES is lower, i.e. the total number of subjects. This could explain why the performances of miSVM-Q are higher overall. This is also consistent with previous results on similar COPD data [19] and non-medical datasets [35]. Therefore, MILES appears to be more prone to overfitting. It would be of interest to investigate how increasing the number of patches per scan would affect the results—for miSVM this would mean an increased sample size, but for MILES it would mean an increased dimensionality.

We did not investigate the use of deep learning methods. This could be done in three ways: training a network from scratch, fine-tuning a pretrained network or using “off-the-shelf” features. Based on recent results in medical imaging [36], we expect that training from scratch would not lead to good results due to the small dataset. We expect fine-tuning and “off-the-shelf” method (which could be used together with the miSVM-Q or MILES classifiers) to be more successful than training from scratch. However, both methods would depend on the data that is used to pretrain the networks, as well as several other parameters. Combining traditional features with features extracted by deep learning has also shown to be effective [37], and would be interesting direction for future work.

5.2 Clinical validation

Spirometry has been widely used as an indicator of COPD severity due to the correlation between FEV_1/FVC and airway obstruction. However, FEV_1 does not reflect structural changes in the lung parenchyma, and therefore, it is not a reliable indicator of emphysema lesions. In contrast, DLCO is a good indicator of the level of anatomic emphysema. In this study, a Spearman correlation analysis between the best classifier from the classifier evaluation, miSVM-Q, and the PFTs is computed. The results in the present study show, as presented in Table 3, that the classifier using both labels has a higher correlation with DLCO values than with FEV_1 . This result is comparable to the result in [38], where the emphysema segmentation using a texture-based approach had a better correlation with DLCO than with values from FEV_1 . This is explained because FEV_1 measures airflow obstruction; however, this is only partially reflected in emphysema lesions. The classifier that is trained on the COPD label based on FEV_1 values, likely detects mostly signs of emphysema and therefore, still correlates better with DLCO than with FEV_1 .

The correlations from the classifier with the PFTs are also compared with the correlations between the PFTs and a density mask method that has been widely used to quantify

emphysema lesions in CT scans. The results show that the density based method correlates moderately better with FEV₁ than both the classifiers and the expert evaluations, and the same behaviour can be observed in [38]. This may be explained by the inability of the density mask to discriminate between air trapping and emphysema due to the nature of their thresholds [39]. However, other studies from the literature that aim to quantify emphysema [16, 22] show a better correlation between FEV₁ with their proposed texture analysis methods than a traditional density based method.

This study uses the PFTs as the most reliable measurement to validate the results of the classifier despite manual annotations by two independent experts being available. This is due to the weak agreement between experts in the annotations of emphysema, as shown by the Spearman correlation results in Section 4.3 and the Dice similarity coefficient. This is in agreement with [40], who showed a low inter-observer agreement in a task of quantifying emphysema in whole lungs. As shown in Fig 3, ClassCOPD overestimates the amount of emphysema in comparison with the manual annotations. This result can be produced by the small dataset of non-COPD. However, ClassDLCO generally tends to agree with experts on the size of emphysema areas. The scatter plots show that ClassDLCO has a good agreement with the experts' annotations in severe and very severe cases, but the agreement is fair in moderate patients and poor in mild patients. ClassDLCO overestimates the emphysema in these cases.

However, these findings in conjunction with the improved correlation with DLCO compared to manual annotations could indicate that ClassDLCO is more sensitive to early changes in the lung parenchyma and can detect emphysema even before these changes are able to be detected visually. To confirm this result, future studies should investigate the progression of emphysema in the areas where the classifier finds emphysema, but that were not assessed visually. This will also help to reduce inter-observer variability, which is a major limitation in visual assessment, as other studies have reported [4, 41]. Furthermore, the correlation results between the ClassDLCO and DLCO values show that quantitative assessment of emphysema with the presented method provides an important measurement of the reduction in the alveolar area. In addition, as suggested in [42], a better detection of emphysema in HRCT scans can also be used in refining the prediction of the 6 minute walking distance test.

5.3 Limitations

A limitation of this study is the size and balance of the datasets. The Fre dataset is very small, and the Aal dataset does not contain any controls. Due to this imbalance, the DLCO threshold used is lowered to 60%; thus, patients with lower DLCO values are included in the high DLCO class (which could be seen as healthy, although note that the COPD diagnosis is based on spirometry). The texture features extracted from the non-COPD group could have a similar representation as the features extracted from COPD patients because different lung diseases could appear in CT scans as LAA as emphysema does, i.e. cystic lung disease. Therefore, it would be desirable to include scans without any pathology.

A related problem is the fact that the Fre and Aal datasets have different acquisition parameters, which can negatively affect classification performance [43]. An improvement in performance will be expected if the appearance of healthy tissue could be learned from both datasets rather than only Fre. An alternative would be to use techniques such as intensity normalization or transfer learning classifiers to reduce the differences between datasets.

6 Conclusion

This study presented two new versions of multiple instance classifiers which identify emphysema regions in patients suffering from COPD without requiring manual annotations. On a

clinical dataset with 88 subjects in total, the proposed method showed a good correlation with the pulmonary function tests, particularly with DLCO. The proposed method had a moderate correlation with manual annotations of emphysema, however, this correlation was higher than that of a density based method, which was also moderately correlated with manual annotations. Therefore it could be considered as a reliable tool to support radiologists in the assessment of emphysema to reduce the inter- and intra-observer variability. As future work, validating the results on a larger and more balanced dataset, and investigating the effect of different acquisition parameters, are recommended.

Author Contributions

Conceptualization: Isabel Pino Peña, Veronika Cheplygina, Sofia Paschaloudi, Morten Vuust, Jesper Carl, Ulla Møller Weinreich, Lasse Riis Østergaard, Marleen de Bruijne.

Data curation: Isabel Pino Peña, Sofia Paschaloudi, Morten Vuust, Jesper Carl, Ulla Møller Weinreich.

Formal analysis: Isabel Pino Peña, Veronika Cheplygina, Marleen de Bruijne.

Investigation: Isabel Pino Peña, Veronika Cheplygina, Sofia Paschaloudi, Lasse Riis Østergaard, Marleen de Bruijne.

Methodology: Isabel Pino Peña, Veronika Cheplygina, Lasse Riis Østergaard, Marleen de Bruijne.

Software: Isabel Pino Peña, Veronika Cheplygina.

Supervision: Jesper Carl, Ulla Møller Weinreich, Lasse Riis Østergaard, Marleen de Bruijne.

Validation: Isabel Pino Peña, Veronika Cheplygina, Sofia Paschaloudi, Morten Vuust, Jesper Carl, Lasse Riis Østergaard, Marleen de Bruijne.

Visualization: Isabel Pino Peña, Veronika Cheplygina, Sofia Paschaloudi, Morten Vuust, Ulla Møller Weinreich.

Writing – original draft: Isabel Pino Peña, Veronika Cheplygina.

Writing – review & editing: Isabel Pino Peña, Veronika Cheplygina, Sofia Paschaloudi, Morten Vuust, Jesper Carl, Ulla Møller Weinreich, Lasse Riis Østergaard, Marleen de Bruijne.

References

1. Mathers C, Loncar D. Projections of Global Mortality and Burden of Disease from 2002 to 2030. *PLOS Medicine*. 2006; 3:1–20. <https://doi.org/10.1371/journal.pmed.0030442>
2. WHO, The top 10 causes of death;. Available from: <http://www.who.int/mediacentre/factsheets/fs310/en/>.
3. GOLD. Global Strategy for the Diagnosis, Management, and Prevention of Chronic Obstruction Pulmonary Disease. 2015;(January).
4. Ginsburg SB, Lynch DA, Bowler RP, Schroeder JD. Automated Texture-based Quantification of Centrilobular Nodularity and Centrilobular Emphysema in Chest CT Images. *Academic Radiology*. 2012; 10:1241–1251. <https://doi.org/10.1016/j.acra.2012.04.020>
5. Matsuoka S, Yamashiro T, Washko GR, Kurihara Y, Nakajima Y, Hatabu H. Quantitative CT assessment of chronic obstructive pulmonary disease. *Radiographics*. 2010; 30:55–66. <https://doi.org/10.1148/rg.301095110> PMID: 20083585
6. Lynch DA, Newell JD. Quantitative imaging of COPD. *Journal of Thoracic Imaging*. 2009; 3:189–194. <https://doi.org/10.1097/RTI.0b013e3181b31cf0>
7. Nakano Y, Muro S, Sakai H, Hirai T, Chin K, Tsukino M, et al. Computed tomographic measurements of airway dimensions and emphysema in smokers. Correlation with lung function. *American journal of*

- respiratory and critical care medicine. 2000; 162:1102–8. <https://doi.org/10.1164/ajrccm.162.3.9907120> PMID: 10988137
8. Nishimura K, Murata K, Yamagishi M, Itoh H, Ikeda A, Tsukino M, et al. Comparison of Different Computed Tomography Scanning Methods for Quantifying Emphysema. *Journal of Thoracic Imaging*. 1998; 13:193–198. <https://doi.org/10.1097/00005382-199807000-00006> PMID: 9671422
 9. Mets OM, Jong PA, Ginneken B, Gietema HA, Lammers JWW. Quantitative Computed Tomography in COPD: Possibilities and Limitations. *Lung*. 2012; 190:133–145. <https://doi.org/10.1007/s00408-011-9353-9> PMID: 22179694
 10. Sørensen L, Shaker SB, Bruijine M. Quantitative Analysis of Pulmonary Emphysema Using Local binary Patterns. *IEEE Transactions on medical imaging*. 2010; 29:559–569. <https://doi.org/10.1109/TMI.2009.2038575> PMID: 20129855
 11. Nagao J, Aiguchi T, Mori K, Suenaga Y, Toriwaki J, Mori M, et al. A CAD system for quantifying COPD Based on 3-D CT images. *Medical Imaging Computing and Computer Assisted Interventions*. 2003; 2878:730–737.
 12. Yao J, Dwyer A, Summers RM, Mollura DJ. Computer-aided Diagnosis of Pulmonary Infections Using Texture Analysis and Support Vector Machine Classification. *Academic Radiology*. 2011; 18:306–314. <https://doi.org/10.1016/j.acra.2010.11.013> PMID: 21295734
 13. Bagci U, Bray M, Caban J, Yao J, Mollura DJ. Computer-Assisted Detection of Infectious Lung Diseases: A review. *Computerized Medical Imaging and Graphics*. 2012; 36:72–84. <https://doi.org/10.1016/j.compmedimag.2011.06.002> PMID: 21723090
 14. Prasad M, Sowmya A, Wilson P. Multi-level classification of emphysema in HRCT lung images. *Pattern Analysis and Applications*. 2007; 12:9–20. <https://doi.org/10.1007/s10044-007-0093-7>
 15. Uppaluri R, Hoffman EA, Sonka M, Hunninghake GW, McLennan G. Computer Recognition of Regional Lung Disease Patterns. *American Journal of Respiratory and Critical Care Medicine*. 1999; p. 648–654. <https://doi.org/10.1164/ajrccm.160.2.9804094> PMID: 10430742
 16. Park YS, Seo JB, Kim N, Chae EJ, Oh YM, Lee SDL, et al. Texture-Based Quantification of Pulmonary Emphysema on High-Resolution Computed Tomography: Comparison with Density-Based Quantification and Correlation with Pulmonary Function Test. *Investigative radiology*. 2008; 43:395–402. <https://doi.org/10.1097/RLI.0b013e31816901c7> PMID: 18496044
 17. Kim N, Seo JB, Lee Y, Lee JG, Kim SS, Kang SH. Development of an Automatic Classification System for Differentiation of Obstructive Lung Disease using HRCT. *Journal of Digital Imaging*. 2009; 2:136–148. <https://doi.org/10.1007/s10278-008-9147-7>
 18. Chabat F, Yang GZ, Hansell DM. Obstructive Lung Diseases: Texture Classification for Differentiation at CT. *Radiology*. 2003; 228:871–877. <https://doi.org/10.1148/radiol.2283020505> PMID: 12869685
 19. Cheplygina V, Sørensen L, Tax DMJ, Pedersen JH, Loog M, de Bruijine M. Classification of COPD with Multiple Instance Learning. *Proceedings International Conference on Pattern Recognition*. 2014; p. 1508–1513.
 20. Melendez J, van Ginneken B, Maduskar P, Philipsen RH, Reither K, Breuninger M, et al. A novel multiple-instance learning-based approach to computer-aided detection of tuberculosis on chest x-rays. *IEEE Transactions on Medical Imaging*. 2015; 34(1):179–192. <https://doi.org/10.1109/TMI.2014.2350539> PMID: 25163057
 21. Kandemir M, Hamprecht FA. Computer-aided diagnosis from weak supervision: A benchmarking study. *Computerized Medical Imaging and Graphics*. 2015; 42:44–50. <https://doi.org/10.1016/j.compmedimag.2014.11.010> PMID: 25475486
 22. Sørensen L, Nielsen M, Lo P, Ashraf H, Pedersen JH, de Bruijine M. Texture-Based Analysis of COPD: A Data-Driven Approach. *Medical Imaging, IEEE Transactions on*. 2012; 31(1):70–78. <https://doi.org/10.1109/TMI.2011.2164931>
 23. Albrechtsen F. *Statistical Texture Measures Computed from Gray Level Cooccurrence Matrices*; 2008.
 24. Ojala T, Pietikäinen M, Harwood D. A comparative study of texture measures with classification based on featured distributions. *Pattern recognition*. 1996; 29(1):51–59. [https://doi.org/10.1016/0031-3203\(95\)00067-4](https://doi.org/10.1016/0031-3203(95)00067-4)
 25. Zhou ZH, Zhang ML. Multi-instance multi-label learning with application to scene classification. In: *Advances in Neural Information Processing Systems*; 2006. p. 1609–1616.
 26. Amores J. Multiple instance classification: Review, taxonomy and comparative study. *Artificial Intelligence*. 2013; 201:81–105. <https://doi.org/10.1016/j.artint.2013.06.003>
 27. Andrews S, Tsochantaridis I, Hofmann T. Support vector machines for multiple-instance learning. In: *Advances in Neural Information Processing Systems*; 2002. p. 561–568.

28. Chen Y, Bi J, Wang JZ. MILES: Multiple-instance learning via embedded instance selection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2006; 28(12):1931–1947. <https://doi.org/10.1109/TPAMI.2006.248> PMID: 17108368
29. Sørensen L L Pechin, Ashraf H, Sparring J, Nielsen M, Bruijine M. Learning COPD sensitive filters in pulmonary CT. *Lecture Notes in Computer Science*. 2009;5762:699–706.
30. Wang Z, Gu S, Leader JK, Kundu S, Tedrow JS, Scirba FC, et al. Optimal threshold in CT quantification of emphysema. *European Radiology*. 2013; p. 975–984. <https://doi.org/10.1007/s00330-012-2683-z> PMID: 23111815
31. Maaten Lvd, Hinton G. Visualizing data using t-SNE. *Journal of Machine Learning Research*. 2008; 9 (Nov):2579–2605.
32. Cheplygina V, Sørensen L, Tax DMJ, de Bruijine M, Loog M. Label Stability in Multiple Instance Learning. In: *Medical Image Computing and Computer-Assisted Interventions*. Springer; 2015.
33. Tragante do O V, Fierens D, Blockeel H. Instance-level accuracy versus bag-level accuracy in multi-instance learning. In: *Benelux Conference on Artificial Intelligence*; 2011.
34. Vanwinckelen G, Tragante do O V, Fierens D, Blockeel H. Instance-level accuracy versus bag-level accuracy in multi-instance learning. *Data Mining and Knowledge Discovery*. 2016; p. 313–341. <https://doi.org/10.1007/s10618-015-0416-z>
35. Cheplygina V, Tax DMJ, Loog M. Multiple instance learning with bag dissimilarities. *Pattern Recognition*. 2015; 48(1):264–275. <https://doi.org/10.1016/j.patcog.2014.07.022>
36. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. *Medical image analysis*. 2017; 42:60–88. <https://doi.org/10.1016/j.media.2017.07.005> PMID: 28778026
37. Nanni L, Ghidoni S, Brahmam S. Handcrafted vs. non-handcrafted features for computer vision classification. *Pattern Recognition*. 2017; 71:158–172. <https://doi.org/10.1016/j.patcog.2017.05.025>
38. Tan J, Zheng B, Wang X, Lederman D, Pu J, Scirba FC, et al. Emphysema quantification in a multi-scanner HRCT cohort using local intensity distributions. In: *Medical Imaging 2011: Biomedical Applications in Molecular, Structural, and Functional Imaging*. SPIE; 2011. p. 1–7.
39. Voelkel NF, MacNee W. *Chronic Obstructive Lung Disease*. PMPH USA, Ltd.; 2008.
40. Mascalchi M, Diciotti S, Sverzellati N, Camiciottoli G, Ciccotosto C, Falschi F, et al. Low agreement of visual rating for detailed quantification of pulmonary emphysema in whole-lung CT. *Acta Radiologica*. 2012; 53:53–60. <https://doi.org/10.1258/ar.2011.110419> PMID: 22114019
41. Barr RG, Berkowitz E, Bigazzi F, Bode F, Bon J, Bowler RP, et al. A combined pulmonary-radiology workshop for visual evaluation of COPD: study design, chest CT findings and concordance with quantitative evaluation. *COPD*. 2012; 9:151–159. <https://doi.org/10.3109/15412555.2012.654923> PMID: 22429093
42. Diaz AA, Pinto-Plata V, Hernandez C, Peña J, Ramos C, Diaz JC, et al. Emphysema and DLCO predict a clinically important difference for 6MWD decline in COPD. *Respiratory Medicine*. 2015; 109:882–889. <https://doi.org/10.1016/j.rmed.2015.04.009> PMID: 25952774
43. de Bruijine M. Machine learning approaches in medical image analysis: from detection to diagnosis. *Medical Image Analysis*, in press. 2016;. <https://doi.org/10.1016/j.media.2016.06.032> PMID: 27481324