

Near-Affine-Invariant Texture Learning for Lung Tissue Analysis Using Isotropic Wavelet Frames

Adrien Depeursinge, Dimitri Van de Ville, Alexandra Platon, Antoine Geissbuhler,
Pierre-Alexandre Poletti, and Henning Müller

Abstract—We propose near-affine-invariant texture descriptors derived from isotropic wavelet frames for the characterization of lung tissue patterns in high-resolution computed tomography (HRCT) imaging. Affine invariance is desirable to enable learning of nondeterministic textures without *a priori* localizations, orientations, or sizes. When combined with complementary gray-level histograms, the proposed method allows a global classification accuracy of 76.9% with balanced precision among five classes of lung tissue using a leave-one-patient-out cross validation, in accordance with clinical practice.

Index Terms—High-resolution computed tomography (HRCT), interstitial lung diseases (ILDs), isotropic wavelet frames, lung tissue analysis, texture analysis.

I. INTRODUCTION

THE appearance and quantification of the types of lung tissue patterns in high-resolution computed tomography (HRCT) are very informative for establishing the differential diagnosis of interstitial lung diseases (ILDs). ILDs regroup more than 150 disorders of the lung parenchyma leading to breathing dysfunction [1]. The diagnosis of these pathologies is established based on the complete history of the patient, a physical examination, laboratory tests, pulmonary function testing, as well as visual findings on radiological recordings. The primary imaging used is the chest radiograph because of its low cost and weak radiation exposure. It also provides a quick overview of the entire chest. However, chest radiographs are normal in

more than 10% of the patients with some forms of ILD and can provide a confident diagnosis in only 23% of the cases with lung diseases in general [2]. When the synthesis of this information arouses suspicions toward an ILD, HRCT imaging of the chest is often required to acquire an accurate visual assessment of the lung tissue as the 3-D form of HRCT data avoids superposition of organs and provides precise representations of the pattern and distribution of the lung tissue with a submillimetric resolution. It quickly became the gold standard imaging protocol for the diagnosis of diffuse pulmonary parenchymal diseases.

Interpreting HRCT images of the chest represents a challenge even for trained radiologists and lung specialists [1], [2]. The 3-D form requires significant reading time, effort, and experience for a correct interpretation [3]. In emergency radiology, for instance, radiologists have recourse to a large diversity of imaging modalities applied to different organs. They have to provide a first radiological report with preliminary diagnosis rapidly without being specialized in a single domain. This may result in errors by omission or confusion of diverse pathologic lung tissue types [4].

Owing to the intrinsic complexity of the interpretation of HRCT image series, a real-time image-based computerized assistance appears useful for radiologists as a second opinion. The taxonomy used by radiologists to interpret patterns in HRCT images often relates to texture properties, which suggest that texture analysis is relevant to the characterization of an ILD that is typically diffuse.

A. Lung Texture Analysis in the Literature: State of the Art

A wide range of features have been proposed for characterizing various lung tissue patterns associated with chronic obstructive pulmonary diseases (COPDs) and ILDs [5]:

- 1) gray-level histograms (GLH) [3], [6]–[21];
- 2) mathematical morphology and shape [3], [7], [8], [11];
- 3) gray-level cooccurrence matrices (GLCM) [6]–[12], [14], [15];
- 4) run length (RLE) [6], [9], [11], [12], [14], [15];
- 5) filter banks and wavelets [7], [8], [13], [16]–[24];
- 6) others [such as fractals and local binary patterns (LBPs)] [12]–[15], [25].

The heterogeneous feature group composed of GLH, RLE, and GLCM was used in most of the studies [6]–[11]. A more comprehensive review of the techniques used for lung tissue categorization in the literature along with qualitative evaluations can be found in [5]. Although being able to well describe the lung tissue patterns, the performance of features derived from GLCM and RLE strongly depends on the underlying parameters

Manuscript received July 7, 2011; revised November 6, 2011; March 12, 2012; accepted April 28, 2012. Date of publication May 11, 2012; date of current version July 5, 2012. This work was supported by the Swiss National Science Foundation under Grant 205321–130046 and Grant PP00P2–123438, the Center for Biomedical Imaging, and the European Union Seventh Framework Program in the context of the Khresmoi Project FP7–257528.

A. Depeursinge is with the MedGIFT Group, Business Information Systems, University of Applied Sciences Western Switzerland, Sierre 3960, Switzerland, and also with the University Hospitals of Geneva, Geneva 1211, Switzerland (e-mail: adrien.depeursinge@hevs.ch).

D. Van de Ville is with the Department of Radiology and Medical Informatics, University of Geneva, Geneva 1211, Switzerland, and also with Ecole Polytechnique Fédérale de Lausanne, Lausanne CH-1015, Switzerland (e-mail: dimitri.vandeville@epfl.ch).

A. Platon, A. Geissbuhler, and P.-A. Poletti are with the Department of Radiology and Medical Informatics, University Hospitals of Geneva, Geneva 1211, Switzerland (e-mail: alexandra.platon@hcuge.ch; antoine.geissbuhler@hcuge.ch; pierre-alexandre.poletti@hcuge.ch).

H. Müller is with the MedGIFT Group, Business Information Systems, University of Applied Sciences Western Switzerland, Sierre 3960, Switzerland (e-mail: henning.mueller@hevs.ch).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TITB.2012.2198829

(i.e., scales and directions). Most of the texture features used in the literature in lung texture analysis are modeling similar information: the spatial periodicity and scales contained in the images (GLCM, Gabor filters, wavelets, LBPs, and so on). The question is, which one characterizes the patterns best and is the most adaptable to the needs of lung tissue analysis in HRCT imaging.

Although not being the most common in the state of the art of computer-aided diagnosis (CAD) in HRCT imaging of the chest, filtering techniques have several desirable properties. First, they are providing continuous responses to transient patterns in images. This is not the case for GLCMs that are based on the sum of occurrences of pixel pairs. GLCMs are as a consequence not able to characterize the similarity between series of contiguous pixels, which are often carrying important information in medical images. Filtering allows us to seek for specific features in the images (i.e., edge or ridge detection) by modeling the shape of the filters either in the spatial or in the frequency domain. Wavelet transforms (WTs), which are implemented as filter banks, have the desirable property of being multiscale and thus allow covering the frequency domain (scale covariance). Filtering techniques and translation-invariant wavelets offer an overcomplete feature set able to fit most of the texture functions under the condition that they efficiently derive features from the coefficients, being perfectly complementary to the measures of density using GLH.

The specific texture signatures of the lung tissue patterns can hardly be described by deterministic methods as intraclass variations are very high due to the influence of factors such as the age of the patient, smoking history, and extend of the disease. Highly flexible texture modeling is required to catch subtle texture signatures of a given lung tissue pattern. In particular, invariance of the texture descriptors to affine properties (i.e., translation, rotation, and scale) is desirable to obtain a system that is able to impartially learn any texture appearance independently of prevailing localizations, orientations, or sizes.

The main research contribution of this paper is the development of a near-affine-invariant set of texture features (translation- and rotation-invariant as well as scale covariant) based on the WT. We combine the isotropic properties of polyharmonic B -spline wavelets [26], the scale-covariant properties of the quincunx subsampling, and the translation invariance of redundant frame transforms [27], [28]. We investigate the ability of the proposed wavelet-based texture features to discriminate among the classes of healthy and pathological lung tissue types in HRCT images. GLHs model a complementary information corresponding to the density of the structures in CT imaging and are thus used along with WTs.

II. METHODS

Our approach to obtain near-affine-invariant descriptions and to learn the lung tissue textures is described in this section. The importance of the gray-level values for the characterization of lung parenchyma densities is highlighted in Section II-A. Complementary texture descriptors based on specific designs of WTs are detailed in Section II-B.

A. GLHs

Grayscale values in medical images contain valuable information for the characterization of objects and textures and are complementary to texture features. In HRCT imaging, scanners deliver 12-bit digital imaging and communications in medicine images with voxel values in Hounsfield units (HU) in the range $[-1024, 3071]$. These values correspond univoquely to densities of the anatomic organs and thus allow the identification of lung tissue components [29]. To extract this information, GLHs with N_{bins} bins are built. The value of each bin is directly used as visual feature. The air value given by the number of pixels with values ≤ -1000 HU is computed as additional feature [30]. Preliminary results show that using the values of the bins directly allows better categorization of the lung tissue patterns compared to statistical measures of the distributions (i.e., mean, variance, etc.). Features from histograms showed to be significantly different for six lung tissue patterns from 38 patients using a Mann–Whitney U test in [31].

B. WTs for Lung Tissue Analysis

This section details custom designs of redundant WTs that can provide near-affine-invariant texture features. WTs have been applied to various domains such as signal compression, denoising, pattern recognition, and others. They have received attention in biomedical applications [32]–[34]. WTs are able to detect transients in signals and contain all the information to perfectly reconstruct the signal. They allow signal approximation with the error decreasing inversely proportional to the number of wavelet coefficients used, which is suitable to obtain compact feature representations for either compression or classification purposes.

1) *Translation, Rotation Invariance, and Scale Covariance:* The WT has several shortcomings for texture analysis in the classical form that are partly due to the sampling of digital images. It is shown how the standard WT can be modified slightly to obtain near-affine-invariant texture features.

a) *Translation invariance: wavelet bases versus redundant wavelet frames:* The WT exists in two flavors: bases and frames. Wavelet bases provide a one-to-one decomposition (that may be orthogonal or not), making them ideally suited for image compression. They have a fast implementation but their drawback is the lack of translation invariance. In the standard dyadic discrete WT (basis), multiple scales are obtained by downsampling the image I by a factor of 2 at each iteration j , which implements the pyramidal image decomposition. Although bases provide translation-invariant representations in the continuous domain, sampling the translation factor b over the uniform grid of digital images removes this property. The image I may be translated by a continuous factor τ , which is not *a priori* equal to the grid interval $a^j b_0$ at scale a^j (see Fig. 1 for an example in 1-D). The coefficients obtained with a wavelet ψ_j of I are different from those obtained with an image I_τ translated by $\tau \neq a^j b_0$

$$W_{\psi_j}^I(a^j, b) \neq W_{\psi_j}^{I_\tau}(a^j, b - \tau). \quad (1)$$

To obtain translation-invariant features, one solution is to keep the original resolution of I and to oversample ψ so that the minimum shift τ_{min} is equal to $a^j b_0$, and all values of τ are




Fig. 1. WT is not shift invariant when sampled on a uniform sampling grid because the grid step $a^j b_0$ may differ from a continuous translation τ of the signal $s(x)$ [35].

multiples of $a^j b_0$. The discrete wavelet frame (DWF) decomposition yields highly redundant image representations offering more flexibility for image analysis and performing well for texture analysis. It was first proposed in 1995 [27] and has been widely used [36]–[40], also with biomedical images [41]. When compared to the WT, the translation invariance of DWF tends to decrease the variability of the estimated texture features, improving classification performance. The use of features based on the DWF to classify lung textures is described in Section II-B3 and evaluated in Section III-B.

b) Scale covariance and rotation invariance: isotropic polyharmonic B-spline wavelets and the quincunx lattice: The combination of isotropic polyharmonic B-spline wavelets along with quincunx subsampling offers several desirable properties for scale invariance and rotation invariance [26]. Inspired by Marr’s model for the mammalian early visual system [42], we propose the use of zero-crossing detection of the Laplacian-of-Gaussian filtered image [43] for the analysis of lung tissue patterns in HRCT. This is supported by the following hypotheses.

- 1) The orientation of the detector has to be aligned with the locally dominant orientation, the orientation of the underlying segment of zero crossings.
- 2) This orientation is also the one at which the zero crossing has the maximum slope.
- 3) The lines of zero crossings correspond to the orientation-independent differential operator: the Laplacian $\Delta = \frac{\partial^2}{\partial x_1^2} + \frac{\partial^2}{\partial x_2^2}$.

In other words, ψ_{Δ_G} is able to “face” image edges in any direction as being isotropic. It also offers an appropriate tradeoff between spatial localization and bandwidth, which allows to locally examine edges present in a portion of the spectrum of the image. ψ_{Δ_G} can thus be the ideal filter for texture analysis and the question is how to best implement a multiscale version of ψ_{Δ_G} on the Cartesian grid imposed by digital images.

A new family of wavelets has been proposed for an isotropic localization operator that makes the polyharmonic B-spline scaling function converge to a Gaussian as their order increases [26]. These isotropic polyharmonic B-spline wavelets are derived from elementary m -harmonic cardinal B-splines, where the second moment is well defined. It was demonstrated that these wavelets behave as a γ -th-order Laplacian Δ for low frequencies. The isotropic polyharmonic B-spline β_γ is a candidate to approximate Gaussian-like basis functions as it con-

verges to the Gaussian when the order increases

$$\beta_\gamma(\mathbf{x}) \approx \frac{6}{\pi\gamma} \exp\left(-\frac{6\|\mathbf{x}\|^2}{\gamma}\right) \quad (2)$$

with $\mathbf{x} \in \mathbb{R}^2$, and a standard deviation $\sigma = \sqrt{\gamma/12}$. The convergence is fast: the normalized squared differences between β_γ and its corresponding Gaussian function are inferior to 5% for $\gamma \geq 3$. Thus, isotropic polyharmonic B-spline wavelets implement a multiscale smoothed version of the Laplacian from which the initial scale can be tuned through γ . This wavelet ψ_γ , at the first decomposition level, can be characterized as

$$\psi_\gamma(\mathbf{M}^{-1}\mathbf{x}) = \Delta^{\frac{\gamma}{2}}\{\phi\}(\mathbf{x}) \quad (3)$$

where ϕ is an appropriate smoothing (low pass) function and \mathbf{M} is the subsampling matrix.¹ γ tunes the iterate of the Laplacian operator (comparable to the traditional vanishing moments). Large values of γ reduce the energy of the wavelet coefficients but increase the ringing effect [44]. ψ_γ yields symmetric filters that have a linear phase response, where the delay (shift) is predictable and no phase distortion is introduced. Thus, ψ_γ allows us to preserve the shift invariance provided by the redundant frame transform.

The improved isotropy of the polyharmonic scaling and wavelet functions allows a more flexible choice than for traditional tensor-product B-spline wavelets of the subsampling scheme defined by the matrix \mathbf{M} . The choice of the subsampling scheme defined by \mathbf{M} in itself has an important influence both on scale covariance and rotation invariance. Two main subsampling schemes are found in the literature for the 2-D case: dyadic and quincunx, the former being by far the most widely used. The quincunx subsampling removes one of two samples, while the dyadic one removes three of four samples as shown in Fig. 2. It corresponds to a similarity transform in 2-D.²

With particular wavelets, dyadic subsampling enables separability of the WT allowing computational efficiency because wavelet coefficients within each subband can be obtained by successive 1-D convolutions along the columns and the rows of the image. However, the dyadic lattice has two drawbacks to perform texture analysis. On the one hand, the scale progression is coarse as images are downsampled by a factor of 2 (in each dimension) between two decomposition levels. Relevant information might be padded out when having major energy contained in a narrow subband located between two successive levels of the dyadic transform. Subtle changes in the scale of lung tissue patterns (e.g., micronodules versus bronchovascular structures) can be neglected by the dyadic scale progression [17]. On the other hand, although separability allows fast WTs, this process tends to favor the vertical and horizontal directions, and produces a so-called diagonal wavelet component [46], which does not have a straightforward directional interpretation.

The scale progression of the quincunx subsampling scheme defined by $\mathbf{M}_{\text{quincunx}} = [1 \ 1; 1 \ -1]$ is slower than for dyadic

¹In the frame transform, \mathbf{M} corresponds to the upsampling of the wavelet filter.

²It is important to note that the extension of the quincunx subsampling in more than two dimensions does not correspond to a similarity transform for two-channel filter banks [45].

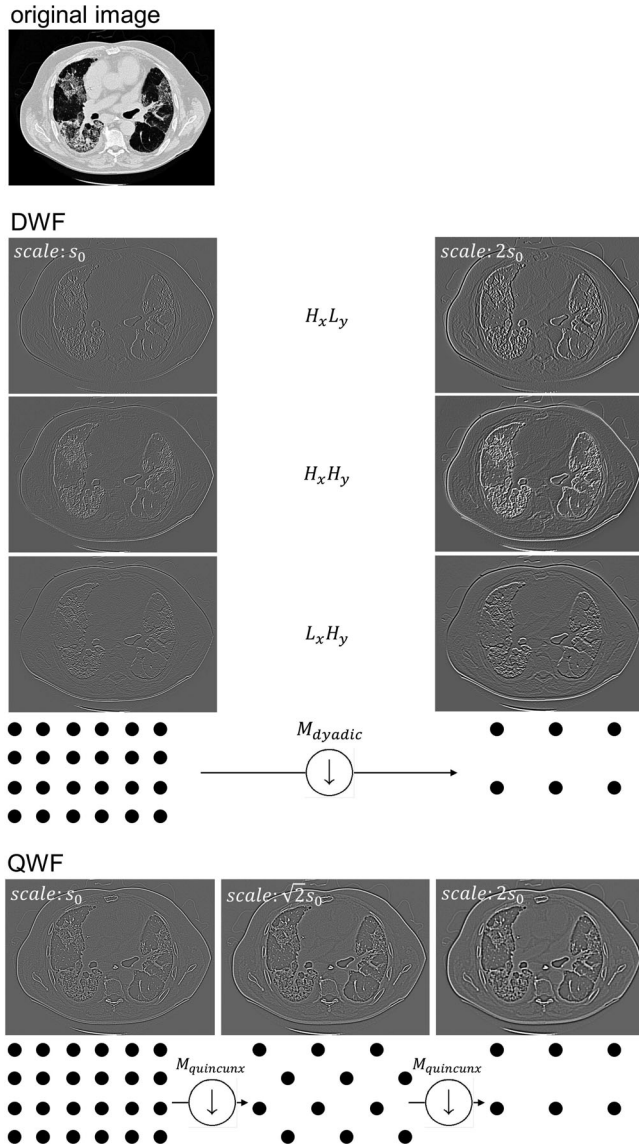


Fig. 2. (Top) Two iterations of the DWF are shown, which corresponds to three iterations of the QWF. At each decomposition level, the filter is upsampled by a factor of 2 and $\sqrt{2}$ for DWF and QWF, respectively, to obtain a multiscale representation. The DWF yields three subbands per scale, whereas QWF only yields one.

subsampling, with an equivalent 1-D downsampling factor of $\sqrt{2}$ instead of 2 for dyadic [47]. Compared to the dyadic separable case, M_{quincunx} preserves isotropy and the wavelet space is spanned by only one wavelet subband S per decomposition level (versus three for separable dyadic) [48]. This leads to a direct and easy interpretation of the subbands; the small number of subbands also breeds small feature spaces (preferable for classification).

The desirable properties of the isotropic polyharmonic B -spline wavelets combined with the quincunx subsampling and the redundant frame transform are summarized in Table I. This combination yields isotropic quincunx wavelet frames (QWFs) that are near-affine invariant: QWF are translation- and rotation-invariant as well as scale covariant. QWFs are evalu-

TABLE I
DESIRABLE PROPERTIES OF THE COMBINATION OF ISOTROPIC POLYHARMONIC B -SPLINE WAVELETS AND THE QUINCUNX LATTICE ALONG WITH A REDUNDANT FRAME TRANSFORM

	translation invariance	scale covariance	rotation invariance
isotropic polyharmonic B -spline wavelets	symmetric filters: no phase distortion	initialization of the scale progression tunable with γ	near isotropic: implement a multiscale smoothed version of Δ
quincunx subsampling	–	finer scale progression compared to dyadic: $\sqrt{2}$ instead of 2	preserves isotropy
redundant frame transform	no image downsampling	multiscale	–

ated and compared to DWF for the classification of the lung tissue patterns in Section III-B.

2) *Features From the Wavelet Coefficients*: For the classification of lung tissue patterns, a subset of features able to characterize the coefficients of the wavelet filterbanks is required. The wavelet frame transform yields a number of wavelet coefficients N_c that is directly proportional to the number of pixels N_{pix} of the input image, the number of iterations J , and the number of subbands per iteration N_{subband} : $N_c = N_{\text{pix}} \times J \times N_{\text{subband}}$. Four iterations of DWF (three subbands per iteration) of an HRCT slice of 512×512 lead to more than three million coefficients. Classification algorithms can hardly learn from such a high-dimensional feature space and compact feature representations are required. In the literature, localized energy measurements were often used to characterize the coefficients in the subbands [27], [36], [49]. As an alternative, Portilla *et al.* showed that Gaussians scale mixtures (GSM) are appropriate to model the distributions of the wavelet coefficients through the successive subbands as well as in a spatial neighborhood within the same subband [50]. This is related to the admissibility condition of wavelets ($\psi(\mathbf{0}) = 0$), which imposes that the global mean of the wavelet coefficients is close to zero. Based on this assumption, the distributions of the wavelet coefficients in each subband S were characterized through the parameters of a simple GSM model of two Gaussians with a fixed mean $\mu_1, \mu_2 = \mu$ and two standard deviations $\sigma_{1,2}$, which are estimated using the expectation-maximization (EM) algorithm [51]. For each subband, μ is initialized with the mean of the wavelet coefficients S_j . $\sigma_{1,2}$ are initialized using the range $r_{S_j} = \max(S_j) - \min(S_j)$ as follows: $\sigma_1 = r_{S_j}, \sigma_2 = r_{S_j}/10$. A high stability of the EM algorithm was also observed with various initializations of σ_1 and σ_2 . The two Gaussians have a weight of 0.5 each. Thirty iterations of the EM algorithm were used, although a convergence of the parameters was reached after ten iterations in most cases.

3) *DWF*: The ability of the DWF combined with GLH to discriminate among five lung tissue patterns is evaluated and compared with QWF. A family of B -splines of various degrees is used as wavelet basis in a separable transform based on the

dyadic lattice for upsampling the filters. B -spline wavelets are compactly supported by smooth piecewise polynomials from which the degree α modulates their scale-space properties [52]. This is convenient for analyzing medical images where the varying scales of the objects require adjustable wavelets. The symmetric B -spline is expressed in the Fourier domain as

$$\hat{\beta}^\alpha(\omega) = \left(\frac{\sin(\omega/2)}{(\omega/2)} \right)^{\alpha+1}. \quad (4)$$

The ℓ_2 -norm of composite diagonal coefficients $C_j(\mathbf{x})$ is computed for each decomposition level j as follows:

$$C_j(\mathbf{x}) = \sqrt{(G_x H_y)_j^2(\mathbf{x}) + (G_y H_x)_j^2(\mathbf{x})} \quad (5)$$

where $(G_x H_y)_j(\mathbf{x})$ and $(G_y H_x)_j(\mathbf{x})$ are the coefficients resulting from the convolution with the high-pass filter on x and with the low-pass filter on y and vice versa. The norm of both is computed because little directionality is contained in lung tissue textures [17].

C. Blockwise Classification of the Lung Parenchyma

This section describes the steps of the blockwise classification with texture (either DWF or QWF) and GLH features. A semiautomatic segmentation of the lung volumes is used to locate the lung parenchyma. Then, the segmented regions are entirely categorized using combined texture and GLH features in overlapping blocks.

1) *Semiautomatic Segmentation of the Lung Volumes*: Segmentation of the lung volumes is a required preliminary step to lung tissue categorization [53]. The result of this step is a binary mask M_{lung} that indicates the regions to be analyzed by the texture analysis routines. Since the geometries of the lungs are subject to large variations among cases, semiautomatic segmentation based on region growing and mathematical morphology is used. Starting from a seed point $p(x, y, z)$ defined by the user, each 26-connected neighbor is added to the region M_{lung} if the summed value of its own neighbors differs less than a given variance defined by the user. The resulting binary mask M_{lung} describes the global lung regions well but contains many holes where the region growing algorithm was stopped by dense regions such as vessels or consolidations of the lung parenchyma. To fill these holes, a closing operation is applied to M_{lung} using a spherical structuring element. Two parameters need to be set: the radius r of the spherical structuring element in millimeters and N_{op} , which defines the number of successive closing operations. Although parameter values $r = 5$ and $N_{\text{op}} = 1$ allowed satisfactory lung segmentation, a separated tuning was required for a few cases to obtain optimal segmentation results [54].

2) *3-D lung tissue categorization*: In order to automatically categorize every pixel of M_{lung} , each 2-D slice is divided into overlapping blocks. Preliminary results using block sizes of $\{8 \times 8; 16 \times 16; 24 \times 24; 32 \times 32; 40 \times 40; 48 \times 48; 56 \times 56; 64 \times 64\}$ showed that blocks of size 32×32 have the best tradeoff between classification performance and localization. $N_{\text{bins}} = 22$ showed to be an appropriate tradeoff between discriminatory properties and dimensionality with a subset of 52 patients in [17]. For each block, 22 bins bin_j of GLH in $[-1050,$




Fig. 3. Construction of the feature vector \mathbf{v} for each block using QWF and GLH.

600] and the number of air pixels pix_{air} are concatenated into a single feature vector \mathbf{v} with GSM parameters of each iteration of either DWF or QWF. An example of \mathbf{v} for eight iterations of QWF is depicted in Fig. 3. Table II summarizes the GLH, DWF, and QWF features. Each feature is normalized using a linear mapping between 0 and 1 from each realization to give an equal weight to each of them. No feature weighting or selection is used for the combination of heterogeneous attributes. A support vector machine (SVM) classifier with a Gaussian kernel learns from the space spanned by \mathbf{v} to find the decision boundaries among five classes of lung tissue types using a *one-versus-one* approach. SVMs have shown to be effective to categorize texture in wavelet feature spaces [39] and in particular lung tissue [21]. The optimal cost of the errors C and the width of the Gaussian kernel σ_K are found with a grid search with $C \in [0, 10, \dots, 100]$ and $\sigma_K \in [10^{-2}, 10^{-1.5}, \dots, 10^2]$ using a fivefold cross validation on the training set. Paired $(C; \sigma_K)$ values that enabled best average classification accuracy among the fivefolds are kept. For each case, the whole lung parenchyma is categorized using a distance between the centers of the blocks equal to 4 pixels, leading to an overlap of 87.5%.

D. Dataset

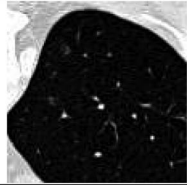
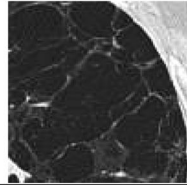

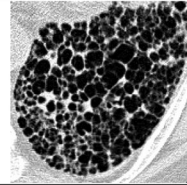
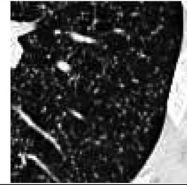
The dataset used is part of a publicly available database of ILD cases [55] containing HRCT images with a slice thickness of 1 mm. 1448 hand-drawn regions of interest (ROIs) were annotated in 2-D HRCT slices in a collaborative fashion by two radiologists with 15 and 20 years of experience at the University Hospitals of Geneva (HUG). The in-plane resolution of the images is 512×512 pixels with an interslice gap of 10 mm. A complete description of the database can be found in [55]. HRCT image series of 85 patients are used to evaluate the performance of the proposed approach. The hand-drawn ROIs are subdivided into 32×32 blocks for evaluating the methods. For blocks to be part of one of the tissue classes, at least 75% of the pixels need to be in the annotated region. The distributions and visual aspects of the five lung tissue classes are detailed in Table III. A total of 17848 blocks were used for the evaluation. The diagnosis of each case was confirmed either by pathology (biopsy and bronchoalveolar washing) or by a laboratory/specific test. In some cases with an early stage of ILD, healthy tissue was annotated in normal parts of the lungs to increase the amount

TABLE II
SUMMARY OF THE FEATURE GROUPS

feature group	\mathbf{v}	$\dim(\mathbf{v})$
GLH	$\mathbf{v} = (bin_0 \ bin_1 \ \dots \ bin_{21} \ pix_{air})$	23
DWF	$\mathbf{v} = (\mu, \sigma_{1,2}(G_1) \ \mu, \sigma_{1,2}(C_1) \ \mu, \sigma_{1,2}(H_1) \ \dots \ \mu, \sigma_{1,2}(G_J) \ \mu, \sigma_{1,2}(C_J) \ \mu, \sigma_{1,2}(H_J))$	$3 \times 3 \times J$
QWF	$\mathbf{v} = (\mu, \sigma_{1,2}(S_1) \ \dots \ \mu, \sigma_{1,2}(S_J))$	$3 \times J$

J is the number of iterations of the WTs.

TABLE III
VISUAL ASPECT AND DISTRIBUTION OF THE 32×32 BLOCKS PER CLASS OF LUNG TISSUE PATTERN

visual aspect					
tissue type	healthy	emphysema	ground glass	fibrosis	micronodules
hand-drawn ROIs	150	101	427	473	297
32×32 blocks	5167	1127	2313	3113	6133
patients	7	6	32	37	16

Note that a patient may have several types of lung disorders.

of healthy annotated tissue, as finding HRCT scans of healthy patients was difficult.

E. Validation Strategy

In order to estimate the generalization performance of the classification of 32×32 blocks, a leave-one-patient-out cross validation (LOPO CV) is used. LOPO CV splits training and testing sets based on patients and has several advantages when compared to other validation methods [56]. First, when compared to the classical leave-one-out (LOO) CV, it ensures that all ROIs belonging to the same patient are contained in the same fold and thus do not allow to train and test with identical patients. This situation corresponds to the clinical routine where the CAD system is trained using the entire database and unseen ROIs from an unknown patient are classified. Second, LOPO has the advantage of LOO where global experience is perfectly reproducible when compared to N_{folds} -fold CV because no random draw is carried out with LOO to create the folds. This is desirable when searching for optimal parameters where the variation of the evaluation conditions introduced by a random draw of the folds in N_{folds} -fold CV can lead to an inappropriate choice of parameters, especially when the number of patients is fairly low for some classes. At last, the computational cost is affordable with N_{folds} equals to the number of cases N_{cases} . The number of classifiers to train is equal to N_{cases} .

III. RESULTS

The visual features are evaluated and compared through experiments to identify groups of visual features (GLH and texture) that allow the best discrimination among the patterns for further integration into a CAD.

A. QWF and GLH

The normalized histograms of the most frequent patterns and their respective mean, standard deviation, along with third and fourth moments (skewness and kurtosis) are depicted in Fig. 4.

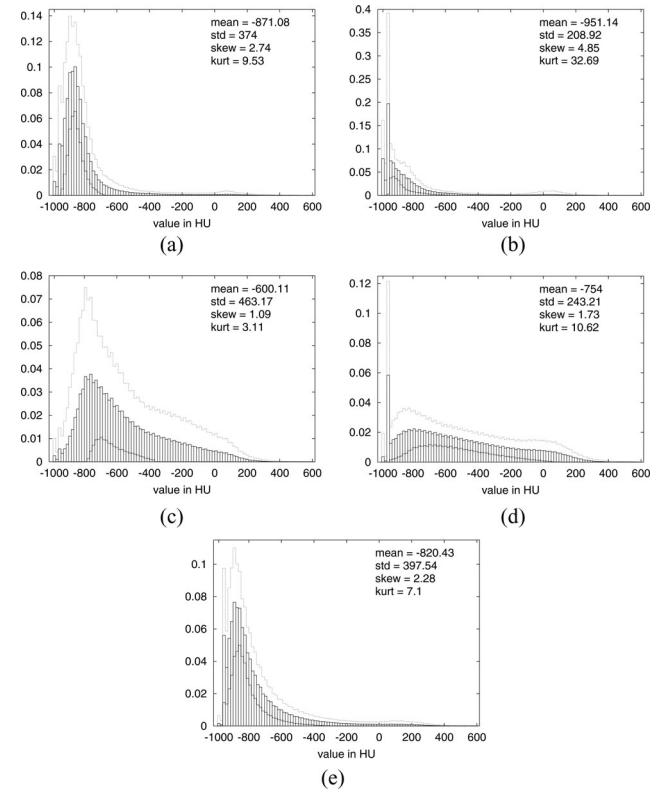


Fig. 4. (a)–(e) GLHs of the most frequent lung tissue patterns. The standard deviation of each bin is shown with dashed lines. Each histogram is obtained from the average of histograms computed over 32×32 blocks of the full dataset, so single cases include a much larger variability. (a) Healthy. (b) Emphysema. (c) Ground glass. (d) Fibrosis. (e) Micronodules.

The confusion matrix of a LOPO CV with all 85 patients using QWF and GLH is shown in Table IV. The associated performance measures are listed in Table V.

The complementarity of QWF and GLH features is investigated in Fig. 5.

TABLE IV
CONFUSION MATRIX OF THE BLOCKWISE CLASSIFICATION USING QWF WITH ORDER $\gamma = 2$ AND GLH FEATURES AND LOPO CV IN %

	healthy	emphysema	ground glass	fibrosis	micronodules
healthy	67.28	6.87	7.22	1.22	17.41
emphysema	4.91	78.72	4.52	7.25	4.6
ground glass	11.79	0.75	71.36	12.87	3.24
fibrosis	2.28	2.75	8.03	82.73	4.21
micronodules	10.96	0.67	3.09	3.71	81.56

TABLE V
PERFORMANCE MEASURES OF THE BLOCKWISE CLASSIFICATION USING QWF WITH ORDER $\gamma = 2$ AND GLH FEATURES AND LOPO CV IN %

	recall	precision	specificity	F-measure	accuracy
healthy	58.92	67.28	94.31	62.83	88.45
emphysema	77.57	78.72	97.87	78.14	96
ground glass	80.92	71.36	90.55	75.84	88.38
fibrosis	80.8	82.73	93.97	81.75	90.5
micronodules	80.88	81.56	93.78	81.22	90.51

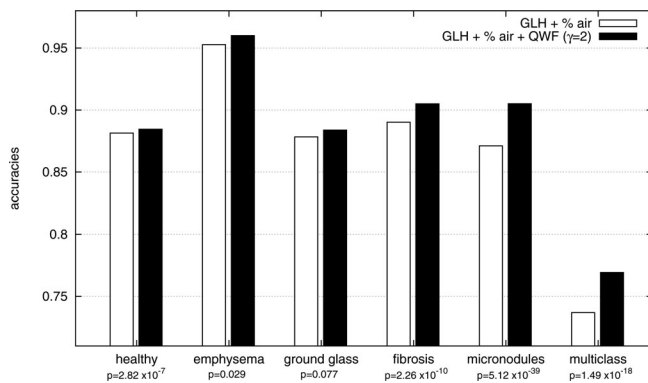


Fig. 5. Contribution of the QWF features with $\gamma = 2$ when compared to GLH with air percentage only. The statistical significance of each comparison is assessed using a paired t-test (two-tailed). The comparisons are all statistically significant ($p \leq 0.05$) except for ground glass.

B. DWF Versus QWF

The abilities of DWF and QWF for classifying the lung tissue patterns are compared. Various orders of the isotropic polyharmonic B -spline wavelets β_γ are evaluated with $\gamma = 2, 3, 4$. To compare performance, four levels of DWFs were performed using B -spline wavelets of degree $\alpha = 1, 2, 3$. The equivalent order of derivatives γ corresponds to $\alpha + 1$ [see (4)]. Pairwise comparisons of classification accuracies using DWF versus QWF for several degrees and orders ($\alpha + 1, \gamma$) are shown in Table VI.

IV. DISCUSSION

The properties of GLH features are discussed in Section IV-A, and their complementarity with QWF features is analyzed in Section IV-B. The comparison of QWF and DWF is dissected in Section IV-C. The blockwise categorization of the lung tissue is discussed in Section IV-D considering its usage as a CAD system. Global remarks on visual features used to characterize lung tissue patterns in HRCT imaging are provided in Section IV-E.

TABLE VI
MEAN ACCURACIES IN % OF COMBINED WAVELET AND GLH FEATURES USING A LOPO CV WITH ALL 85 CASES

class		$\alpha=1$ $\gamma=2$	$\alpha=2$ $\gamma=3$	$\alpha=3$ $\gamma=4$	p
healthy	DWF	<u>87.64</u>	86.64	87.63	$p = 0.008$
	QWF	88.45	88.4	87.26	
emphysema	DWF	95.86	95.76	<u>96.03</u>	$p = 0.297$
	QWF	96	96.07	95.81	
ground glass	DWF	<u>87.64</u>	86.52	86.64	$p = 1.94 \times 10^{-5}$
	QWF	88.38	87.29	86.62	
fibrosis	DWF	89.96	90.9	88.65	$p = 0.446$
	QWF	<u>90.51</u>	88.9	88.53	
micronodules	DWF	<u>87.91</u>	87.2	87.63	$p = 0.008$
	QWF	90.51	87.39	87.09	
multiclass	DWF	<u>74.51</u>	73.52	73.3	$p = 1.26 \times 10^{-8}$
	QWF	76.93	74.03	72.66	

Best accuracies are marked in bold. The statistical significances of the comparisons of best performing orders between DWF and QWF are assessed using a paired t-test (two-tailed). Best orders are underlined.

A. GLH of the Lung Tissue Patterns

GLHs show high variability of their distributions among the five patterns. The healthy pattern is constituted of soft tissue with HU values in the range $[-1050, -300]$ and has a mean value of -871 HU. Emphysema patterns have the lowest mean value of -951 HU. Distribution of the ground glass voxels is close to those of fibrosis, which highlights the need for texture features that characterize the spatial organization of the voxels to separate the two patterns. According to the notable variation among the gray-level distributions, features such as histogram bins, mean, variance, skewness, and kurtosis of the distribution along with air components may have strong discriminative potential for the classification of the patterns. An accurate description of the gray levels can be obtained by reducing the number of bins to 22 corresponding to pixel values in $[-1050, 600]$ because the bins outside this interval are very sparsely populated. $[-1050, 600]$ corresponds to an extended interval of the lung HU values that includes pathological tissue of high density (e.g., calcified nodules [57]). Because the majority of the information is contained in low HU values, a nonlinear binning with bins of exponential width was tested but did not lead to better results. A reason for this is that the information allowing to differentiate among the patterns is mostly contained in HU values of $[-700, -200]$ (see Fig. 4). It is important to note that although the shapes of the gray-level distributions are distinctive in Fig. 4, they are averaged over a large number of cases and thus correspond to the center of mass of each class in the GLH feature space. Large intraclass variations are observed as depicted in Fig. 6, which makes the classification task difficult. Additional features along with machine learning algorithms able to draw intricate decision boundaries are required for an accurate recognition of the lung tissue patterns. This is especially true with healthy tissue [58]. The blockwise analysis also tends to increase the intraclass variability.

B. QWF and GLH

The complementarity of QWF and GLH features is investigated in Fig. 5. The combination of QWF and GLH features shows statistically significant improvement in classification




Fig. 6. (a) and (b) Intraclass variations for the class healthy. Histograms of the two blocks show two outliers that mark the boundaries of the class. The block in (b) contains several bronchovascular structures increasing the global density.

accuracy for all patterns ($p \leq 0.029$) except ground glass ($p = 0.077$) and is particularly effective when carrying out multiclass classification ($p = 1.49 \times 10^{-18}$). The GLH and percentage of air describe the composition of the tissue through its densities whereas QWF complementarily characterizes its structure. For instance, fibrosis is characterized by sharp transitions between high density tissue and small air bubbles yielding a specific spectrum signature that is successfully learned by QWF features.

C. DWF Versus QWF

Pairwise comparisons shown in Table VI indicate that QWF outperforms DWF in 61% of the comparisons (11 among 18) for all degrees/orders. High statistical significance of the differences of accuracy in the multiclass configuration is observed ($p = 1.26 \times 10^{-8}$). Isotropic polyharmonic B -spline wavelets (QWF) with order $\gamma = 2$ allowed an arithmetic mean accuracy of 76.9% of correct predictions among the five lung tissue classes with balanced precision over the classes (see Table V). This global increase in performance is primarily due to the better isotropy properties of the nonseparable WT, thanks to their close connection to the isotropic Laplacian operator. On the contrary, the favored directions of the separable transform lead to noisy features creating nonhomogeneous clusters of instances belonging to the same class in the feature space, which decreases global classification performance due to a lack of regularization. Although having an influence on global accuracy, the finer scale progression allowed by the quincunx subsampling increases the precision of the classification, e.g., by reducing confusions between patterns with well-defined object size such as micronodules and bronchovascular structures. The highest improvements using QWF are achieved for classes healthy ($p = 0.008$), ground glass ($p = 1.94 \times 10^{-5}$), and micronodules ($p = 0.008$). The reduced number of subbands with QWF (one per scale for QWF and three for DWF) is also an important clue to avoid the curse of dimensionality. QWF features allow to characterize two times more scales with even a reduced number of features when compared to DWF. In this study, the DWF requires 36 features ($4 \times 3 \times 3$) for characterizing four scales when the QWF uses no more than 24 features (8×3) to characterize eight scales. This highlights that albeit separable transforms being faster, because the convolutions are carried out in one dimension, they yield several subbands per scale, which is less adapted for multidimensional signal analysis where the amount of information increases exponentially. Degree/order

$\alpha = 1$ and $\gamma = 2$ allow for best classification for all patterns except for fibrosis although being very close ($p = 0.446$) to the best accuracy as well. A first reason for this is that the scale progression is well initialized and fits best the characteristic structures of each pattern leading to increased classification accuracy. These degrees/orders are best tradeoffs between spatial and spectral localization. Second, this suggests that β_γ begins to have acceptable isotropic properties already starting from $\gamma = 2$ and lower values favor the x - and y -directions.

D. Blockwise Classification With QWF

Experience with the segmentation of 85 lung volumes shows that the 3-D region growing with closing allows an almost fully automatic segmentation. The trachea is included as lung tissue in many cases but automatic solutions exist to remove it. Manual corrections are required when the closing operation cannot fill large regions of consolidated tissue.

The automatic segmentation of the lung tissue is crucial for the success of the CAD. The accuracies obtained in Table V show that the SVM classifier can learn effectively from the hybrid feature space. The recurrent confusion between healthy and micronodule patterns in Table IV suggests that the decision boundaries are not trivial in some cases. These confusions have an impact on the performance measures of healthy patterns in Table V. The small nodules in micronodules are mixed with bronchovascular structures contained in healthy tissue. This is a major limitation of the 2-D QWF used in this study. The use of 3-D isotropic polyharmonic B -spline wavelets [26] in near isotropic multidetector CT (MDCT) image series may be able to overcome the confusions between healthy and micronodules as the bronchovascular structures are cylindrical, whereas small nodules are roughly round [59]. Unfortunately, the HRCT protocol used at the HUG does not allow for a true 3-D analysis as the spacing between the slices is 10 mm, yielding extremely anisotropic volumetric images but limiting the radiation dose delivered to the patients when compared to the MDCT protocol. Using the clinical context of the images such as the age of the patient shows allows clarifications between visually similar patterns [20]. Micronodules in a 20-year old subject are very visually similar to healthy tissue surrounded by vessels of an 80-year old person. The clinical context allows significant improvements particularly for the characterization of micronodules [20].

Table IV shows recurrent confusion between ground glass and fibrosis. This may be partially explained by the fact that fibrosis patterns are most often accompanied by small regions of ground glass because of the redistribution of the perfusion to the functional tissue remaining. This has the effect to overload the healthy tissue, which thus has the visual appearance of ground glass because of increased attenuation. However, during the annotation sessions, the label fibrosis was assigned to the entire ROI leading to classification errors when the system correctly detects the small ground glass regions. At the border of the lungs, misclassifications occur due to the response of the wavelets to the sharp change of intensity. A solution to this is to use a symmetry of the lung tissue using the tangent to the lung border as axis. To remove noise in the blockwise classification,

a 3-D voting of the classifier may avoid small inconsistent isolated regions. The anatomical regions of the lung do not have the same prior probabilities of having abnormal tissue. For instance, fibrosis is prevailing at the periphery of the lung bases. Integrating the anatomical regions for improving the classification performance requires to automatically segment the lung regions and showed significant improvement of classification accuracy in [10] and [60].

E. Global Remarks on the Visual Features

Wavelets and filter banks are not the most popular feature extraction techniques for the characterization of parenchymal textures in the literature [5]. Although a direct comparison of the performance reported does not make sense as the types of tissue included, the validation methods and the dataset used vary much among the studies; we believe that wavelet-based texture features are more adapted to analyze texture patterns in medical images than many other methods. Filtering techniques and wavelets offer an overcomplete feature set able to fit most of the textures to efficiently derive features from the coefficients. Wavelet-based texture features are covering the entire spectrum, which truly allows detecting the spectral signatures of the patterns, being complementary to the measures of density using GLH. Among the experiments in this study, the combination of QWF with GLH showed to best discriminate lung tissue patterns, which is consistent with our hypothesis on the requirement of affine invariance of the texture characterization. No comparison with directional analysis is carried out. Preliminary tests using the methods described in [49] did not lead to better results suggesting that isotropic analysis is adapted to lung tissue analysis that may not contain prevailing orientation. This is contradictory to results obtained in [16] and [24], and further investigations are required to learn about the orientations in the lung tissue patterns.

As highlighted in the literature, we observed that in this dataset, the selected classes of lung tissue and the validation methods have all a strong influence on the measured performance of the CAD. The performance shown in Tables IV and V has to be considered as the baseline performance in accordance with actual clinical situations as the LOPO cross validation simulates true clinical routine usage of the CAD. The block-wise classification allows us to detect abnormalities in HRCT images series and provides a 3-D map of the lung tissue (see Fig. 7) that indicates the suspicious regions to unexperienced radiologists [54].

V. CONCLUSION AND PERSPECTIVES

We have developed a near-affine-invariant texture feature set extracted using a tailored WT, i.e., isotropic polyharmonic B -spline wavelets deployed as a redundant frame transform with quincunx subsampling. The quantification of the wavelet coefficients using parameters of a mixture of Gaussians combined with GLH bins in HUs efficiently characterizes five lung tissue patterns in HRCT imaging with statistically significant improvements when compared to density-based features ($p = 1.49 \times 10^{-18}$) and standard DWFs ($p = 1.26 \times 10^{-8}$). While limitations occur as no directional analysis is carried out,

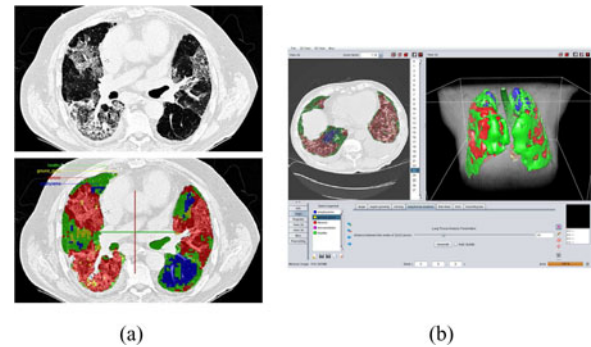


Fig. 7. Screenshots of the 2-D/3-D diagnosis aid tools for the radiologists. (a) 2D map of the lung tissue. (b) 2D/3D tissue map displayed in a graphical user interface.

the assumption was that no prevailing directions are contained in lung tissue patterns projected on 2-D axial slices in HRCT. The LOPO cross validation used to assess the performance of the visual features is in accordance with a clinical use of the system.

An N -dimensional extension of the WT is desirable for the characterization of organs and tissue in volumetric medical images. The custom WT is easily extendable to three dimensions with the limitation that the equivalent of quincunx subsampling will not yield similarity transforms [45]. It is important to note that 3-D WTs are not appropriate for analyzing HRCT image series because filtering along the z -axis with a very low axial resolution (20 to 50 slices with 10-mm distance) leads to coarse blurring of the relevant information.

We believe that the proposed methods provide accurate descriptions of textures that do not contain prevailing orientations. The characterization of biomedical tissue is often related to texture [61], [62], opening a wide variety of applications for the texture analysis framework proposed. Modifications concerning the GLH (based on HU) and tuning of the parameters of QWF (i.e., the order γ and number of iterations) are required.

REFERENCES

- [1] W. R. Webb, N. L. Müller, and D. P. Naidich, *High-Resolution CT of the Lung*. Baltimore, MD: Williams Wilkins, 2001.
- [2] P. Stark, *High Resolution Computed Tomography of the Lungs*, D. S. Basow, Ed. Waltham, MA: UpToDate, 2008.
- [3] K. Doi, "Current status and future potential of computer-aided diagnosis in medical imaging," *Brit. J. Radiol.*, vol. 78, pp. 3–19, 2005.
- [4] R. Kakinuma, H. Ohmatsu, M. Kaneko, K. Eguchi, T. Naruke, K. Nagai, Y. Nishiwaki, A. Suzuki, and N. Moriyama, "Detection failures in spiral CT screening for lung cancer: Analysis of CT findings," *Radiology*, vol. 212, no. 1, pp. 61–66, 1999.
- [5] A. Depeursinge, "Affine-invariant texture analysis and retrieval of 3D medical images with clinical context integration," Ph.D. dissertation, Serv. Med. Informat. Univ. of Geneva, 2010.
- [6] S. Delorme, M.-A. Keller-Reichenbecher, I. Zuna, W. Schlegel, and G. V. Kaick, "Usual interstitial pneumonia: Quantitative assessment of high-resolution computed tomography findings by computer-assisted texture-based image analysis," *Investigat. Radiol.*, vol. 32, no. 9, pp. 566–574, 1997.
- [7] C. Brodley, A. Kak, C. Shyu, J. Dy, L. Broderick, and A. M. Aisen, "Content-based retrieval from medical image databases: A synergy of human interaction, machine learning and computer vision," in *Proc. 10th Nat. Conf. Artif. Intell.*, Orlando, FL, 1999, pp. 760–767.
- [8] C.-R. Shyu, C. E. Brodley, A. C. Kak, A. Kosaka, A. M. Aisen, and L. S. Broderick, "ASSERT: A physician-in-the-loop content-based retrieval system for HRCT image databases," *Comput. Vis. Image Understand.*, vol. 75, no. 1–2, pp. 111–132, 1999.

- [9] F. Chabat, G.-Z. Yang, and D. M. Hansell, "Obstructive lung diseases: Texture classification for differentiation at CT," *Radiology*, vol. 228, no. 3, pp. 871–877, 2003.
- [10] T. Zrimec and J. S. J. Wong, "Improving computer aided disease detection using knowledge of disease appearance," in *Proc. MEDINFO*, 2007, vol. 129, pp. 1324–1328.
- [11] N. Kim, J. B. Seo, Y. Lee, J. G. Lee, S. S. Kim, and S. H. Kang, "Development of an automatic classification system for differentiation of obstructive lung disease using HRCT," *J. Digit. Imag.*, vol. 22, no. 2, pp. 136–148, 2009.
- [12] R. Uppaluri, E. A. Hoffman, M. Sonka, G. W. Hunninghake, and G. McLennan, "Interstitial lung disease: A quantitative study using the adaptive multiple feature method," *Amer. J. Respirat. Crit. Care Med.*, vol. 159, no. 2, pp. 519–525, 1999.
- [13] L. Sørensen, S. B. Shaker, and M. D. Bruijne, "Quantitative analysis of pulmonary emphysema using local binary patterns," *IEEE Trans. Med. Imag.*, vol. 29, no. 2, pp. 559–569, Feb. 2010.
- [14] Y. Xu, M. Sonka, G. McLennan, J. Guo, and E. A. Hoffman, "MDCT-based 3D texture classification of emphysema and early smoking related lung pathologies," *IEEE Trans. Med. Imag.*, vol. 25, no. 4, pp. 464–475, Apr. 2006.
- [15] V. A. Zavaletta, B. J. Bartholmai, and R. A. Robb, "Nonlinear histogram binning for quantitative analysis of lung tissue fibrosis in high-resolution CT data," in *Medical Imaging 2007: Physiology, Function, and Structure from Medical Images*, vol. 6511, A. Manduca and X. P. Hu, Eds. Bellingham, WA: SPIE, 2007, p. 65111Q.
- [16] K. Vo and A. Sowmya, "Directional multi-scale modeling of high-resolution computed tomography (HRCT) lung images for diffuse lung disease classification," in *Proc. 13th Int. Conf. Comput. Anal. Images Patterns*, vol. 5702. New York: Springer-Verlag, 2009, pp. 663–671.
- [17] A. Depeursinge, D. Sage, A. Hidki, A. Platon, P.-A. Poletti, M. Unser, and H. Müller, "Lung tissue classification using Wavelet frames," in *Proc. 29th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2007, pp. 6259–6262.
- [18] A. Depeursinge, D. Van De Ville, M. Unser, and H. Müller, "Lung tissue analysis using isotropic polyharmonic B-spline wavelets," in *Proc. MIC-CAI Workshop Pulmonary Image Anal.*, New York, 2008, pp. 125–134.
- [19] A. Depeursinge, A. Vargas, A. Platon, A. Geissbuhler, P.-A. Poletti, and H. Müller, "3D case-based retrieval for interstitial lung diseases," in *Medical Content-Based Retrieval for Clinical Decision Support* (Lecture Notes in Computer Science Series). New York: Springer-Verlag, 2010, pp. 39–48.
- [20] A. Depeursinge, D. Racoceanu, J. Iavindrasana, G. Cohen, A. Platon, P.-A. Poletti, and H. Müller, "Fusing visual and clinical information for lung tissue classification in high-resolution computed tomography," *Artif. Int. Med.*, vol. 50, no. 1, pp. 13–21, 2010.
- [21] A. Depeursinge, J. Iavindrasana, A. Hidki, G. Cohen, A. Geissbuhler, A. Platon, P.-A. Poletti, and H. Müller, "Comparative performance analysis of state-of-the-art classification algorithms applied to lung tissue categorization," *J. Digit. Imag.*, vol. 23, no. 1, pp. 18–30, 2010.
- [22] C.-T. Liu, P.-L. Tai, A. Y.-J. Chen, C.-H. Peng, T. Lee, and J.-S. Wang, "A content-based CT lung retrieval system for assisting differential diagnosis images collection," in *Proc. 2nd Int. Conf. Multimedia Expo.*, Tokyo, Japan, 2001, pp. 172–177.
- [23] I. C. Sluimer, P. F. van Waes, M. A. Viergever, and B. van Ginneken, "Computer-aided diagnosis in high resolution CT of the lungs," *Med. Phys.*, vol. 30, no. 12, pp. 3081–3090, 2003.
- [24] A. Tlouee, H. Abrishami-Moghaddam, R. Garnavi, M. Forouzanfar, and M. Giti, "Texture analysis in lung HRCT images," in *Proc. Digital Imaging Comput.: Tech. Appl.*, 2008, pp. 305–311.
- [25] M. Gangeh, L. Sørensen, S. Shaker, M. Kamel, M. D. Bruijne, and M. Loog, "A texon-based approach for the classification of lung parenchyma in CT images," in *Proc. Med. Image Comput. Comput.-Assisted Intervention* (Lecture Notes in Computer Science Series), vol. 6363. Berlin, Germany: Springer-Verlag, 2010, pp. 595–602.
- [26] D. Van De Ville, T. Blu, and M. Unser, "Isotropic polyharmonic B-Splines: Scaling functions and wavelets," *IEEE Trans. Image Process.*, vol. 14, no. 11, pp. 1798–1813, Nov. 2005.
- [27] M. Unser, "Texture classification and segmentation using wavelet frames," *IEEE Trans. Image Process.*, vol. 4, no. 11, pp. 1549–1560, Nov. 1995.
- [28] G. van de Wouwer, P. Scheunders, and D. V. Dyck, "Rotation-invariant texture characterization using isotropic wavelet frames," in presented at the 14th Int. Conf. Pattern Recognit. Brisbane, Qld., Australia, 1998.
- [29] P. G. Hartley, J. R. Galvin, G. W. Hunninghake, J. A. Merchant, S. J. Yagla, S. B. Speakman, and D. A. Schwartz, "High-resolution CT-derived measures of lung density are valid indexes of interstitial lung disease," *J. Appl. Phys.*, vol. 76, no. 1, pp. 271–277, 1994.
- [30] A. Depeursinge, H. Müller, A. Hidki, P.-A. Poletti, A. Platon, and A. Geissbuhler, "Image-based diagnostic aid for interstitial lung disease with secondary data integration," in *Medical Imaging 2007: Computer-Aided Diagnosis*, vol. 6514, M. L. Giger and N. Karssemeijer, Eds. Bellingham, WA: SPIE, 2007, p. 65143P.
- [31] H. Sumikawa, T. Johkoh, S. Yamamoto, K. Takahei, T. Ueguchi, Y. Ogata, M. Matsumoto, Y. Fujita, J. Natsag, A. Inoue, M. Tsubamoto, N. Mihara, O. Honda, N. Tomiyama, S. Hamada, and H. Nakamura, "Quantitative analysis for computed tomography findings of various diffuse lung diseases using volume histogram analysis," *J. Comput. Assist. Tomography*, vol. 30, no. 2, pp. 244–249, 2006.
- [32] M. Unser and A. Aldroubi, "A review of wavelets in biomedical applications," *Proc. IEEE*, vol. 84, no. 4, pp. 626–638, Apr. 1996.
- [33] M. Fernández and A. Mavilio, "Texture analysis of medical images using the wavelet transform," in *MEDICAL PHYSICS: 6th Mexican Symp. Med. Phys.*, 2002, vol. 630, pp. 164–168.
- [34] B. van Ginneken and B. M. ter Haar Romeny, "Multi-scale texture classification from generalized locally orderless images," *Pattern Recognit.*, vol. 36, no. 4, pp. 899–911, 2003.
- [35] S. G. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. San Diego, CA: Academic, 1999.
- [36] A. Laine and J. Fan, "Frame representations for texture segmentation," *IEEE Trans. Image Process.*, vol. 5, no. 5, pp. 771–780, May 1996.
- [37] A. V. Nevel, "Texture classification using wavelet frame decompositions," in *Proc. 31st Asilomar Conf. Signals, Syst. Comput. Conf. Rec.*, 1997, vol. 1, pp. 311–314.
- [38] T. Randen and J. H. Husoy, "Filtering for texture classification: A comparative study," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 4, pp. 291–310, Apr. 1999.
- [39] S. Li, J. T. Kwok, H. Zhu, and Y. Wang, "Texture classification using the support vector machines," *Pattern Recognit.*, vol. 36, no. 12, pp. 2883–2893, 2003.
- [40] S. Liapis and G. Tziritas, "Color and texture image retrieval using chromaticity histograms and wavelet frames," *IEEE Trans. Multimedia*, vol. 6, no. 5, pp. 676–686, Oct. 2004.
- [41] A. Chebira, "Adaptive multiresolution frame classification of biomedical images," Ph.D. dissertation, Dept. Biomed. Eng., Carnegie Mellon Univ., 2008.
- [42] D. Marr and E. Hildreth, "Theory of edge detection," *Proc. Roy. Soc. Lond.*, vol. 207, no. 1167, pp. 187–217, 1980.
- [43] D. Van De Ville and M. Unser, "Complex wavelet bases, steerability, and the Marr-like pyramid," *IEEE Trans. Image Process.*, vol. 17, no. 11, pp. 2063–2080, Nov. 2008.
- [44] M. P. Feilner, D. Van De Ville, and M. Unser, "An orthogonal family of quincunx wavelets with continuously adjustable order," *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 499–510, Apr. 2005.
- [45] D. Van De Ville, T. Blu, and M. Unser, "On the multidimensional extension of the quincunx subsampling matrix," *IEEE Signal Process. Lett.*, vol. 12, no. 2, pp. 112–115, Feb. 2005.
- [46] C.-L. Chang, A. Maleki, and B. Girod, "Adaptive wavelet transform for image compression via directional quincunx lifting," in *Proc. 7th IEEE Workshop Multimedia Signal Process.*, 2005, pp. 1–4.
- [47] M. F. Toubin, F. Truchetet, E. P. Verrecchia, C. Dumont, and M. A. Abidi, "Multiresolution description of range images through 2D quincunx wavelet analysis," in *Wavelet Appl. VI*, vol. 3723, H. H. Szu, Ed. Bellingham, WA: SPIE, 1999, pp. 350–360.
- [48] A. Mojsilovic, M. Popovic, S. Markovic, and M. Krstic, "Characterization of visually similar diffuse disease from B-scan liver images using non-separable wavelet transform," *IEEE Trans. Med. Imag.*, vol. 17, no. 4, pp. 541–549, Aug. 1998.
- [49] D. Van De Ville, D. Sage, K. Balać, and M. Unser, "The Marr wavelet pyramid and multiscale directional image analysis," presented at the 16th Eur. Signal Process. Conf., Lausanne, Switzerland, 2008.
- [50] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.
- [51] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc. Series B Methodol.*, vol. 39, no. 1, pp. 1–38, 1977.
- [52] M. Unser, "Splines: A perfect fit for medical imaging," in *Progress in Biomedical Optics and Imaging*, vol. 4684, M. Sonka and M. Fitzpatrick, Eds. Bellingham, WA: SPIE, 2002, pp. 225–236.
- [53] P. Korfiatis, S. Skiadopoulos, P. Sakellariopoulos, C. Kalogeropoulou, and L. Costaridou, "Automated 3D segmentation of lung fields in thin slice CT exploiting wavelet preprocessing," *Comput. Anal. Images Patterns*, vol. 4673, pp. 237–244, 2007.

- [54] A. Depeursinge, A. Vargas, F. Gaillard, A. Platon, A. Geissbuhler, P.-A. Poletti, and H. Müller, "Case-based lung image categorization and retrieval for interstitial lung diseases: Clinical workflows," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 7, no. 1, pp. 97–110, 2012.
- [55] A. Depeursinge, A. Vargas, A. Platon, A. Geissbuhler, P.-A. Poletti, and H. Müller, "Building a reference multimedia database for interstitial lung diseases," *Comput. Med. Imag. Graph.*, vol. 36, no. 3, pp. 227–238, 2012.
- [56] M. Dundar, G. Fung, L. Bogoni, M. Macari, A. Megibow, and B. Rao, "A methodology for training and validating a CAD system and potential pitfalls," in *Proc. 18th Int. Congr. Exhib. Comput. Assist. Radiol. Surg.*, 2004, vol. 1268, pp. 1010–1014.
- [57] E. M. Marom, "CT of the solitary pulmonary nodule—A commentary," *Amer. J. Roentgenol.*, vol. 190, no. 5, pp. 1154–1155, 2008.
- [58] H. Adams, M. S. Bernard, and K. McConnochie, "An appraisal of CT pulmonary density mapping in normal subjects," *Clin. Radiol.*, vol. 43, no. 4, pp. 238–242, 1991.
- [59] U. Studler, T. Gluecker, G. Bongartz, J. Roth, and W. Steinbrich, "Image quality from high-resolution CT of the lung: Comparison of axial scans and of sections reconstructed from volumetric data acquired using MDCT," *Amer. J. Roentgenol.*, vol. 185, no. 3, pp. 602–607, 2005.
- [60] A. Depeursinge, T. Zrimec, S. Busayarat, and H. Müller, "3D lung image retrieval using localized features," in *Proc. Med. Imag. 2011: Comput.-Aided Diagnosis*, SPIE, 2011, vol. 7963, p. 79632E.
- [61] G. Castellano, L. Bonilha, L. M. Li, and F. Cendes, "Texture analysis of medical images," *Clin. Radiol.*, vol. 59, no. 12, pp. 1061–1069, 2004.
- [62] R. A. Lerski, K. Straughan, L. R. Schad, D. Boyce, S. Blüml, and I. Zuna, "VIII. MR image texture analysis—an approach to tissue characterization," *Magn. Reson. Imag.*, vol. 11, no. 6, pp. 873–887, 1993.



Adrien Depeursinge received the B.Sc. and M.Sc. degrees in electrical engineering from the Swiss Federal Institute of Technology, Lausanne, Switzerland, in 2000 and 2005, respectively, with a specialization in signal and image processing.

From 2006 to 2010, he performed his Ph.D. thesis on medical image analysis with a focus on texture analysis and content-based image retrieval at the University Hospitals of Geneva (HUG). He is currently a Research Fellow at the University of Applied Sciences Western Switzerland, Sierre, Switzerland and

the HUG.

Dr. Depeursinge was the recipient of the 2011 German Association for Medical Informatics, Biometry and Epidemiology Award in medical informatics for his Ph.D. thesis.



Dimitri Van De Ville (M'02) received the M.S. and Ph.D. degrees in computer science and the from Ghent University, Ghent, Belgium, in 1998, and 2002, respectively.

After a Postdoctoral stay (during 2002–2005) in the Biomedical Imaging Group of Prof. M. Unser, Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, he joined Signal Processing Unit, University Hospitals of Geneva, as part of the Centre d'Imagerie Biomedicale. In 2009, he was the recipient of a Swiss National Science Foundation Professorship and currently holds a joint position at the University of Geneva, Geneva, Switzerland, and the EPFL. His research interests include wavelets, sparsity, pattern recognition, and their applications in biomedical imaging, such as functional magnetic resonance imaging.

Dr. Van De Ville is currently the Chair of the Bio Imaging and Signal Processing TC of the IEEE Signal Processing Society (2012–2013). He is the recipient of the Pfizer Research Award in 2012 in the category neurosciences for his research on scale-free dynamics of EEG microstates.



Alexandra Platon was born in Iasi, Romania, in 1966. She currently is an attending radiologist in the Emergency Radiology unit, Geneva University Hospital, Switzerland. Her area of research is emergency radiology (trauma and non-trauma).



Antoine Geissbuhler is currently a Professor of Medicine, Chairman of the Department of Radiology and Medical Informatics at Geneva University, Geneva, Switzerland, Director of the Division of eHealth and Telemedicine at Geneva University Hospitals, Geneva, and President of the Health-On-the-Net Foundation, Geneva. He is also the President of the International Medical Informatics Association and Vice-chair of the HIMSS Europe Governing Council.

He trained as a physician at Geneva University where he specialized in internal medicine, then, after a post-doctoral fellowship, became Associate Professor of biomedical informatics at Vanderbilt University. In 1999, he returned to Geneva to take the responsibility of medical information systems at Geneva University Hospitals. In 2005, his efforts in developing telemedicine and tele-education are recognized by the creation of the UNESCO chair for telemedicine and multidisciplinary teaching, and, in 2010, of the World Health Organization collaborating center for eHealth and telemedicine.



Pierre-Alexandre Poletti was born in 1964. He completed medical studies at the Faculty of Medicine, University of Geneva, Geneva, Switzerland.

After a board in medical radiology at the University Hospital of Geneva (HUG), he completed a fellowship in emergency radiology at the University of Maryland Medical System, Baltimore. He is currently responsible of the Emergency Radiology Unit, Geneva University Hospital, Geneva. His research interests include the development and improvement of emergency medical imaging, especially by

CT-scanner.

Mr. Poletti is also an expert for the Radiation Safety Committee of HUG and the Advisory Committee of atomic–biological–chemical protection coordinates of the State Council in Geneva.



Henning Müller studied medical informatics at the University of Heidelberg, Heidelberg, Germany, from 1992 to 1997, with a specialization in signal and image processing. He received the Ph.D. degree on content-based image retrieval at the University of Geneva, Geneva, Switzerland, in 2002, with a research stay at Monash University, Melbourne, Vic., Australia, in 2001.

After a diploma thesis in the telemedicine project Chili, he was at Daimler-Benz Research and Technology North America, Portland, OR, for six months, with a scholarship of the CarlDuisberg Society. In 2002, he joined Medical Informatics Service, University Hospitals and the University of Geneva. He started the medical image retrieval project medGIFT and initiated the medical image retrieval benchmark ImageCLEFmed. Since 2007, he has been a Full Professor at the University of Applied Sciences Western Switzerland, Sierre, Switzerland, while keeping a part-time research position in medical informatics at Geneva. He has published more than 300 scientific articles. He is in the Editorial Boards of several journals and in the committees of various conferences. He has initiated several national and international research projects and currently coordinates the EU project Khresmoi.