

Rotation-covariant tissue analysis for interstitial lung diseases using learned steerable filters: Performance evaluation and relevance for diagnostic aid



Ranveer Joyseeree^{a,*}, Henning Müller^a, Adrien Depeursinge^{a,b}

^a Institute of Information Systems, University of Applied Sciences Western Switzerland (HES-SO), Rue de Technopôle 3, 3960 Sierre, Switzerland

^b Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), Lausanne 1015, Switzerland

ARTICLE INFO

Keywords:

Texture signatures
One-versus-one
Classification
ILD

ABSTRACT

A novel method to detect and classify several classes of diseased and healthy lung tissue of interstitial lung diseases is presented, as these diseases are hard to diagnose and differentiate. Local organizations of image directions at several scales drive the process of creating discriminative lung tissue texture signatures using spatial and Fourier domain information extracted from the images. The signatures are generated for four diseased tissue classes and healthy tissue, all of which appear in the Interstitial Lung Disease (ILD) database, using a novel one-versus-one approach for learning discriminative filter signatures. A multiclass tissue classification accuracy of 80.31% is observed using Radial Basis Function (RBF) Support Vector Machines (SVMs). The presented method compares well against a variety of state-of-the-art approaches. Another strong feature of our approach is the ability to access the individual class probabilities before a final classification decision is made. This enables an analysis of the causes of misclassification in this paper. We also make the case against total reliance on the accuracy of the ground truth given that the ILD database only contains a single label for a specific region and sometimes more than one pattern can be present, particularly for regions classified as healthy tissue. Measures to address misclassifications in this context are also proposed.

1. Introduction

Respiratory diseases are a leading cause of premature mortality according to the White Book of the European Respiratory Society (ERS) (European, 2016). In 2008, 9.5 million deaths were attributed to them, corresponding to 1 in every 6 deaths worldwide. In addition, they accounted for an annual cost of 380 billion Euros for the European Union (EU) alone in terms of direct medical care, prescription of drugs, and loss of productive output. Tackling such diseases is therefore a priority in the healthcare sector and is also of economic significance. To prevent avoidable deaths and costs, it is vital that an accurate diagnosis is found as early as possible. Clinicians can then administer the correct treatment directly. This can provide the patient either with the highest achievable quality of life while living with the disease and with the best chance of survival.

Among such diseases, those affecting lung parenchyma are very common. Interstitial lung diseases (ILDs) form a large and diverse group of such diseases. They consist of more than 200 pathologies affecting the small lung airways, the pulmonary interstitium, and the alveoli (Kreuter et al., 2015). Their accurate diagnosis is based on information

gathered from clinical, radiological, and pathological analyses.

High Resolution Computed Tomography (HRCT) images are typically used for the detection of ILDs due to the superior amount of details they provide over other imaging modalities (Barr et al., 2016). Radiologists are subject to a certain degree of subjectivity when interpreting the content of the images. As many diseases are rare and clinical diagnoses rely on experience, some ailments can easily be misdiagnosed. Ideally, an objective and accurate approach for detecting interstitial lung pathologies is required and we contribute towards it in this paper through an automatic interstitial lung disease detection and classification method that leverages the advances made in the field of visual pattern recognition. In addition, our approach leads towards a quantitative analysis of the prevalence of diseased tissue in a patient. This can facilitate clinical decision-making by giving probabilities for different tissue types.

To fully exploit visual pattern recognition in the detection of lung tissue pathologies, we note that pattern discrimination relies heavily on local organization of scales and directions in the image being analyzed (Blakemore and Campbell, 1969; ter Haar Romeny, 2010). An analysis of scales and directions in tissue patterns needs to be done together

* Corresponding author.

E-mail address: ranveer@icloud.com (R. Joyseeree).

URL: <http://medgift.hevs.ch/wordpress/> (R. Joyseeree).

because co-occurrence of the two provides vital discriminatory properties. It is also important to avoid the loss of local variations of pattern properties such as local anisotropy, as they might be key in discriminating between tissue texture types (Depeursinge et al., 2014a). The spatial domain representation of images provides insufficient information to properly examine the local organization of scales and directions. To obtain a more complete overview of the relationships between them, we complement intensity information with frequency domain information.

Many techniques that explicitly exploit the local organization of scales and directions exist. For the case of directional analysis, separable and non-separable wavelets (Jeng-Shyang and Jing-Wein, 1999), grey-level co-occurrence matrices (GLCM) (Haralick et al., 1973), run-length encoding (RLE) (Xu et al., 2004), histograms of gradients (HOG) (Dalal and Triggs, 2005) used in the scale-invariant feature transform (SIFT) (Lowe, 2004), and oriented filterbanks and wavelets (Gaussian, Gabor, Leung-Malik, Maximum Response (Cula and Dana, 2004; Leung and Malik, 2001; Porter and Canagarajah, 1997; Randen and Husoy, 1999; Xu et al., 2010) have been proposed. Separable wavelets unfortunately present bias along the horizontal and vertical directions (Mallat, 1989) and the remaining methods all require an arbitrary choice of image directions for directional analysis (Depeursinge et al., 2014a). Local binary patterns (LBP) (Ojala et al., 2002) are also commonly used and they perform multidirectional analysis using a sequence of pixels along perimeters of radius r but do not come with multiresolution capabilities and r has to be determined through manual optimization. To overcome such challenges, methods based on brushlets (Meyer and Coifman, 1997), curvelets (Candès and Donoho, 2000), ridgelets (Donoho, 2001) and contourlets (Do and Vetterli, 2002; Po and Do, 2006) were developed. They allow comprehensive multidirectional analysis for neighborhoods of any radius.

Other techniques in the literature leverage the local organization of scales and directions indirectly. Such methods include convolutional neural networks (CNN) (LeCun et al., 2004, 2010), the scattering transform (Ablowitz et al., 1974; Ablowitz and Segur, 1981), and topographic independent component analysis (TICA) (Hyvärinen et al., 2001). Another approach is based on steerable filterbanks (Freeman and Adelson, 1991; Simoncelli and Farid, 1996) that allow continuous directional characterization (Depeursinge et al., 2014a). For scale analysis, (Mallat, 1989) provides an elegant multiresolution approach whereby a wavelet representation is defined after study of the information difference between successive resolutions of a signal. Its usefulness was demonstrated through its application to texture discrimination.

Another point to note is that the same texture pattern can exist in different local orientations. The approach described in this paper needs to be able to deal with such variations. Rotation invariance or covariance in the features extracted from images is a solution. Invariance implies that the features are the same even if the images undergo rotation while covariance implies that the features have a bilinear relationship with the rotational transformation. Digging deeper into this notion, however, one discovers that, in actual fact, rotation invariance (implemented in Schmid, 2001; Van De Ville et al., 2005) cannot deal with certain important directional properties such as local anisotropy. That is the case because one of the most common ways of achieving invariance is to average the output of directional operators over all directions (Depeursinge et al., 2014a). We clearly require rotation-covariance, which retains information relating to local directional variations. Several methods exhibiting rotation covariance exist, such as rotation-covariant LBP (Ojala et al., 2002) and rotation-covariant SIFT (also known as RIFT) (Lazebnik et al., 2005). These, however, require exhaustive calculations and do not model discriminative patterns specifically (Depeursinge et al., 2014a). Steerable filters are able to generate rotation-covariant features without these difficulties (Depeursinge et al., 2011a; Do and Vetterli, 2002).

Once suitable features for optimal discrimination between diseased

tissue types can be reliably found, an appropriate classification technique is necessary. For each region of interest being analyzed, a vector of discriminative features can be created. The basic idea from here on is to compute the vector of features for regions with unknown disease label and then automatically assign the label with highest probability to each of them. Machine-learning-based classification is perfect for this type of inference. Machine learning techniques exist in many different forms. Supervised learning methods employ labeled training data while unsupervised techniques do not use any labeled data. The goal in the latter is usually to find clusters of similar data so they can be grouped together as belonging to the same category. Semi-supervised learning algorithms are a mix of the above two approaches and work well on partially labeled datasets or data containing errors in the labels. Many classifiers are in regular use in the field to achieve automatic identification of the labels of image regions. Random Forests (RF) (Breiman, 2001) are an example of decision-tree learning. They are an ensemble of weak classifiers and often yield excellent results in practice. In Zarzar et al. (2015), for instance, they were successfully used in cancer diagnosis. Artificial neural networks are also often used as classifiers. They were designed to mimic the network of neurons in the human brain. One example is the Multilayer Perceptron (MLP) (Ruck et al., 1990). It has, for instance, successfully been used for heart disease diagnosis (Yan et al., 2006). The main problem with neural networks is that they operate as black boxes with intermediate levels between image input and classification output that are often difficult to interpret. A Convolutional Neural Network (CNN) has been used on a lung disease database in Gao et al. (2016) where classification performance was not assessed explicitly and in Li et al. (2014) with a reasonable level of accuracy, precision, and recall. Support Vector Machines (SVM) (Cortes and Vapnik, 1995) are very popular as well and represent an instance of a non-probabilistic binary linear classifier. They aim at finding an optimal separation hyperplane between two classes. In the literature, they have been used for cancer diagnosis with success (Guyon et al., 2002). Although very accurate, they have high computational complexity and large memory demands. In addition, it is not straightforward to set the optimal parameters and performance can vary if not well optimized.

Since the ILD dataset is being used to evaluate performance in this paper, it has to be noted that it was annotated by two radiologists in agreement with each other. As a result, the ground truth can contain a level of subjectivity and inter-rater disagreement cannot be measured. Many publications have, however, made use of that database and results therein provide a useful indication of relative performance to compare with. It can be pointed out that classification methods that abstain from classifying data with limited annotation quality exist and are important to consider (Condessa et al., 2017). Common classification performance measures are not adapted to evaluate the performance of such classifiers and three other evaluation measures are proposed in Condessa et al.: non-rejected accuracy, classification quality and rejection quality. In our work, we tackle this problem by moving a threshold on the classification likelihood before making a final decision.

The rest of this paper is structured as follows. We begin by presenting the database used to validate the proposed approach. Work done by other researchers using the same data and similar validation schemes are also covered. This is followed by an in-depth description of the theoretical foundations of the proposed method. Subsequently, the results are displayed. These are then analyzed thoroughly to interpret all observed behavior. Finally, we draw conclusions from the work done and propose some measures for further improvement in future work.

2. Methods

An overview of the proposed method is provided in Fig. 1. A suitable database for training and testing is first chosen. A steerable Riesz filterbank is then applied to the training set to obtain a one-versus-one

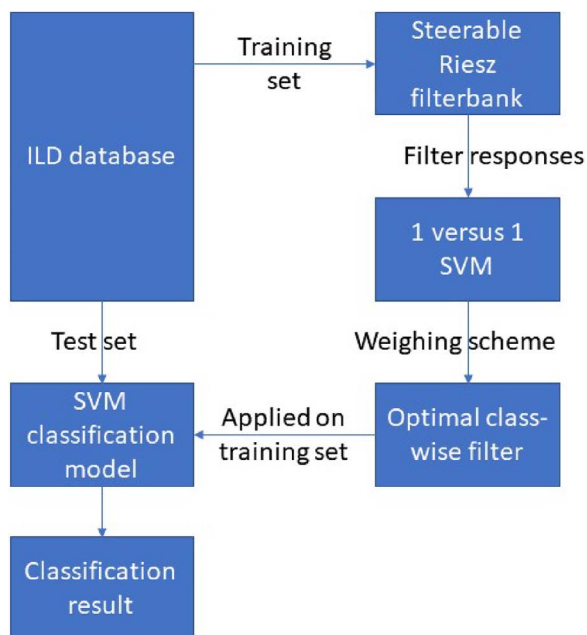


Fig. 1. The workflow for the proposed method first identifies the database used for training and testing. A steerable Riesz filterbank is then used on the training set to obtain a one-versus-one SVM that yields a weighing scheme for the above filterbank. That is used to generate optimal class-wise filters that are then applied on the training set to get an SVM classification model. Applying this on the test set yields the desired classification results.

SVM model. This is used to generate a weighing scheme for the above filterbank.

A weighted sum of the components of the filterbank lead to an optimal filter for each class of tissue considered. Applying these optimal filters on the training set allows for the creation of an SVM classification model that is subsequently used to classify lung tissue.

More details on each step of the above process is provided in the rest of this section.

2.1. Data used

To assess the effectiveness of our method in diagnosing ILDs, we use a publicly available ILD database for validation (Depeursinge et al., 2012b). It consists of 108 HRCT volumes corresponding to 13 histological diagnoses of ILDs. A total of 1946 ROIs were annotated by two radiologists in agreement, which translates to more than 41 liters of annotated tissue. Fig. 2 illustrates one slice from the database showing a cross-section of the thorax of a patient along with an expert-generated annotation highlighted in the mid-left of the slice. Several publications have used the same dataset for validation and they will be described




Fig. 2. A sample slice from an image in the ILD database with highlighted annotation is presented here. The annotated tissue is the bright area in the mid-left of the slice.

briefly. In the results section, the reported performance values are compared to ours. The exact evaluation scenario can vary between publications and thus not all of the results are directly comparable. The radiologists explicitly aimed to find healthy tissue in all patients eventhough this can have some abnormal patters mixed in the tissue.

Many authors tried to classify five lung tissue types: healthy, emphysema, ground glass, fibrosis, and micronodules. Early on, Depeursinge et al. (2007, 2008) worked on an incomplete version of the ILD dataset with fewer images and combined clinical parameters with image data. In further work, Depeursinge et al. (2012c) use isotropic wavelet frames for near-affine-invariant texture learning. They also used the Riesz transform for multiscale texture signature learning (Depeursinge et al., 2012a). In yet further work, they apply a method based on steerable Riesz filterbanks on the ILD database (Depeursinge et al., 2011a). Low-level localized features are used in Depeursinge et al. (2011b) on the same set of classes. The evaluation methodologies slightly varied over the years switching to leave-one-patient-out early.

Song et al. (2013) use feature-based image patch approximation for classification while in Song et al. (2015), a locality-constrained sub-cluster representation ensemble is used. Li et al. (2014) use a customized CNN approach to classify ILD images. Gao et al. (2016) use a deep CNN approach. As for Li et al. (2013), they use automatic feature learning for image patch classification. Some authors extended the set of classes considered to six by including the consolidation class. For instance, Shin et al. (2016) use a deep CNN for classification on this set. Similarly, Foncubierta-Rodríguez et al. (2012) use the six classes to test an implementation of multiscale visual words for classification and retrieval. Finally, it is also possible to use almost completely different classes for classification. For example, Anthimopoulos et al. (2016) apply a deep CNN to the following classes: healthy, ground glass opacity (GGO), micronodules, consolidation, reticulation, honeycombing, and a combination of GGO and reticulation. Thus, even though the data set is the same there can be several ways the evaluation can be done and for this reason absolutely performance comparisons need to be taken with care.


We use supervised learning (the ILD dataset has regions of interest annotated) on the five classes that are most commonly encountered in the literature, namely healthy, emphysema, ground glass, fibrosis, and micronodules. This facilitates performance comparisons with many of the methods above. Furthermore, these classes have a reasonably high representation, with other classes containing fewer examples. Fig. 3 illustrates their respective appearance and it is clear that the difference in appearance can occasionally be subtle. This is especially noticeable for healthy tissue, emphysema and micronodules.

For training and testing on the dataset, a leave-one-patient-out (LOPO) cross-validation strategy is employed. This means that the patches belonging to one patient are removed from the pool for training and are then used for testing. The remaining patches are taken to train a classification model. The model is then used to classify the patches set aside. The predictions on the test set are compared with the actual class values in order to evaluate the classification accuracy. If only one patch is left out instead of a full patient, bias can lead to much better results and not in all publications this is explicitly mentioned. Finally, the whole process is repeated such that each patient in the ILD dataset is set aside once and all patches are classified.

We do not use 3D regions of interest as the data base has anisotropic data, meaning that the slice thickness is 1 mm and pixel size is 1 mm in all three directions, but there is a 10 mm gap between two slices. This missing information cannot easily be reconstructed.

2.2. Tissue characterization

Through this work, we aim to generate texture features that are multiscale, multidirectional, and rotation-covariant. From the methods presented in this paper so far, steerable filterbanks based on wavelets seem to be the most appropriate approach to follow. The use of



Healthy Emphysema Ground glass Fibrosis Micronodules

Fig. 3. Five classes were selected for classification from the ones available in the ILD database with a high representation. They represent healthy tissue as well as tissue representative of emphysema, ground glass, fibrosis, and micronodules.

multiresolution wavelets caters to the need for multiscaleability while the steerable filterbanks allow for continuous directional characterization along with rotation-covariance. Additionally, no assumption is made on the scales and directions and there is no need for arbitrary choices of neighborhood radii. The wavelet-based approach is also easy to adapt to the problem being resolved (Castellano et al., 2004). Specifically, we concentrate on the use of Riesz filterbanks (Depeursinge et al., 2014b,a, 2013, 2012a, 2011a; Held et al., 2010) as much recent work has been done with this technique with very good results on texture.

2.2.1. Building Riesz filters with desirable properties

We begin the process of detecting and classifying lung tissue types by observing that the information contained in any 2D patch can be represented with pixel position (x and y coordinates) and grayscale value. The information contained in a patch can therefore be mathematically modeled as a function having finite boundaries and, as a result, finite energy content. Let $f(\mathbf{x})$ represent the finite energy function that models the content of a patch where \mathbf{x} encodes the pixel coordinates within the patch in the two-dimensional space. In other words, $f: \mathbf{x} \rightarrow f(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^2$, where $\mathbf{x} = (x_1, x_2)$.

Since texture information is solely encoded in the spatial transitions between the pixel values, the characterization of the imaging features is best described in the Fourier domain in terms of spatial frequencies. We choose the Fourier domain as it is fast and straightforward to compute using the Fast Fourier Transform (FFT) (Brigham, 1988). The Fourier domain representation of $f(\mathbf{x})$ is defined as:

$$f(\mathbf{x}) \xleftrightarrow{\mathcal{F}} \hat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^2} f(\mathbf{x}) e^{-j\langle \boldsymbol{\omega}, \mathbf{x} \rangle} dx_1 dx_2, \quad (1)$$

where $\boldsymbol{\omega} = (\omega_1, \omega_2)$ and $\langle \cdot, \cdot \rangle$ denotes the scalar product.

We first define the Riesz filterbanks needed for our work, which are based on the real Riesz transform (Unser et al., 2011). For a 2D signal, the $N + 1$ components of the N th-order Riesz transform are defined as

$$\mathbf{R}^N \{f\}(\mathbf{x}) = \begin{pmatrix} \mathcal{R}^{(0,N)} \{f\}(\mathbf{x}) \\ \vdots \\ \mathcal{R}^{(n,N-n)} \{f\}(\mathbf{x}) \\ \vdots \\ \mathcal{R}^{(N,0)} \{f\}(\mathbf{x}) \end{pmatrix}, \quad (2)$$

with $n = 0, 1, \dots, N$. A kernel $\mathcal{R}^{(n,N-n)} \{f\}(\mathbf{x})$ that represents the action of a generic member of the filterbank on the signal is defined in the spatial and Fourier domains as:

$$\mathcal{R}^{(n,N-n)} \{f\}(\mathbf{x}) \xleftrightarrow{\mathcal{F}} \mathcal{R}^{(n,N-n)} \{f\}(\boldsymbol{\omega}),$$

where

$$\mathcal{R}^{(n,N-n)} \{f\}(\boldsymbol{\omega}) = \frac{\sqrt{\frac{N}{n!(N-n)!}}}{\frac{(-j\omega_1)^n (-j\omega_2)^{N-n}}{\|\boldsymbol{\omega}\|^N}} \hat{f}(\boldsymbol{\omega}). \quad (3)$$

In Eq. (3), multiplication by $j\omega_1$ and/or $j\omega_2$ in the numerator followed by division with the norm of $\boldsymbol{\omega}$ produces filterbanks that behave as allpass filters that only retain the phase information defining directions according to partial image derivatives of maximal order N (Depeursinge et al., 2014a; Unser and Van De Ville, 2010). N determines the angular selectivity of the Riesz kernels. Fig. 4 illustrates a representation of the Riesz kernels for $N = 1, \dots, 5$.




Fig. 4. A representation of the Riesz filterbanks for orders 1 to 5 are illustrated here. The Riesz transform was applied to an isotropic Gaussian function to represent the filters on a finite spatial support.

One of the most interesting properties of the Riesz filterbanks is steerability (Freeman and Adelson, 1991; Unser and Van De Ville, 2010). What this essentially means is that a linear combination of the filterbanks can be used to model any local rotation and this represents an advantage over such techniques as CNNs where rotation-invariance is achieved through the expensive process of augmenting the original training images with several rotated versions.

For a given rotation angle θ , a steering matrix A_θ can be used to obtain the response of any orientation of the kernels in the filterbank to an image $f(\mathbf{x})$ for a rotation around the position $\mathbf{0}$ as

$$\mathbf{R}^N \{f_\theta\}(\mathbf{0}) = A_\theta \mathbf{R}^N \{f\}(\mathbf{0}), \quad (4)$$

where f_θ denotes the rotation of f as $f_\theta(\mathbf{x}) = f(R_\theta \mathbf{x})$ and R_θ is a 2D rotation matrix parameterized by θ .

Multiscale analysis proceeds by partitioning the Fourier domain into several dyadic bands of decreasing sizes. This is done using Simoncelli's isotropic multiresolution framework (Simoncelli and Freeman, 1995) and it controls the spatial support (i.e., scale) of the Riesz filters. Class-specific texture signatures for several tissue classes can be obtained by determining a weighting scheme for the Riesz filterbanks for every scale. This is illustrated in Fig. 5 where a Riesz filterbank of order 5 is used to generate a signature for the micronodule class using a suitable weighting scheme for the individual Riesz filters. More formally, we are looking for an optimal texture signature Γ_c^N of the class c from a linear combination of the individual Riesz filters in the filterbank as:

$$\Gamma_c^N = \mathbf{w}^T \mathbf{R}^N = w_1 \mathcal{R}^{(0,N)} + w_2 \mathcal{R}^{(1,N-1)} + \dots + w_{N+1} \mathcal{R}^{(N,0)}, \quad (5)$$

where \mathbf{w} contains the weights of the respective Riesz filters. For our specific case, a multiscale texture signature is obtained by extending Eq. (5) using multiscale Riesz filterbanks (Depeursinge et al., 2012a) as:

$$\Gamma_c^N = w_1 (\mathcal{R}^{(0,N)})_{s_1} + w_2 (\mathcal{R}^{(1,N-1)})_{s_2} + \dots + w_J (\mathcal{R}^{(N,0)})_{s_J}, \quad (6)$$

where s_j , for $j = 1, \dots, J$ is the scale index. The image to be analyzed is mapped onto the Fourier domain. Then, different bands of frequencies from this domain are extracted using Simoncelli's multi-resolution




Fig. 5. The process of building a texture signature for the micronodule class is illustrated here through an appropriate weighing scheme for a Riesz filterbank of order 5.

framework for domain partitioning to obtain different filtered versions of the original image. The weighting scheme described above is determined for each scale such that we obtain texture signatures at different frequency-domain scales for the same image.

2.2.2. Determining optimal texture signatures using a novel one-versus-one approach

The weighing scheme itself may be determined by using an SVM-based one-versus-one classification technique. The energies of the filter responses $E(\mathcal{R}^{(n,N-n)}\{f\}(\mathbf{x}))$ need to be calculated for one class c versus each one from the remaining classes and then fed to an SVM that finds the optimal separation (in the sense of structural risk minimization Vapnik, 1995; Guyon et al., 2002) between c and each of the remaining classes. Each class is thus uniquely characterized against the remaining classes, meaning that this approach produces highly discriminative classification variables. Since there are five classes, the approach leads to $5 \times 4 = 20$ different optimal signatures. However, the optimal separation between a class A and a class B is the same as separation between class B and class A, the number of optimal signatures reduces to 10. The optimal weights $\mathbf{w} = (w_1, \dots, w_{N+1})$ are directly derived from the support vectors of the optimal separation found above as in Depeursinge et al. (2014a).

2.3. Classification

The images are filtered using the Riesz template and steerability is used to obtain the responses of class-wise locally aligned texture signatures Γ_c^N . The average energy of these responses within a patch is used to build a feature space for classification. This feature space is complemented by a histogram of Hounsfield Units (HU) from the spatial domain as it is important to include the so-called DC component of image signals that is neglected in the Fourier domain approach we have presented so far. The average air content of lung tissue in a patch under investigation is also an important characteristic of the tissue class. Therefore, the number of pixels representing air per block is also included as a variable in the final feature vector for a patch. As a result, the feature space for each patch consists of the histogram of HU values, the average air content, and 10 filter responses (one for each optimal one-versus-one Riesz filter).

The features are then used to train an SVM model as part of a supervised learning approach. First, the vector containing the histogram of HU values is combined with the vector containing the response of the healthy-versus-emphysema Riesz filter (the first of the ten filters used) for training instances of the class ‘healthy’ (first of the five tissue classes being investigated here) and all other classes, collectively labeled as ‘all others’. An SVM is used to find the optimal separation between them. The vectors within the feature space of the test patches are similarly

partitioned and combined. Another strong feature of our method is that it allows the analysis of classwise probabilities that are used to make a final classification decision. Using the SVM model just trained, the relative probabilities of a patch belonging to each one of the five classes are determined. This is followed by a similar calculation for the training instances of the remaining four classes in turn against all others. The process is repeated for the remaining nine Riesz filters. The class probabilities are subsequently aggregated and the class having the highest probability for each of the patches is chosen as the predicted class for the patch.

2.4. Parallel computing

Calculating Riesz energies is a computationally expensive process since this has to be done by iterating over individual pixels in each patch of the ILD dataset. The required timescale is too large for our method to be viable on typical workstations. Thanks to the power of dedicated Graphical Processing Units (GPUs), filter responses are calculated in parallel within a workstation using techniques developed in previous work (Vizitiu et al., 2016). Several classification tasks on different sets of training and test images are handled in parallel on separate workstations through the use of Hadoop on a local network of workstations.

3. Results

3.1. Experimental results

MATLAB was used for the calculations described above. Square patches with a length of 33 pixels were extracted from the annotated region in each of the slices of the training and test images. A patch is extracted only if at least 75% of the 33-by-33 square lies inside the annotated region and its center is at least half a patch length away from the respective centers of other extracted patches in both the horizontal and vertical axis directions. The specific patch length was chosen after preliminary investigations of appropriate patch lengths that give good results and have a sufficiently large number of patches for classification. Patch extraction is illustrated in Fig. 6 where several overlapping square patches are identified on a particular slice with a known label. 14,594 patches are obtained in this way from the ILD dataset. Their distribution according to tissue class: healthy (H), emphysema (E), ground glass (G), fibrosis (F), and micronodules (M) is presented in Table 1.

For each of the patches generated, the respective grey level histogram is computed. Since we are dealing with HRCT images, the intensity values are expressed in HU. The histogram used in computing the feature space for a patch has 22 bins and is restricted to the value




Fig. 6. The extraction of several overlapping patches is exemplified here on an annotated slice from the ILD database that was shown in Fig. 2.

Table 1
The classwise distribution of patches extracted from the ILD database is shown here.

Class	Number of patches
Healthy (H)	3011
Emphysema (E)	407
Ground glass (G)	2226
Fibrosis (F)	2962
Micronodules (M)	5988
Total	14,594

range – 1000 (HU value for air) to 650 (typical HU value for bone). A variable for air content in the patch (HU less than or equal to – 1000) is also included. The feature space is completed using energy values computed using the multiscale approach presented in the previous section for $J = 4$ scales. This number of scales was chosen to optimally cover spatial frequencies contained in 33×33 image patches.

After a series of initial investigations (not detailed in this paper) into an appropriate value for the Riesz order N , a value of 5 was chosen. For $N = 5$, 10 texture signatures of length $(N + 1) * J = 24$ each are obtained. To create the final feature vector, the $24 \times 10 = 240$ variables obtained in the previous step are concatenated with the 22 variables from the histogram and one variable for air content.

As a result, a feature vector of length 263 is obtained for each patch. Classification is then carried out on the feature vectors for several values of Radial Basis Function (RBF) kernel SVM cost parameters: gamma, and cost. Gamma controls how far the influence of a single training instance extends in the determination of support vectors; low gamma corresponds to a far-reaching influence. The cost parameter controls the penalty for misclassification; high cost forces the algorithm to explain the input information better, leading to a higher risk of overfitting. A maximum absolute accuracy of 80.31% is achieved for gamma = 1.0e0 and cost = 1.0e – 2. The confusion matrix for this result is shown in Table 2. Using the confusion matrix, the true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN) are identified for each individual tissue class. Subsequently, the per-class specificity, precision, recall, F -score, and accuracy are calculated and shown in Table 3. In particular, the accuracy is calculated using Eq. (7). Finally, Table 4 shows the precision, recall, and F -score for the proposed method and comparable work for which those figures are available.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (7)$$

Table 2
The confusion matrix for the highest classification accuracy observed using our method is shown here.

True label	Predicted label				
	H	E	G	F	M
H	2186	18	21	19	767
E	48	233	7	59	60
G	124	2	1619	205	276
F	20	22	250	2441	229
M	441	9	152	144	5242

Table 3
The per-class specificity, precision, recall, F -score, and accuracy for the experiment yielding maximum observed classification accuracy are shown in the table.

Class	Specificity	Precision	Recall	F -Score	Accuracy
H	0.9454	0.7755	0.7260	0.7499	0.9001
E	0.9964	0.8204	0.5725	0.6744	0.9846
G	0.9652	0.7901	0.7273	0.7574	0.9289
F	0.9633	0.8511	0.8241	0.8373	0.9350
M	0.8452	0.7974	0.8754	0.8346	0.8576

Table 4
The overall precision, recall, and F -score for the proposed method and comparable work are shown below. The exact evaluation scenario might not always exactly be the same, with differences in patch size and training/test splits.

Method	Precision	Recall	F -Score
Ours	0.807	0.745	0.771
Song et al. (2013)	0.807	0.826	0.815
Song et al. (2015)	0.825	0.841	0.833
Depeursinge et al. (2012c)	0.763	0.758	0.760
Depeursinge et al. (2011b)	0.649	0.716	0.667
Li et al. (2014)	0.738	0.770	n/a
Li et al. (2013)	0.702	0.744	n/a

4. Discussion

The results obtained using our method are compared with the results obtained by other authors who used the same ILD dataset but with possibly slight variations in terms of the evaluation methodology. Table 5 displays the accuracies for the different tissue types obtained by our method and others. Fig. 7 allows a comparison of classwise recall, Fig. 8 allows a comparison of classwise precision, and Fig. 9 allows a comparison of classwise F -scores.

It can be seen that the performance of our approach compares and sometimes improves over a variety of current methods (including deep

Table 5
The diagonal of the confusion matrices obtained for our method versus others is shown below.

Method	Class				
	H	E	G	F	M
Ours	0.7260	0.5725	0.7273	0.8241	0.8754
Song et al. (2013)	0.8760	0.8060	0.8270	0.8120	0.8110
Shin et al. (2016)	0.6800	0.9100	0.7000	0.8300	0.7900
Depeursinge et al. (2012c)	0.6728	0.7872	0.7136	0.8273	0.8156
Depeursinge et al. (2012a)	0.8270	0.7270	0.6840	0.8420	0.8350
Foncubierta-Rodríguez et al. (2012)	0.0530	0.7450	0.4960	0.7460	0.519
Depeursinge et al. (2011a)	0.7750	0.7330	0.7230	0.8450	0.8050
Li et al. (2013)	0.7600	0.6700	0.7000	0.7400	0.8400
Depeursinge et al. (2011b)	0.7900	0.6920	0.5930	0.8050	0.7020
Gao et al. (2016)	0.9142	0.8270	0.8151	0.8910	0.8799
Song et al. (2015)	0.8850	0.7960	0.8000	0.8540	0.8720




Fig. 7. Classwise recall obtained with our Riesz-based approach is compared with the approaches in the literature that use the ILD dataset.

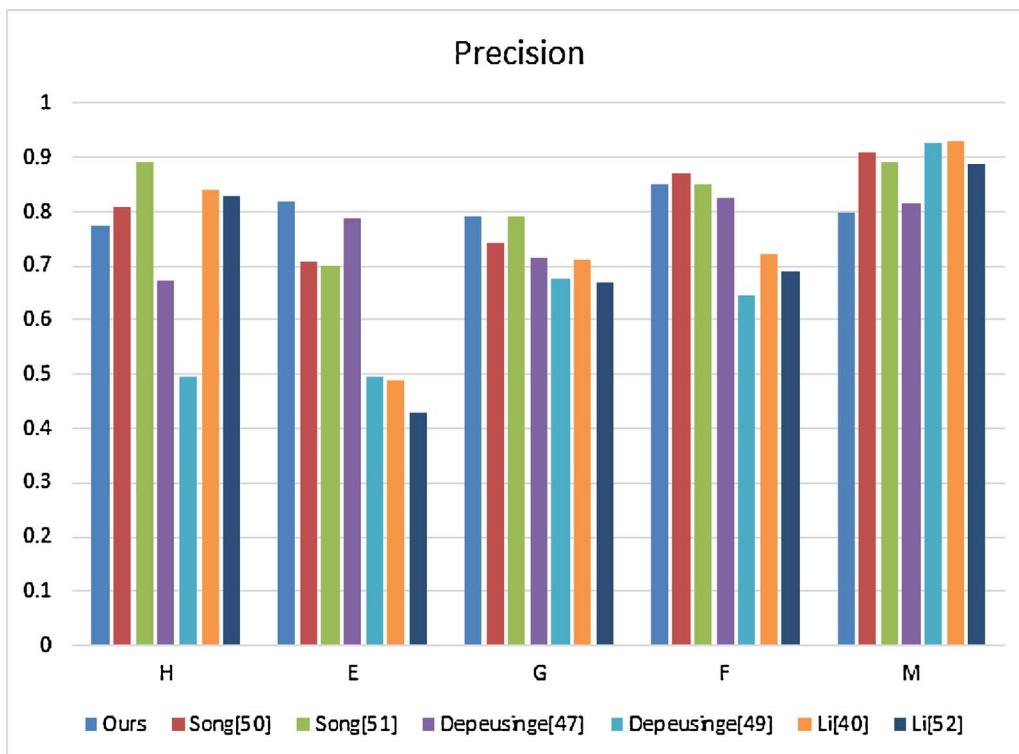


Fig. 8. Classwise precision obtained with our Riesz-based approach is compared with the approaches in the literature that use the ILD dataset.

CNNs and feature-based learning) in the classification of ILDs while using a novel one-versus-one method of generating classification signatures. It is important to note that although the dataset used is the same, the exact validation scheme differs from one method to another according to patch size, selection of patches (percentage contained in the ROI), distribution of the classes, and cross-validation. We aim to reuse the same methodology of our prior work. Our performance for the healthy class is not as good as other approaches and this requires further analysis. It is appropriate to question the existence of latent lung tissue damage in the patients who contributed to the healthy class given the fact that they had been recruited while they were under examination for suspected lung problems and all of them have some regional abnormal texture. The radiologists who annotated the images were

specifically ask to find healthy tissue in each patient if possible and this may have led to some healthy tissue annotations containing other underlying patterns.

In fact, it might be wise to consider whether the entire initial ground-truth annotations are fully clean in the annotated areas. In Aziz et al. (2004), a study of inter-observer disagreement between thoracic radiologists during the diagnosis of diffuse parenchymal lung disease (DPLD) using HRCT images was carried out. The study found that observer agreement was only moderate on the whole. It even advised the use of the expertise of a reference panel for cases diagnosed with low confidence. This is relevant here since annotation of the ILD database was carried out by a consensus of two clinicians and no inter-observer disagreement was measured. Inter-annotator disagreement thus




Fig. 9. Classwise *F*-score obtained with our Riesz-based approach is compared with the approaches in the literature that use the ILD dataset.

presents a non-negligible risk of causing under- or over-estimation of the classification accuracy observed in this paper and various others that used the ILD database.

In an attempt to find further causes of misclassification, a detailed analysis of the patches that were wrongly classified was conducted. Fig. 10 illustrates an instance of a common occurrence of misclassification on an HRCT image slice from the ILD database. The correct class for all the patches extracted from this slice was emphysema (class 2 in internal calculations). Out of 11 patches present, nine were correctly classified, one was classified as fibrosis (class 4, shown in red and labeled accordingly in Fig. 10), and the remaining one was classified as micronodules (class 5, also shown in red and labeled accordingly in Fig. 10). The error seems to occur at the border of the ROI where a significant portion (up to 25%) of the patch contains tissue not

part of the annotated region, most frequently the healthy class. In fact, of the 14,594 patches, 14,550 (or 99.7%) are at the edges. That percentage is large because of the way edge patches are defined in our approach: a patch is at the edge of an ROI when at least one pixel in it lies outside that ROI. Their classification accuracy, standing at 80.27%, is lower than the overall accuracy. The accuracy for the 44 patches fully within the ROI stands at 93.18%. This is an indication that non-annotated tissue in edge patches have a nefarious influence on the classification algorithm, leading to erroneous behavior.

A further explanation for error at the edges of an ROI is the influence of transitions at the lung boundaries on the response of wavelets at large scales (Depeusinge et al., 2015b). Fig. 11 shows the respective distribution of probabilities for belonging to a certain class versus all others for the two misclassified patches: the left one misclassified as




Fig. 10. Typical misclassifications of HRCT regions from the ILD database are shown here. They occur often at the boundary of the annotated region. Emphysema was wrongly classified as fibrosis and micronodules.




Fig. 11. Bar charts of individual probabilities for a patch belonging to a certain class versus all others is shown here for a patch misclassified as fibrosis (left) and one misclassified as micronodules (right). The correct class was emphysema.

fibrosis and the right one misclassified as micronodules. The correct class was emphysema and it was a strong candidate (appearing in the top two most likely) for the final classification decision but was erroneous.

Taking into account the fact that the presence of more than one tissue class in a patch creates noise in the learned signatures and that ground-truth annotation used to assess classification accuracy may partly be subjective. We subsequently explored the possibility of predicting a class only if the individual probability of belonging to that class versus all others (vertical axis in Fig. 11) is high. Fig. 12 shows the effect on absolute accuracy of choosing only the best candidate classes displaying a high value for the probability. As can be seen, the classification accuracy improves significantly for each class and overall. However, the number of patches conforming to the probability threshold and classified in this way is reduced accordingly as is shown in Fig. 13. An analysis of probability distributions for each candidate class for the remaining patches may be used to suggest the two most likely tissue types along with relevant probability instead of proposing a single tissue type. We could also increase the amount of a patch required to be in the region of interest when he patches are created but this would make results uncomparable to our own prior work. The above analysis of probability distributions before classification is a way to quantify the likelihood of the presence of tissue types and may help clinicians become aware that a certain outcome might not be solid. This

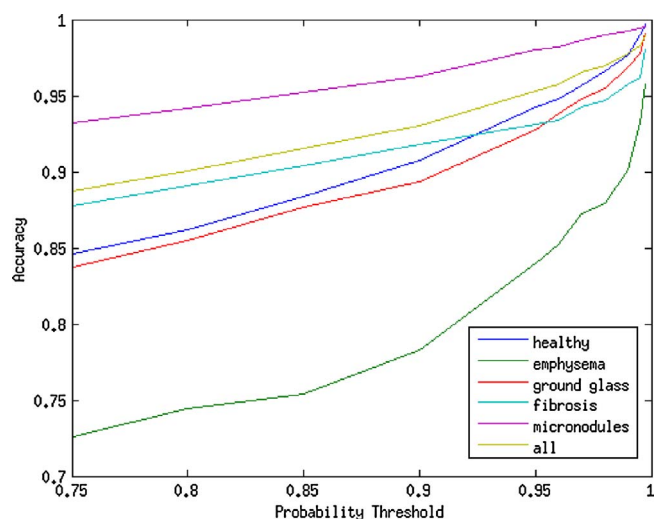


Fig. 12. The classification accuracy overall and for each individual class is seen to increase as the probability threshold utilized to enact classification is raised.




Fig. 13. As the probability threshold used to enact classification is raised, the percentage of patches classified in this way drops.

could substitute the role of a reference panel (recommended in Aziz et al., 2004) providing reproducible advice regarding possible tissue types and diagnoses.

Finally, Tables 2 and 3 show that there is some room for improvement in the classification of the emphysema class, especially regarding the recall obtained. Indeed, only 407 patches with identified emphysema are encountered in the ILD database while the next least frequent disease class is ground glass with 2226 patches. That is a large disparity and we would argue that our Riesz filter for emphysema is consequently less well trained as compared to the other classes due to a much lower number of patches used for training. In addition, emphysema has very large intra-class variations and would require learning several steerable models or signatures per class. We contend that the use of more patches belonging to that class for training and the use of more than one signature for emphysema in subsequent work would significantly improve the overall classification accuracy.

In more future work, we plan to address the outstanding challenges regarding classification accuracy by taking into account an important aspect that has been neglected in most approaches seen so far: the use of a-priori information regarding the physical location of a patch within a slice and within an HRCT volume (Zrimec and Wong, 2007; Depeursinge et al., 2015a; Gao et al., 2016). From Fig. 10, it is clear that patches lying adjacent to each other more than likely belong to the same tissue class. The exploitation of this property is expected to help

us with incorrect classifications at the border of annotations and improve overall classification accuracy.

5. Conclusion

In this paper, a method to characterize several tissue types in chest CTs of patients with ILDs using Riesz-wavelet-based texture signatures is presented. A maximum classification accuracy of 80.31% was observed for texture classification of 5 tissue classes in the ILD database. This is better than previous own work with the same evaluation methodology. The approaches with slightly higher performance may have slightly different patch selection and evaluation techniques, which makes a direct comparison hard.

Accuracy improves if classification is only carried out when we have high confidence in the outcome. For patches with lower confidence, we can propose two of the most likely tissue classes along with the corresponding probability for the clinicians' consideration much in the same way as a reference panel of experts could do. In addition, existing methods classify image patches independently from each other, neglecting the fact that adjacent patches often come from the same annotated area within an HRCT slice. Future work aims at leveraging this in order to correct for the misclassification of patches at the border of annotated areas.

It is important to note that differences in the visual appearance of the classes encountered in this paper were often only visible as subtle changes to normal lung tissue texture. The methods presented have promising potential for future applications of other organ textures.

Acknowledgements

This work was supported by the Swiss National Science Foundation (under grant PZ00P2_154891) and partially funded by the European Commission under the CP-Collaborative Project funding scheme through the FP7-ICT MD-PAEDIGREE project (ID: 600932).

References

- Ablowitz, M.J., Segur, H., 1981. *Solitons and the Inverse Scattering Transform*, vol. 4 SIAM.
- Ablowitz, M.J., Kaup, D.J., Newell, A.C., 1974. The inverse scattering transform-Fourier analysis for nonlinear problems. *Stud. Appl. Math.* 53 (4), 249–315.
- Anthimopoulos, M., Christodoulidis, S., Ebner, L., Christe, A., Mougiakakou, S., 2016. Lung pattern classification for interstitial lung diseases using a deep convolutional neural network. *IEEE Trans. Med. Imaging* 35 (5), 1207–1216.
- Aziz, Z.A., Wells, A.U., Hansell, D.M., Bain, G.A., Copley, S.J., Desai, S.R., Ellis, S.M., Gleeson, F.V., Grubnic, S., Nicholson, A.G., Padley, S.P., Pointon, K.S., Reynolds, J.H., Robertson, R.J., Rubens, M.B., 2004. HRCT diagnosis of diffuse parenchymal lung disease: inter-observer variation. *Thorax* 59 (6), 506–511.
- Barr, R., Parr, D., Vogel-Claussen, J., 2016. *Imaging*. European Respiratory Society.
- Blakemore, C., Campbell, F.W., 1969. On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *J. Physiol.* 203 (1), 237–260.
- Breiman, L., 2001. Random forests. *Mach. Learn.* 45 (1), 5–32.
- Brigham, E.O., 1988. *The Fast Fourier Transform and its Applications*.
- Candès, E.J., Donoho, D.L., 2000. Curvelets – a surprisingly effective nonadaptive representation for objects with edges. *Curves and Surface Fitting*. Vanderbilt University Press, Nashville, pp. 105–120.
- Castellano, G., Bonilha, L., Li, L.M., Cendes, F., 2004. Texture analysis of medical images. *Clin. Radiol.* 59 (12), 1061–1069.
- Condessa, F., Bioucas-Dias, J., Kovačević, J., 2017. Performance measures for classification systems with rejection. *Pattern Recognit.* 63, 437–450.
- Cortes, C., Vapnik, V., 1995. Support-vector networks. *Mach. Learn.* 20 (3), 273–297.
- Cula, O.G., Dana, K.J., 2004. 3D texture recognition using bidirectional feature histograms. *Int. J. Comput. Vis.* 59 (1), 33–60.
- Dalal, N., Triggs, B., 2005. Histograms of oriented gradients for human detection. *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1 of CVPR '05. IEEE Computer Society, Washington, DC, USA, pp. 886–893.
- Depeursinge, A., Sage, D., Hidki, A., Platon, A., Poletti, P.-A., Unser, M., Müller, H., 2007. Lung tissue classification using Wavelet frames. 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2007, EMBS 2007. IEEE Computer Society, Lyon, France, pp. 6259–6262.
- Depeursinge, A., Iavindrasana, J., Cohen, G., Platon, A., Poletti, P.-A., Müller, H., 2008. Lung tissue classification in HRCT data integrating the clinical context. In: 21th IEEE Symposium on Computer-Based Medical Systems (CBMS). Jyväskylä, Finland. pp. 542–547.
- Depeursinge, A., Foncubierta-Rodríguez, A., Van De Ville, D., Müller, H., 2011a. Lung texture classification using locally-oriented Riesz components. In: Fichtinger, G., Martel, A., Peters, T. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2011*, vol. 6893 of Lecture Notes in Computer Science, Springer Berlin/Heidelberg, pp. 231–238.
- Depeursinge, A., Zrimec, T., Busayarat, S., Müller, H., 2011b. 3D lung image retrieval using localized features. In: *Medical Imaging 2011: Computer-Aided Diagnosis*, vol. 7963. SPIE. pp. 79632E.
- Depeursinge, A., Foncubierta-Rodríguez, A., Van De Ville, D., Müller, H., 2012a. Multiscale lung texture signature learning using the Riesz transform. *Medical Image Computing and Computer-Assisted Intervention MICCAI 2012*, vol. 7512 of Lecture Notes in Computer Science, Springer Berlin/Heidelberg 517–524.
- Depeursinge, A., Vargas, A., Platon, A., Geissbuhler, A., Poletti, P.-A., Müller, H., 2012b. Building a reference multimedia database for interstitial lung diseases. *Comput. Med. Imaging Graph.* 36 (3), 227–238.
- Depeursinge, A., Van De Ville, D., Platon, A., Geissbuhler, A., Poletti, P.-A., Müller, H., 2012c. Near-affine-invariant texture learning for lung tissue analysis using isotropic wavelet frames. *IEEE Trans. Inf. Technol. BioMed.* 16 (4), 665–675.
- Depeursinge, A., Foncubierta-Rodríguez, A., Müller, H., Van De Ville, D., 2013. Rotation-covariant visual concept detection using steerable Riesz wavelets and bags of visual words. In: *SPIE Wavelets and Sparsity XV*, vol. 8858. SPIE. pp. 885816–885816-11.
- Depeursinge, A., Foncubierta-Rodríguez, A., Van De Ville, D., Müller, H., 2014a. Three-dimensional solid texture analysis and retrieval in biomedical imaging: review and opportunities. *Med. Image Anal.* 18 (1), 176–196.
- Depeursinge, A., Foncubierta-Rodríguez, A., Van De Ville, D., Müller, H., 2014b. Rotation-covariant texture learning using steerable Riesz wavelets. *IEEE Trans. Image Process.* 23 (2), 898–908.
- Depeursinge, A., Chin, A.C., Leung, A.N., Terrone, D., Bristow, M., Rosen, G., Rubin, D.L., 2015a. Automated classification of usual interstitial pneumonia using regional volumetric texture analysis in high-resolution CT. *Invest. Radiol.* 50 (4), 261–267.
- Depeursinge, A., Pad, P., Chin, A.C., Leung, A.N., Rubin, D.L., Müller, H., Unser, M., 2015b. Optimized steerable wavelets for texture analysis of lung tissue in 3-D CT: classification of usual interstitial pneumonia. *IEEE 12th International Symposium on Biomedical Imaging, ISBI 2015* 403–406.
- Do, M.N., Vetterli, M., 2002. Rotation invariant texture characterization and retrieval using steerable wavelet-domain hidden Markov models. *IEEE Trans. Multimed.* 4 (4), 517–527.
- Donoho, D.L., 2001. Ridge functions and orthonormal ridgelets. *J. Approx. Theory* 111 (2), 143–179.
- European Lung White Book, <http://www.erswhitebook.org/chapters/the-burden-of-lung-disease/> (accessed: 2016-07-06).
- Foncubierta-Rodríguez, A., Depeursinge, A., Müller, H., 2012. Using multiscale visual words for lung texture classification and retrieval. In: Greenspan, H., Müller, H., Syeda Mahmood, T. (Eds.), *Medical Content-based Retrieval for Clinical Decision Support*, vol. 7075 of MCBR-CDS 2011, Lecture Notes in Computer Sciences (LNCS), pp. 69–79.
- Freeman, W.T., Adelson, E.H., 1991. The design and use of steerable filters. *IEEE Trans. Pattern Anal. Mach. Intell.* 13 (9), 891–906.
- Gao, M., Bagci, U., Lu, L., Wu, A., Buty, M., Shin, H.-C., Roth, H., Papadakis, G.Z., Depeursinge, A., Summers, R.M., et al., 2016. Holistic classification of CT attenuation patterns for interstitial lung diseases via deep convolutional neural networks. *Comput. Methods Biomech. Biomed. Eng. Imaging Vis.* 1–6.
- Guyon, I., Weston, J., Barnhill, S., Vapnik, V., 2002. Gene selection for cancer classification using support vector machines. *Mach. Learn.* 46 (1–3), 389–422.
- Haralick, R.M., Shanmugam, K., Dinstein, I., 1973. Textural features for image classification. *IEEE Trans. Syst. Man Cybern.* 3 (6), 610–621.
- Held, S., Storath, M., Massopust, P., Forster, B., 2010. Steerable wavelet frames based on the Riesz transform. *IEEE Trans. Image Process.* 19 (3), 653–667.
- Hyvärinen, A., Hoyer, P., Inki, M., 2001. Topographic independent component analysis. *Neural Comput.* 13 (7), 1527–1558.
- Jeng-Shyang, P., Jing-Wein, W., 1999. Texture segmentation using separable and non-separable wavelet frames. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* 82 (8), 1463–1474.
- Kreuter, M., Herth, F.J., Wacker, M., Leidl, R., Hellmann, A., Pfeifer, M., Behr, J., Witt, S., Kauschka, D., Mall, M., et al., 2015. Exploring clinical and epidemiological characteristics of interstitial lung diseases: rationale, aims, and design of a nationwide prospective registry – the exciting-ILD registry. *BioMed Res. Int.*
- Lazebnik, S., Schmid, C., Ponce, J., 2005. A sparse texture representation using local affine regions. *IEEE Trans. Pattern Anal. Mach. Intell.* 27 (8), 1265–1278.
- LeCun, Y., Huang, F.J., Bottou, L., 2004. Learning methods for generic object recognition with invariance to pose and lighting. In: *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004 CVPR*, vol. 2. IEEE. pp. II-97.
- LeCun, Y., Kavukcuoglu, K., Farabet, C., 2010. Convolutional networks and applications in vision. *IEEE International Symposium on Circuits and Systems (ISCAS)* 253–256.
- Leung, T., Malik, J., 2001. Representing and recognizing the visual appearance of materials using three-dimensional textons. *Int. J. Comput. Vis.* 43 (1), 29–44.
- Li, Q., Cai, W., Feng, D.D., 2013. Lung image patch classification with automatic feature learning. In: 2013 35th Annual International Conference of the Engineering in Medicine and Biology Society (EMBC). IEEE. pp. 6079–6082.
- Li, Q., Cai, W., Wang, X., Zhou, Y., Feng, D.D., Chen, M., 2014. Medical image classification with convolutional neural network. In: 2014 13th International Conference on Control Automation Robotics & Vision (ICARCV). IEEE. pp. 844–848.
- Lowe, D.G., 2004. Distinctive image features from scale-invariant keypoints. *Int. J.*

- Comput. Vis. 60 (2), 91–110.
- Mallat, S.G., 1989. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (7), 674–693.
- Meyer, F.G., Coifman, R.R., 1997. Brushlets: a tool for directional image analysis and image compression. *Appl. Comput. Harm. Anal.* 4 (2), 147–187.
- Ojala, T., Pietikäinen, M., Mäenpää, T., 2002. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (7), 971–987.
- Po, D.D.Y., Do, M.N., 2006. Directional multiscale modeling of images using the contourlet transform. *IEEE Trans. Image Process.* 15 (6), 1610–1620.
- Porter, R., Canagarajah, N., 1997. Robust rotation-invariant texture classification: wavelet, Gabor filter and GMRF based schemes. *IEE Proc. Vis. Image Signal Process.* 144 (3), 180–188.
- Randen, T., Husoy, J.H., 1999. Filtering for texture classification: a comparative study. *IEEE Trans. Pattern Anal. Mach. Intell.* 21 (4), 291–310.
- Ruck, D., Rogers, S., Kabrisky, M., Oxley, M., Suter, B., 1990. The multilayer perceptron as an approximation to a Bayes optimal discriminant function. *IEEE Trans. Neural Netw.* 1 (4), 296–298. <http://dx.doi.org/10.1109/72.80266>.
- Schmid, C., 2001. Constructing models for content-based image retrieval. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2 of CVPR 39–45.
- Shin, H.-C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R.M., 2016. Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Trans. Med. Imaging* 35 (5), 1285–1298.
- Simoncelli, E.P., Farid, H., 1996. Steerable wedge filters for local orientation analysis. *IEEE Trans. Image Process.* 5 (9), 1377–1382.
- Simoncelli, E.P., Freeman, W.T., 1995. The steerable pyramid: a flexible architecture for multi-scale derivative computation. *Proceedings of International Conference on Image Processing*, 1995, vol. 3 444–447.
- Song, Y., Cai, W., Zhou, Y., Feng, D.D., 2013. Feature-based image patch approximation for lung tissue classification. *IEEE Trans. Med. Imaging* 32 (4), 797–808.
- Song, Y., Cai, W., Huang, H., Zhou, Y., Wang, Y., Feng, D.D., 2015. Locality-constrained subcluster representation ensemble for lung image classification. *Med. Image Anal.* 22 (1), 102–113.
- ter Haar Romeny, B.M., 2010. Multi-scale and multi-orientation medical image analysis. In: *Biomedical Image Processing*. Springer, pp. 177–196.
- Unser, M., Van De Ville, D., 2010. Wavelet steerability and the higher-order Riesz transform. *IEEE Trans. Image Process.* 19 (3), 636–652.
- Unser, M., Chenouard, N., Van De Ville, D., 2011. Steerable pyramids and tight wavelet frames in $L_2(\mathbb{R}^d)$. *IEEE Trans. Image Process.* 20 (10), 2705–2721.
- Van De Ville, D., Blu, T., Unser, M., 2005. Isotropic polyharmonic B-splines: scaling functions and wavelets. *IEEE Trans. Image Process.* 14 (11), 1798–1813.
- Vapnik, V.N., 1995. *The Nature of Statistical Learning Theory*. Springer, New York.
- Vizitiu, A., Itu, L., Joyseeree, R., Depeursinge, A., Müller, H., Suci, C., 2016. GPU-accelerated texture analysis using steerable Riesz wavelets. *24th IEEE Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP2016*.
- Xu, D.-H., Kurani, A.S., Furst, J., Raicu, D.S., 2004. Run-length encoding for volumetric texture. In: *The 4th IASTED International Conference on Visualization, Imaging, and Image Processing – VIIP 2004*. Marbella, Spain.
- Xu, Y., Yang, X., Ling, H., Ji, H., 2010. A new texture descriptor using multifractal analysis in multi-orientation wavelet pyramid. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 161–168.
- Yan, H., Jiang, Y., Zheng, J., Peng, C., Li, Q., 2006. A multilayer perceptron-based medical decision support system for heart disease diagnosis. *Expert Syst. Appl.* 30 (2), 272–281.
- Zarzar, M., Razak, E., Htike, Z.Z., Yusof, F., 2015. Classification of immunosignature using random forests for cancer diagnosis. *Adv. Sci. Lett.* 21 (11), 3449–3452.
- Zrimec, T., Wong, J.S.J., 2007. Improving computer aided disease detection using knowledge of disease appearance. In: *MEDINFO 2007. Proceedings of the 12th World Congress on Health (Medical) Informatics*, vol. 129. IOS Press, pp. 1324–1328.