

Lung Tissue Classification Using Wavelet Frames

Adrien Depeursinge*, Daniel Sage, Asmâa Hidki, Alexandra Platon, Pierre–Alexandre Poletti, Michael Unser and Henning Müller

Abstract—We describe a texture classification system that identifies lung tissue patterns from high–resolution computed tomography (HRCT) images of patients affected with interstitial lung diseases (ILD). This pattern recognition task is part of an image–based diagnostic aid system for ILDs. Five lung tissue patterns (*healthy, emphysema, ground glass, fibrosis and microdules*) selected from a multimedia database are classified using the overcomplete discrete wavelet frame decomposition combined with grey–level histogram features. The overall multiclass accuracy reaches 92.5% of correct matches while combining the two types of features, which are found to be complementary.

I. INTRODUCTION

Diagnosing Interstitial Lung Disease (ILD) is regarded as a difficult task, even for specialists, as many forms of the disease are rare and thus little experience exists. The diagnosis of ILD is established from the interpretation of several clinical parameters of the patient in addition to radiological findings [1]. Detection of pathologic lung tissue patterns is the first step towards an image–based computer–aided diagnosis system [2]. When the chest x–ray does not carry enough elements to finalize the diagnosis, high–resolution computed tomography (HRCT) is used to provide an accurate assessment of lung tissue patterns [3]. HRCT produces three–dimensional (3D) images of the pulmonary volumes, avoids the superposition of anatomic structures, and is well suited for the assessment of lung tissue texture. 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 characterisation of ILD which is typically diffuse.

Texture analysis in digital image processing has been an active research domain over more than thirty years. In [4], texture in digital images is defined as nonfigurative and cellularly organized areas of pixels. Such patterns can be described by a given spatial organisation of grey levels (e.g., random, periodic). Early examples of texture features are the autocorrelation function, textural edginess, measurements derived from mathematical morphology, run–length and gray–level co–occurrence matrices, the latter being the

most popular of the lot [4,5]. Unfortunately, building co–occurrence matrices from HRCT images where grey–levels are corresponding to Hounsfield Units (H.U.) with values from –1000 H.U. (air) to 1500 H.U. (high density bones) is unrealistic because the number of possible co–occurrences is simply too large to be stored and estimated reliably. This can be handled by compressing the grey scale or by considering an alternative representation in terms of sums and differences [6]. Complementary to the characterization of spatial dependencies, the distribution of grey–level values can be studied through statistical measures of grey–level histograms. The Fourier transform has also been proposed for texture analysis, based on the property that some image patterns (especially periodic ones) are well described in terms of sinusoidal components [5,7]. However, the latter is not appropriate for segmentation because the Fourier transform is global.

An attractive solution for carrying out a more local texture analysis, which also takes into account scale, is to use the wavelet transform (WT) [8]. Wavelets are mathematical analysis functions that decompose signals into different subbands, and then analyse each component with a resolution matching its scale (the analysed image is iteratively subsampled by a factor of 2 in the standard dyadic discrete version). The WT is particularly well suited for the representation of piecewise–smooth signals, as well as stochastic processes with a fractal–like behavior, which partly explains their success in biomedical imaging applications [9,10]. Wavelets come into two flavors: bases or frames. Wavelet bases provide a one–to–one decomposition (that may be orthogonal or not), which makes them ideally suited for image compression. They have a fast implementation, but their main drawback is their lack of translation invariance. Wavelet frames, on the other hand, are redundant and offer more flexibility for image analysis. Moreover, one can easily obtain a translation–invariant frame from a wavelet basis by removing the sub–sampling part of the algorithm; this yields a discrete wavelet frame (DWF) which performs well for texture analysis [11,12]. When compared to the WT, the DWF tends to decrease the variability of the estimated texture features thereby improving classification performance.

A. Texture analysis for lung tissue classification

The majority of papers describing image–based computer–aided diagnosis for lung diseases use texture features to classify lung tissue patterns. In [13], Aisen *et al.* describe a physician–in–the–loop content–based image retrieval system in which the physician delineates a suspicious region of

*This work was supported by the Swiss National Science Foundation (FNS) with grant 205321-109, the equalization fund of University and Hospitals of Geneva (grant 05-9-II) and the EU 6th Framework Program in the context of the KnowARC project (IST 032691).

A. Depeursinge, A. Hidki and H. Müller are with Service of Medical Informatics, University & University Hospitals of Geneva (HUG), Geneva, Switzerland. adrien.depeursinge@sim.hcuge.ch

D. Sage and M. Unser are with the Biomedical Imaging Group, Swiss Federal Institute of Technology Lausanne (EPFL), Lausanne, Switzerland.

A. Platon and P.A. Poletti are with Service of Emergency Radiology, University Hospitals of Geneva (HUG), Geneva, Switzerland.

interest (ROI). The system then matches the ROI against reference ROIs in images in JPEG (and not DICOM) that are already indexed in the database through co-occurrences matrices, grey-level distributions and size of the ROI. The evaluation of this tool in [14] suggests that it improves diagnosis, particularly for nonspecialist radiologists. However, the clustering of grey-level values in H.U. can have a critical influence on the efficiency of co-occurrence matrices to characterise lung tissue patterns.

Non-linear binning of grey-level values for co-occurrence matrices is proposed in [15] in order to quantify lung tissue fibrosis in HRCT data. This approach is less suitable for carrying out quantitative analysis of several classes of lung tissue patterns at once.

In [16], six lung tissue patterns in HRCT data are classified using an adaptive multiple feature method. 22 texture features are extracted from grey-level distribution, run-length, and co-occurrence matrix along with fractal analysis based on the Brownian motion model principle. This approach consists of weighting several features to obtain a classification of the patterns from a training set. The extension of this feature set to 3D in [17] shows noticeable improvement in classification accuracy of lung tissue patterns. However, no features measuring spatial frequencies are used.

A quincunx wavelet transform along with support vector machines (SVM) are used in [18] to classify 5 lung tissue patterns. Nevertheless, the classification task is slightly biased since the training set contains an equal number of instances for each of the five patterns, which is usually not the case in clinical practice.

Despite those works, there has not yet been any direct attempt in addressing the problem of the classification of healthy tissue versus pathological ones, which is the first step for building a diagnosis system for ILD. Here, we present some preliminary results concerning the application of the redundant discrete wavelet frame (DWF) transform to the classification of HRCT images. We consider 5 types of lung tissue patterns associated with ILD.

B. Wavelet frames

The discrete wavelet frame (DWF) decomposition described in [11] consists of analysing the input image $f(\mathbf{x})$ in terms of the overcomplete family of templates

$$\mathcal{S} = \{g_1(\mathbf{x} - \mathbf{l}), \dots, g_I(\mathbf{x} - \mathbf{l}), h_I(\mathbf{x} - \mathbf{l})\}_{\mathbf{l} \in \mathbb{Z}^2}$$

where h_i stands for a lowpass filter at iteration i and g_i a family of highpass filters with $i = 1, \dots, I$. The associated decomposition algorithm is

$$\begin{cases} G_i(\mathbf{x}) := \langle g_i(\mathbf{x} - \mathbf{l}), f(\mathbf{x}) \rangle_{l_2} \\ H_I(\mathbf{x}) := \langle h_I(\mathbf{x} - \mathbf{l}), f(\mathbf{x}) \rangle_{l_2} \end{cases} \quad (1)$$

where $\langle \cdot, \cdot \rangle_{l_2}$ is the l_2 scalar product. G_i contains coefficients generated by the convolution of the image with the highpass filters at iteration i and H_I the convolution of the image with the lowpass filter at the last iteration I . More details together with a justification of the method can be found in [11].

TABLE I
DISTRIBUTION OF THE ROIS PER CLASS OF LUNG TISSUE PATTERN

	healthy	emphysema	ground glass	fibrosis	micronodules
# of ROIs	77	72	113	64	155
# of patients	10	5	13	11	5

II. MATERIAL & METHODS

The dataset used is part of an internal multimedia database of ILD cases [19, 20]. For each case in the database, 99 clinical parameters related to the 15 most frequent ILDs are filled in and relevant non-injected HRCT image series with slice thickness $< 3mm$ are annotated by experienced radiologists through a graphical user interface (GUI) implemented in Java. The interface was adapted in order to meet the needs of the radiologists for the various annotation tasks; it allows high-quality annotations in 3D HRCT data. The data acquisition is still in process; the database actually contains 95 cases with full clinical parameters from which 52 HRCT image series are annotated with 675 ROIs distributed into 12 different classes of lung tissue patterns, including healthy tissue. In this study, the five most frequent patterns in $1mm$ slice thickness (without contrast agent) are selected from the multimedia database in order to test the set of texture features. The selected patterns are *healthy*, *emphysema*, *ground glass*, *fibrosis* and *micronodules*. Distributions of the ROIs are detailed in Table I.

The algorithm for the extraction of the texture features (including wavelet frames) from DICOM files was implemented as an *ImageJ*¹ plugin in Java, and the classification was carried out using the free data mining software *R*².

III. RESULTS

A. Feature extraction

1) *Wavelet frames*: The algorithm used to compute the wavelet frame coefficients $G_i(\mathbf{x})$ and $H_I(\mathbf{x})$ directly implements equation (1). A family of *B-splines* of third order are used as wavelet basis [21]. The coefficients $H_i(\mathbf{x})$ resulting from the convolutions with lowpass filters h_i are kept for each iteration i in order to investigate continuous components of the lung tissue patterns at different scales. Moreover the l_2 -norm of the composite coefficients $C_i(\mathbf{x})$ is computed for each iteration as follows:

$$C_i(\mathbf{x}) = \sqrt{(G_x H_y)_i^2(\mathbf{x}) + (G_y H_x)_i^2(\mathbf{x})} \quad (2)$$

where $(G_x H_y)_i(\mathbf{x})$ and $(G_y H_x)_i(\mathbf{x})$ are the coefficients resulting from the convolution with the highpass filter on x and with the lowpass filter on y , and vice versa. The norm of both is computed because we believe that no directionality is contained in lung tissue textures. To extract higher-frequency features at smaller scales, the input images are upsampled by a factor of 2^n . The images are containing values in H.U.. The mean μ and the variance σ of the coefficients $G_i(\mathbf{x})$, $C_i(\mathbf{x})$

¹<http://rsb.info.nih.gov/ij/>

²<http://www.r-project.org/>

and $H_i(\mathbf{x})$ are computed over all ROIs for each iteration i (scale) to create the feature vector

$$\left(\begin{array}{cccc} \mu, \sigma(G_1(\mathbf{x}_R)) & \mu, \sigma(C_1(\mathbf{x}_R)) & \mu, \sigma(H_1(\mathbf{x}_R)) & \dots \\ \dots & \mu, \sigma(G_I(\mathbf{x}_R)) & \mu, \sigma(C_I(\mathbf{x}_R)) & \mu, \sigma(H_I(\mathbf{x}_R)) \end{array} \right)$$

where \mathbf{x}_R denotes the points belonging to the ROI.

2) *Grey-level histograms & specific features*: Since grey-level values of the pixels in HRCT images are expressed in H.U., each pixel value corresponds univoquely to densities of the anatomic organs and thus allows the identification of lung tissue components. In order to take advantage of this, histograms of pixel values are computed over each ROI. In addition, the percentage of air pixels (i.e. pixels with value < -1000 H.U.) is computed as an additional feature. Features related to grey-level values of HRCT images showed good discriminative properties for the separation of five lung tissue patterns in [20].

B. Classification algorithm

In order to quantify the discriminative properties of our texture features, we use a k -nearest neighbor (k -nn) classifier with euclidean distance computed between normalised feature vectors. Each feature is normalised using a linear mapping between 0 and 1 from each realisation. No weighting is used for the combination of heterogeneous features. A leave-one-out cross-validation is carried out to compute the classification accuracy. We use the function *knn.cv* of the package *class* of the free software *R*.

C. Classification results

We investigated the relevant scales of the DWF by performing a multi-class classification of the five patterns with a 1-nearest neighbor. Global accuracies using only coefficients from the first iteration of the DWF are shown in Figure 1. Since the original images are sampled by a factor 2^n , the sampling factor is inversely proportional to the scales contained in the patterns.




Fig. 1. Overall accuracies at iterative scales. The sampling factor 2^n is inversely proportional to the studied scale.

The optimal number of bins for computing histogram features is iteratively found by maximising the overall accuracy (Figure 2). Only histogram texture features are used.

In order to study which patterns are effectively discriminated by the combination of features from DWF (using the original image upsampled by 4 and 5 iterations) with histograms (40 bins) and percentage of air pixels, a confusion matrix is built (see Table II).




Fig. 2. Overall accuracies for varying number of bins, using the best number of nearest neighbors for each.

TABLE II
CONFUSION MATRIX OF COMBINED FEATURES

	healthy	emphysema	ground glass	fibrosis	micronodules
healthy	71 (92.2%)	0	1	0	5
emphysema	0	72 (100%)	0	0	0
ground glass	1	5	98 (86.7%)	0	9
fibrosis	0	0	4	60 (93.8%)	0
micronodules	6	0	1	4	144 (92.9%)

The comparison of accuracies of DWF, histograms plus percentage of air pixels, and the combination is illustrated in Figure 3.




Fig. 3. Comparison of accuracies of classification using histogram plus air percentage, discrete wavelet frames (DWF) and the combination of all features applied to each pattern versus all and to the multiclass configuration.

IV. INTERPRETATION

The results in Figure 1 show that smaller scales allow better discrimination of patterns. Indeed, in spite of the fact that upsampling images without carrying out interpolation introduces high-frequency artefacts, the best overall accuracy is reached with upsampled original images with a factor of 2. However, we believe that each pattern contains specific scales and the effect of the scale has to be studied for each of them. Higher accuracy can be reached by using the DWF with quincux wavelets where scales increase by a factor of $\sqrt{2}$ instead of 2 at each iteration [22]. Indeed, relevant information can be hidden between two consecutive scales. The optimal number of bins for grey-level histograms (40) is chosen as a trade-off between accuracy and dimensionality of the feature space. Using a large number of bins does not significantly improve the overall classification and increases the dimensionality of the feature space considerably. A non-linear binning can be effective as the majority of the information is contained in low H.U. values.

The confusion matrix in Table II shows how the combined features can separate each pattern. The 72 ROIs showing *emphysema*, which is characterised by the destruction of lung tissue (air), are 100% correctly classified. In this case, features such as percentage of air pixels are much more relevant compared to frequential analysis. *Ground glass* is the most difficult to classify and is confused 9 times with *micronodules*. This can be explained by the lack of intermediate scales with the DWF, where the size of the little nodules in *micronodules* is not correctly matched by the scaled templates. These observations seem to be validated by the respective accuracies in Figure 3 where DWF is not accurate to classify *emphysema* and *micronodules*. For patterns *ground glass* and *fibrosis*, the DWF shows superior discrimination performance. This can be explained by the fact that *fibrosis* is characterised by sharp transitions between high density tissue and small air bubbles.

The combination of histogram and DWF features shows improvement in classification accuracy for all patterns except *healthy*, and is particularly effective when carrying out multiclass classification.

We believe that using extra data associated with the images such as clinical parameters can improve the classification accuracy. Only very rarely can radiologists interpret images without taking into account the medical context. For example, a healthy lung of a 20-year-old and a 70-year-old have completely different characteristics. Moreover, 99 clinical parameters are available for each case in the database.

V. CONCLUSIONS AND FUTURE WORK

A. Conclusions

We used texture features to classify lung tissue patterns in HRCT data. The dataset used is in accordance with the clinical application, and the pattern recognition task is ready to be used in computer-aided diagnosis for ILDs. Texture features extracted from DWF and grey-level histograms show to be complementary and allow a multiclass classification of 5 patterns with an overall accuracy of 92.5%. The main weaknesses of the presented technique are the lack of resolution in scales with the DWF decomposition, along with required feature weighting while merging features from different origins.

B. Future Work

In order to avoid the introduction of high frequencies while upsampling the images, we plan to use interpolation (e.g. Gaussian blurring or spline interpolation). A possible solution to the lack of scales with the DWF is to use the redundant quincunx wavelet transform, where each successive lower resolution level is scaled by a factor of $\sqrt{2}$ [22], or to convolve the image with given templates at selected scales (e.g. convolution with scaled Gabor functions). In order to combine visual features with clinical parameters, we plan to use support vectors machines (SVM). Finally, an extension of 2D texture features to 3D is planned in order to fully exploit information in 3D HRCT data, which showed significant improvement in classification accuracy in [17].

REFERENCES

- [1] T. E. King, "Approach to the adult with interstitial lung disease," *UpToDate*, vol. August, 2004.
- [2] H. Müller, N. Michoux, D. Bandon, and A. Geissbuhler, "A review of content-based image retrieval systems in medicine – clinical benefits and future directions," *International Journal of Medical Informatics*, vol. 73, pp. 1–23, 2004.
- [3] P. Stark, "High resolution computed tomography of the lungs," *UpToDate*, vol. August, 2004.
- [4] R. M. Haralick, "Statistical and structural approaches to texture," *Proceedings of the IEEE*, vol. 67, no. 5, May 1979.
- [5] T. R. Reed and J. M. H. du Buf, "A review of recent texture segmentation and feature extraction techniques," *Computer Vision, Graphics and Image Processing*, vol. 57, no. 3, pp. 359–372, 1993.
- [6] M. Unser, "Sum and difference histograms for texture classification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 1, pp. 118–125, 1986.
- [7] T. Hsu, A. D. Calway, and R. Wilson, "Texture analysis using the multiresolution fourier transform," Bristol, UK, Tech. Rep., 1993.
- [8] S. Arivazhagan and L. Ganesan, "Texture classification using wavelet transform," *Pattern Recogn. Lett.*, vol. 24, no. 9–10, pp. 1513–1521, 2003.
- [9] M. Unser and A. Aldroubi, "A review of wavelets in biomedical applications," *Proceedings of the IEEE*, vol. 84, no. 4, pp. 626–638, April 1996.
- [10] M. Fernández and A. Mavilio, "Texture analysis of medical images using the wavelet transform," in *AIP Conference Proceedings*, vol. 630, October 2002, pp. 164–168.
- [11] M. Unser, "Texture classification and segmentation using wavelet frames," *IEEE Transactions on Image Processing*, vol. 4, no. 11, pp. 1549–1560, 1995.
- [12] A. Van Nevel, "Texture classification using wavelet frame decompositions," *Proceedings of 31st IEEE Asilomar Conference on Signals, Systems and Computers*, vol. 1, no. 11, pp. 311–314, November 1997.
- [13] 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," *Computer Vision and Image Understanding (special issue on content-based access for image and video libraries)*, vol. 75, no. 1/2, pp. 111–132, August 1999.
- [14] A. M. Aisen, L. S. Broderick, H. Winer-Muram, C. E. Brodley, A. C. Kak, C. Pavlopoulou, J. Dy, C.-R. Shyu, and A. Marchiori, "Automated storage and retrieval of thin-section CT images to assist diagnosis: System description and preliminary assessment," *Radiology*, vol. 228, pp. 265–270, 2003.
- [15] V. Zavaletta, B. J. Bartholmai, and R. A. Robb, "Nonlinear histogram binning for quantitative analysis of lung tissue fibrosis in high-resolution CT data," in *SPIE Medical Imaging*, vol. 6511, San Diego, CA, USA, February 2007.
- [16] R. Uppaluri, E. A. Hoffman, M. Sonka, P. G. Hartley, G. W. Hunninghake, and G. McLennan, "Computer recognition of regional lung disease patterns," *American Journal of Respiratory and Critical Care Medicine*, vol. 160, no. 2, pp. 648–654, August 1999.
- [17] Y. Xu, M. Sonka, G. McLennan, J. Guo, and E. A. Hoffman, "MDCT-based 3-D texture classification of emphysema and early smoking related lung pathologies," *IEEE Transactions on Medical Imaging*, vol. 25, no. 4, pp. 464–475, April 2006.
- [18] A. Shamsheyeva and A. Sowmya, "The anisotropic gaussian kernel for svm classification of hrct images of the lung," in *Proceedings of the 2004 Intelligent Sensors, Sensor Networks and Information Processing Conference*, December 2004, pp. 439–444.
- [19] A. Depeursinge, H. Müller, A. Hidki, P.-A. Poletti, T. Rochat, and A. Geissbuhler, "Building a library of annotated pulmonary CT cases for diagnostic aid," in *Swiss conference on medical informatics (SSIM 2006)*, Basel, Switzerland, April 2006.
- [20] 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 *SPIE Medical Imaging*, vol. 6514, San Diego, CA, USA, February 2007.
- [21] M. Unser and T. Blu, "Wavelet theory demystified," *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 470–483, February 2003.
- [22] D. Van De Ville, T. Blu, and M. Unser, "Isotropic polyharmonic B-Splines: Scaling functions and wavelets," *IEEE Transactions on Image Processing*, vol. 14, no. 11, pp. 1798–1813, November 2005.