1. INTRODUCTION

1.1. ARTIFICIAL INTELLIGENCE AND MEDICINE[2-1]

Medicine is a field in which computer help is critically needed. Our increasing expectations of the highest quality health care and the rapid growth of ever more detailed medical knowledge leave the physician without adequate time to devote to each case and struggling to keep up with the newest developments in his field. For lack of time, most medical decisions must be based on rapid judgments of the case relying on the physician's unaided memory. Only in rare situations can a literature search or other extended investigation be undertaken to assure the doctor (and the patient) that the latest knowledge is brought to bear on any particular case. Continued training procedures encourage the physician to keep more of the relevant information constantly in mind, but fundamental limitations of human memory and recall coupled with the growth of knowledge assure that most of what is known cannot be known by most individuals. This is the opportunity for new computer tools: to help organize, store, and retrieve appropriate medical knowledge needed by the practitioner in dealing with each difficult case, and to suggest appropriate diagnostic and therapeutic decisions and decision making techniques.

In a 1970 review article, Schwartz speaks of the possibility that the computer as an intellectual tool could reshape the system of health care, fundamentally alter the role of the physician, and profoundly change the nature of medical manpower recruitment and medical education.

The key technical developments leading to this reshaping will almost certainly involve exploitation of the computer as an 'intellectual' and 'deductive' instrument, a consultant that is built into the very structure of the medical care[1-1].

Artificial Intelligence in Medicine (AIM) is now slowly taking up the challenge of creating and distributing the necessary tools to accomplish the above mentioned tasks. One introductory textbook defines artificial intelligence (AI) as "the study of ideas which enable computers to do the things that make people seem intelligent ... The central goals of Artificial Intelligence are to make computers more useful and to understand the principles which make intelligence possible"[1-2].

AI in Medicine (AIM) is AI specialized to medical applications. Researchers in AIM employ human-like reasoning methods in the programs, justifying that choice either as a commitment to a human-computer equivalence sought by some or as a good engineering technique for capturing the best understood source of existing expertise on medicine, which is the practice of human experts. Most researchers adopt the latter view.

Relying on the knowledge of human experts to build expert computer programs is actually helpful for several additional reasons: First, the decisions and recommendations of a program can be explained to its users and evaluators in terms that are familiar to the experts. Second, because we hope to duplicate the expertise of human specialists, we can measure the extent to which our goal is achieved by a direct comparison of the program's behaviour to that of the experts. Finally, within the collaborative group of computer scientists and physicians engaged in AIM research, basing the logic of the programs on human models supports each of the three somewhat disparate goals that the researchers may hold:

- To develop expert computer programs for clinical use, making possible the inexpensive dissemination of the best medical expertise to geographical regions where that expertise is lacking, and making consultation help available to non-specialists who are not within easy reach of expert human consultants.

- To formalize medical expertise, to enable physicians to understand better what they know and to give them a systematic structure for teaching their expertise to medical students.

- To test AI theories in a "real world" domain and to use that domain to suggest novel problems for further AI research.

1.2. MOTIVATIONS AND OBJETIVES

The motivations for this project, as previously mentioned above, are mainly due to the complication of the diagnostic process, since it requires integration of information about a patient's anatomy, physiology and medical history. The process of diagnosis begins with a physical exam and patient history. In many cases, this step is inconclusive and some form of medical imaging is required to confirm or exclude a diagnosis. Using medical imaging, physicians are able to have qualitative and quantitative information about the anatomy and physiology of the patient. It allows physicians to determine the location and extent of potential abnormalities. With these advantages, medical imaging has become central to medical diagnosis.

Combining digital imaging and computer processing capabilities, it has become possible to approach the problem of automating diagnosis in medical imaging, which will be the main objective of the current project. From this point of view, diagnosis can be thought of as a series of detection and classification tasks. First, an abnormality must be detected. Next, the abnormality must be localized and, finally, the abnormality must be classified and any physical properties must be quantified.

This research focuses on the development of an artificial intelligence tool to assist physicians in the interpretation of CT lung scans, for the diagnosis of abnormalities in the lung tissue. The detection of unhealthy regions from CT scans is approached in two steps. First, texture analysis is performed on the lung images, and a subset of textural features extracted. Second, an artificial neural network is trained for predicting the presence or absence of disease based on the textural parameters extracted in the first step. Therefore, the main goal of this project is to prove that there is a relationship between the texture characteristic of CT lung scans and the presence of disease.

This research is organized as follows. The first two chapters provide background information. Chapter 2 focuses on texture analysis, reviewing the literature about texture models and texture analysis problems, as well as its applications in medical image processing, and specifically in the detection of lung diseases. Chapter 3 outlines the normal lung function and structure, and the clinical problems of lung

diseases and their diagnosis, narrowing the focus to emphysema and fibrosis, which will be the most interesting ones for this investigation. Chapter 4 introduces the program developed for the automatic detection of unhealthy tissues, and the chosen language and environment. Chapter 5 provides the techniques used for the feature selection and attempts to validate them in some CT samples images of emphysematous and fibrous lungs. Here we divide the textural features into four categories: statistical, co-occurrence matrix, Fourier, and binary features. In Chapter 6, artificial intelligence techniques (neural networks) are reviewed and applied to the diagnosis of unhealthy lungs, and they are combined with texture analysis to develop the final tool for automatic diagnosis. Chapter 7 explores an approach to a 3D diagnosis on whole data-sets of lung CT scans, comparing the new results with the previous ones in 2D, and discussing the advantages and disadvantages of both of them. Finally, Chapter 8 summarizes the conclusions reached in the dissertation, as well as suggesting future directions of research.

2. TEXTURE ANALYSIS[1-3, 1-4, 2-2]

2.1. INTRODUCTION

Texture is an important characteristic for the analysis and processing of various types of images, including natural scenes, remotely sensed data and biomedical modalities. It is believed that the texture plays an important role in the visual systems for recognition and interpretation of data. Texture analysis is a significant research field in computer vision, image processing and pattern recognition. A number of techniques have been developed for texture feature extraction, segmentation, classification and synthesis[1-5].

Texture analysis is also relevant in the characterization of images that do not exhibit a clear boundary between different objects within the images. This can be observed in CT images of the lungs, especially if abnormalities are present. For this reason, texture analysis plays a crucial role in this research.

In reviewing texture features for use in classification and discrimination schemes this Chapter has three main objectives:

- To identify research on the different texture models and texture analysis problems

- To review some previous works about texture features applied to medical images, and specifically to lung images

- To select a set of feature measures for further investigation about changes in parenchyma density, i.e., abnormally bright or dark areas in lung images

However, it is not practical to provide an exhaustive survey of all texture measures here. We will review though some of the more popular techniques. Concerning the selection of features, the criteria will be popularity in the literature, ease of implementation and use, efficiency and simplicity. Hence some of the following techniques will not be directly used in this project; however they are named to illustrate further possibilities of research. Later on Chapter 5, a more accurate selection is developed, and some of the features are further outlined.

2.2. TEXTURE MODELS

Identifying the perceived qualities of texture in an image is an important first step towards building mathematical models for texture. Modelling the intensity variations in an image that characterize texture is very difficult, so it is usually characterized by the two-dimensional variations in the intensities present in the image.

The "definition" of texture is formulated by different people depending upon the particular application and there is no generally agreed upon definition. Some are perceptually motivated, and others are driven completely by the application in which the definition will be used. Coggins[1-6] has compiled a catalogue of texture definitions in the computer vision literature. Here we outline some properties of texture which are generally assumed to be true:

- Texture is a property of areas, so texture is a contextual property and its definition must involve grey values in a spatial neighbourhood.

- Texture involves the spatial distribution of grey levels. Thus, the two-dimensional histograms or co-occurrence matrices are reasonable texture analysis tools.

- Texture in an image can be perceived at different scales or levels of resolution

- A texture is perceived when significant individual forms are not present.

Figure 2.1 shows an image that can be segmented into five homogeneous textured regions. The purpose of image segmentation is to divide the input image into homogeneously textured regions, without knowing a priori what the textures are. Texture classification methods, on the other hand, attempt to assign a known texture class to each image region[1-4].

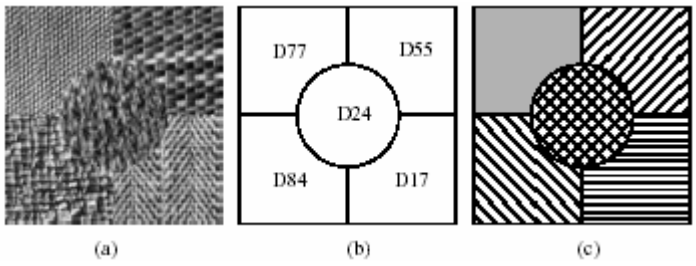

FIGURE 2.1. (a) An image consisting of five different textured regions: cotton canvas (D77), straw matting (D55), raffia (D84), herringbone weave (D17), and pressed calf leather. [8]. (b) The goal of texture classification is to label each textured region with the proper category label: the identities of the five texture regions present in (a). (c) The goal of texture segmentation is to separate the regions in the image which have different textures and identify the boundaries between them. The texture categories themselves need not be recognized. In this example, the five texture categories in (a) are identified as separate textures by the use of generic category labels (represented by the different fill patterns).

These qualities play an important role in describing texture: uniformity, density, coarseness, roughness, regularity, linearity, directionality, direction, frequency and phase. The fact that the perception of texture has so many different dimensions is an important reason why there is no single method of texture representation that is adequate for a variety of textures.

2.2.1. SURVEYS

Haralick provided the classic survey of texture measures[1-7]. He listed and described a number of texture extraction methods which he divided into two types: structural and statistical. The former used primitives to describe texture elements and placement rules to describe the spatial relationship between elements. This approach is better suited to textures that exhibit a regular macro-structure, and will not be discussed further. The statistical approaches are better suited to micro textures, and Haralick identified techniques based upon auto correlation functions, frequency domain analysis, edge operators, grey-level co-occurrence matrices, grey-level run lengths, and autoregressive models. In the last years, there has been an explosion of interest in model-based techniques (Markov fields, fractals etc.), detailed in a survey in 1993 by Reed and Buf[1-22]. Here we will make a differentiation between two main groups: model-based and non-model-based features.

2.2.2. MODEL-BASED FEATURES

A number of random field models (i.e. models of two-dimensional random processes) have been used for modelling and synthesis of texture. If a model is shown to be capable of representing and synthesising a range of textures, then its parameters may provide a suitable feature set for classification and/or segmentation of the textures. For a model based approach to be successful, there must exist a reasonably efficient and appropriate parameter estimation scheme, and the model itself should be parsimonious, i.e. use the minimum number of parameters. Popular random field models used for texture analysis include fractals, autoregressive models and Markov random fields. Only fractal models will now briefly be reviewed, because we will not use the others in this project. A more extensive review of these approaches may be found in the survey by Ahuja and Schachter[1-8].

2.2.2.1. FRACTAL MODELS

Many natural surfaces have a statistical quality of roughness and self-similarity at different scales. Fractal geometry has proven to be a useful tool in quantifying the structure of a wide range of idealized and naturally occurring objects, from pure mathematics, through physics and chemistry, to biology and medicine. In the past few years fractal analysis techniques have gained increasing attention in signal and image processing, especially in medical sciences, e.g. in pathology, neuropsychiatry, cardiology. Mandelbrot[1-9] proposed fractal geometry and is the first one to notice its existence in the natural world.

We first define a deterministic fractal in order to introduce some of the fundamental concepts. Self-similarity across scales in fractal geometry is a crucial concept. A deterministic fractal is defined using this concept of self-similarity as follows. Given a bounded set A in a Euclidean n-space, the set A is said to be self-similar when A is the union of N distinct (non-overlapping) copies of itself, each of which has been scaled down by a ratio of r. The fractal dimension D is related to the number N and the ratio r as follows:

$$D = \frac{\log N}{\log(1/r)} \tag{2.1.}$$

The fractal dimension gives a measure of the roughness of a surface. Intuitively, the larger the fractal dimension, the rougher the texture is. There are a number of methods proposed for estimating the fractal dimension *D*. One method is the estimation of the box dimension as follows. Given a bounded set *A* in Euclidean n-space, consider boxes of size on a side that cover the set *A*. A scaled down version of the set by ratio *r*, will result in similar sets. This new set can be covered by boxes of size *L=rLmax.* The number of such boxes then is related to the fractal dimension by

$$N(L) = \frac{1}{r^D} = \left[\frac{L\max}{L}\right]^D \tag{2.2}$$

The fractal dimension is then estimated from Equation (2.2) by the following procedure. For a given *L*, divide the n-space into a grid of boxes of size *L* and count the number of boxes covering *A*. Repeat this procedure for different values of *L*. Then estimate the value of the fractal dimension *D* from the slope of the line

$$\ln(N(L)) = -D\ln(L) + D\ln(L\max) \tag{2.3}$$

This can be accomplished by computing the least squares linear fit to the data, namely, a plot of ln(*L*) vs. –ln(*N*(*L*)) .

Fractal geometry can be applied in medical image, since fractals can be found within the human body. For example, the pulmonary system is composed of tubes that lead to air sacs called alveoli. The main tube is the trachea, which splits into the two smaller tubes that lead to the separate lungs, the bronchi. The tubes keep splitting until the smallest tubes; the bronchioles lead into the alveoli. This is a similar description to that of fractal canopies, as shown in Figure 2.2.

More evidence to support the idea that the lungs are fractal comes from measuring the alveolar area. When they are lightly magnified this area was found to be $80m^2$, but when magnified with an electron microscope it was found to be $140m^2$. We know that the increase in size with magnification is a property of fractals. So the fractal dimension can be used as a tool for the detection of structure changes and quantification of lung diseases.

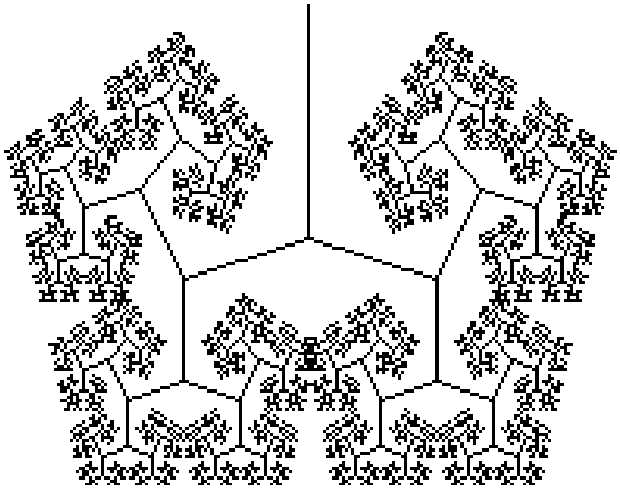

Figure 2.2. A fractal canopy is an approximation of a lung

2.2.3. NON-MODEL-BASED FEATURES

This section briefly reviews Co-occurrence matrices and other related Statistical Features, and Frequency Domain Methods.

2.2.3.1. CO-OCCURRENCE MATRICES AND OTHER STATISTICAL FEATURES

One of the defining qualities of texture is the spatial distribution of grey values. The use of statistical features is therefore one of the early methods proposed in the machine vision literature. A large number of features have been proposed but they are

not independent. There is a relationship between the various statistical texture measures and the input image. Among these features there are some simple ones that are highly popular due to their simplicity and efficiency, such as the *Mean Grey Value, Maximum and Minimum grey values, Range, Standard Deviation,* and *Percentiles,* and other complex ones that will be reviewed here, such *as Co-occurrence matrix* features and *Autocorrelation* features.

The Spatial *Grey Level Dependence Matrix (SGLDM)* describes the spatial distribution and spatial dependence among the grey tones in a local area based on the estimation of second order conditional probability density functions $f(i,j,d,\theta)$. Each of these functions are the probability of going from grey level i to grey level j separated by a distance d and aligned to the angle θ. The estimated values (which are grey-tone spatial-dependence frequencies) can be written in matrix form, the so-called *Co-occurrence matrices*.

Haralick proposed 28 features extracted from 14 equations, but usually only six of them are used. Many researchers have used Haralick's co-occurrence based features. The most popular features include *Entropy, Energy, Inverse Difference Moment, Maximum Probability, Contrast* and *Correlation*, with small displacement vectors e.g. (1,0) and (0.1).

An important property of many textures is the repetitive nature of the placement of texture elements in the image. The *Autocorrelation Function* of an image can be used to see the amount of regularity as well as the fineness and coarseness of the texture present in the image. Formally, the *Autocorrelation Function* of an image $I(x,y)$ is:

$$r(x,y) = \frac{\sum_{u=0}^{N-1}\sum_{v=0}^{N-1} I(u,v)I(u+x,v+y)}{\sum_{u=0}^{N-1}\sum_{v=0}^{N-1} I^2(u,v)} \tag{2.4}$$

The autocorrelation is related to the size of the texture primitive (i.e., the fineness of the texture). If the texture is coarse, then the autocorrelation will drop off slowly. It is also related to the power spectrum of the Fourier transform in the space domain.

2.2.3.2. FREQUENCY DOMAIN METHODS

The frequency analysis of the textured images is best done in the Fourier domain. The human visual system analyses the textured images by decomposing the image into its frequency and orientation components.

Two-dimensional power or magnitude spectra provide information on texture coarseness and directionality from their radial and angular distributions respectively. The most commonly extracted features consist of sums of coefficients within wedges, rings, or sectors of two-dimensional power spectra. However, results derived from tests on a Brodazt set are quite disappointing.

Other frequency domain measures include those derived from the characteristics of spectral peaks. D'Astous and Jernigan (1984) used features that included the frequency, direction, area and relative power of spectral peaks. Other measures contain the *Maximum Peak* of the spectrum, the *Average* value and the *Energy*.

2.3. TEXTURE ANALYSIS PROBLEMS

The various methods for modeling textures and extracting texture features can be applied in the following broad categories of problems: texture segmentation, texture classification, texture synthesis and shape from texture. Only the two first will be described here, because they are the only ones with interest for our study.

- *Texture segmentation* is used to refer to the process of dividing an image up into homogeneous regions according to some homogeneity criteria. This is a difficult problem because one usually does not know a priori what types of textures exist in an image, how many different textures there are, and what regions in the image have which textures. In fact, one does not need to know which specific textures exist in the image in order to do texture segmentation. All that is needed is to tell that two textures (usually in adjacent regions of the images) are different. It is therefore intimately concerned with establishing the boundaries between these regions without regard to the type of class of the regions.

  The two general approaches to performing texture segmentation are analogous to methods for image segmentation: region-based approaches or boundary-based approaches. In a region-based approach, one tries to identify regions of the image that have a uniform texture, and in a boundary-based approach one tries to detect the differences in texture in adjacent regions.

- *Texture classification* involves deciding what texture category an observed image belongs to. In order to accomplish this, one needs to have an a priori knowledge of the classes to be recognized. Once this knowledge is available and the texture features are extracted, one then uses classical pattern classification techniques in order to do the classification. If the classes have not been defined a priori, the task is referred to as unsupervised classification. Alternatively, if the classes have already been defined (normally through the use of training sets of sample textures) then the process is referred to as supervised classification, and this is the kind of classification we will work

with in the current research. In Figure 2.1 the classification and segmentation problems are explained and clearly differentiated.

Before either segmentation or classification can take place, some homogeneity or similarity criterion must be defined. These criteria are normally specified in terms of a set of feature measures, which each provide a quantitative measure of a certain texture characteristic. These feature measures are alternatively referred to as textures measures or just simply features. Groups of features assembled for segmentation or classification purposes are often referred to as feature vectors.

2.4. APPLICATIONS IN MEDICAL IMAGE PROCESSING

Texture analysis methods have been utilized in a variety of applications domains, such as remote sensing, surface inspection, document processing, and medical image processing. Image analysis techniques have played an important role in several medical applications. In general, the applications involve the automatic extraction of features from the image that are then used for a variety of classification tasks, such as distinguishing normal tissue from abnormal tissue. Depending upon the particular classification task, the extracted features capture morphological properties, color properties, or certain textural properties of the image. The textural properties computed are closely related to the application domain to be used. We present here some examples of medical applications, and later we will focus on some applications in the detection of lung diseases.

Harms et al.[1-10] used image texture in combination with color features to diagnose leukemic malignancy in samples of stained blood cells. They extracted texture micro-edges and "textons" between these micro-edges. The textons were regions with almost uniform color. They extracted a number of texture features from the textons including the total number of pixels in the textons that have a specific color, the mean texton radius and texton size for each color and various texton shape features. In combination with color, the texture features significantly improved the correct classification rate of blood cell types compared to using only color features.

Landeweerd and Gelsema[1-11] extracted various first-order statistics (such as mean gray level in a region) as well as second-order statistics (such as gray level co-occurrence matrices) to differentiate different types of white blood cells.

Lundervold[1-12] used fractal texture features in combination with other features (such as response to edge detector operators) to analyze ultrasound images of the heart.

2.4.1. APPLICATIONS IN THE DETECTION OF LUNG DISEASES

There are many studies about automatic detection of lung diseases using texture features, since in the last years applying computers to medical image processing has become essential to medical diagnosis. In this section we present several previous investigations, stating the features used in each one, as well as the results of interest.

Sutton and Hall[1-13] discussed the classification of pulmonary disease using texture features. Some diseases, such as interstitial fibrosis, affect the lungs in such a manner that the resulting changes in the X-ray images are texture changes as opposed to clearly delineated lesions. In such applications, texture analysis methods are ideally suited for these images. Sutton and Hall proposed the use of three types of texture features to distinguish normal lungs from diseased lungs. These features are computed based on an isotropic contrast measure, a directional contrast measure, and a Fourier domain energy sampling. In their classification experiments, the best classification results were obtained using the directional contrast measure.

Douglas *et al.*[2-3] outline the importance of fractal analysis in lung images as a useful method for assisting in the diagnostic interpretation of perfusion lung scans. This work is based on the hypothesis that the fractal dimension can quantify the difference between normally and abnormally perfused lung scan areas. Each lung was divided into three zones (upper, middle, lower) and a small region of interest (ROI) was selected for each lung zone. The ROI approach was chosen to determine the utility of fractal analysis to characterize each disease process individually, and was placed to correspond to normal lung or to the abnormality in the lung zone. Each ROI was characterized as either normal or abnormal with abnormal divided into PE (pulmonary embolism), OPD (obstructive pulmonary disease), OPAC (parenchymal opacity), EFFU (pleural effusion), ATEL (atelectasis), and OTHER. The results show that the average fractal dimension of normal lung ROIs was significantly higher than that of abnormal ones. As an example, the average fractal dimension of abnormal regions with pulmonary embolism is much lower than the fractal dimension of regions with chronic obstructive pulmonary disease present.

Kelly and Cannon[2-4] developed a method for calculating the similarity between two digital images, and applied this algorithm to the problem of search and retrieval for a database containing pulmonary CT imaginery. A global signature describing the texture, shape, or color content is first computed for every image stored in a database, and a normalized distance between probability density functions of feature vectors is used to match signatures. When the database retrieval software is asked to search for images similar to a given target image, it first computes the global signature of the target image and then matches it against the signatures of all images in the database. A handful of images having similar content, i.e. database images having a similar signature to the target image, is returned to the user. A normalized distance between probability density functions of feature vectors is used to match signatures. The general idea is that first several features (local colour, texture, and/or shape) at every pixel in the image are computed, and then a histogram of feature vector (pixel vector) occurrences for that image is made. The features selected for this problem were texture energy measures, which have the advantage of being able to discriminate between different textures, while being quick and easy to compute. Using only four texture features, the system successfully discriminates between different pulmonary diseases, returning images with the same content and resolution as the target images.

Sonka *et al*.[2-5] worked on a texture-based tissue characterization method based upon the training acquired on a set of representative examples. The AFMF has been applied to several different discrimination tasks including normal subjects, subjects with interstitial lung disease, smokers, asbestos-exposed subjects, and subjects with cystic fibrosis, as well as for the analysis of pulmonary parenchyma from X-Ray CT. It basically consists of four steps: preprocessing of the images, feature extraction, optimal feature selection, and classification. The feature computation involved extracting measures based on grey level distribution, percentiles, run length matrices, co-occurrence matrices, geometric fractal dimension and stochastic fractal dimension. The optimal feature selection was performed using the divergence method along with correlation analysis. The classification was performed using a Bayesian approach. The lung diseases under study were emphysema, interstitial lung disease, asbestosis and cystic fibrosis.

2.5. SUMMARY AND CONCLUSIONS

In this Chapter we have developed a review of the most frequent texture analysis techniques and some of their main applications. The textural features that we will finally choose for our project in Chapter 5 are extracted from the models described here.

Other features and models that will not be used in the current investigation have also been briefly described, in order to have a more general view of the complexity of the problem.

3. LUNG FUNCTION AND PATHOLOGY

3.1. INTRODUCTION

A sound knowledge of the physiology and pathophysiology of the lung will be necessary in diagnosis. Therefore this Chapter provides an overall view of the respiratory function and structure of the lungs, in order to identify the syndromes of abnormalities and disorders that are indicated, precisely, by the changes in structure. Some diseases are reviewed, indicating their cause and diagnosis, and narrowing the view to the most relevant ones for our research, emphysema and fibrosis.

The current work is based mainly on extracting measures from CT lung scans. Hence, further description of CT imaginery is outlined, providing details about general procedures for extracting these images, benefits and risks of its use and a comparison between CT techniques and Magnetic Resonance Imaging (MRI).

In such images, the unhealthy regions are presented with abnormally bright or dark grey values, which represent the main characteristic used to detect diseases. In fact, throughout some subsequent cases and diagnostics our goal will consist of detecting these areas as ‘abnormal’ without regard to the kind of abnormality, leaving this task to expert physicians, who can make a more accurate diagnosis. The point of this is that if this research is successful, physicians will save a considerable amount of time, since they will only have to look into the images with ‘unhealthy’ regions, and decide about the type of disorder.

3.2. LUNG FUNCTION[1-14]

The *respiratory system* is composed of the lungs, the conducting airways, the parts of the central nervous system concerned with the control of the muscles of respiration, and the chest wall. The chest wall consists of the muscles of respiration (the diaphragm, the intercostal muscles, and the abdominal muscles), and the rib cage.

Its main function is the exchange of carbon dioxide for oxygen that takes place in the lungs. Fresh air, containing oxygen, is inspired into the lungs through the conducting airways. The forces needed to cause the air to flow are generated by the respiratory muscles, acting on commands initiated by the central nervous system. At the same time, venous blood returning from the various body tissues is pumped into the lungs by the right ventricle of the heart. This mixed venous blood has high carbon dioxide content and low oxygen content. In the pulmonary capillaries, carbon dioxide is exchanged for oxygen.

3.2.1. THE MUSCLES OF RESPIRATION AND THE CHEST WALL

The muscles of respiration and the chest wall are essential components of the respiratory system. The lungs are not capable of inflating themselves; the muscles of respiration must supply the force for this inflation. The chest wall must be intact and able to expand if air is to enter the alveoli normally. The primary components of the chest wall are shown schematically in figure 3.1 (left). These include the rib cage, the external and internal intercostal muscles and the diaphragm, which are the main muscles of respiration involved in the processes of inspiration and expiration.

The normal inspiration and expiration processes are altered by the presence of disease, as we can see in figure 3.1 (centre), where static pressure-volume curves corresponding to the expiratory process are presented. *Emphysema*, as described later in further detail, increases the compliance of the lungs because it destroys the alveolar septal tissue that normally opposes lung expansion. On the other hand, *fibrosis* makes the lungs less compliant, or stiffer, and increases alveolar elastic recoil.

3.2.2. THE AIRWAYS

After passing through the nose or mouth, the pharynx, and the larynx (*the upper airways*), air enters the tracheobronchial tree. Starting with the trachea, the air may pass through as few as 10 or as many as 23 generations, or branchings, on its way to the alveoli. The first 16 generations of airways, the *conducting zone*, contain no alveoli and thus are anatomically incapable of gas exchange with the venous blood. They constitute the *anatomic dead space*. Alveoli start to appear at the seventeenth through the nineteenth generations, in the respiratory bronchioles, which constitute the *transitional zone*. The twentieth to twenty-second generations are lined with alveoli. These *alveolar ducts* and the *alveolar sacs*, which terminate the tracheobronchial tree, are referred to as the *respiratory zone*.

The structural architecture of the airways is composed of cartilage, smooth muscle and fibrous tissue, and this structure is shaped like *natural* or *random fractals* (structures that grow with an element of chance and over a range of magnifications have the same fractal dimension). It varies considerably, depending on their location in the tracheobronchial tree. The trachea is a fibromuscular tube supported ventrolaterally by C-shaped cartilage and completed dorsally by smooth muscle. The cartilage of the large bronchi is semicircular, like that of the trachea, but as the bronchi enter the lungs, the cartilage rings disappear and are replaced by irregularly shaped cartilage plates. By definition, airways with no cartilage are termed *bronchioles*. Because the bronchioles and alveolar ducts contain no cartilage support, they are subject to collapse when compressed.

3.2.3. THE ALVEOLAR-CAPILLARY UNIT

This unit is the site of gas exchange in the lung. The alveoli, estimated to number about 300 million, are almost completely enveloped in pulmonary capillaries. There may be as many as 280 billion pulmonary capillaries. The result of these staggering numbers of alveoli and pulmonary capillaries is a vast area of contact between alveoli and pulmonary capillaries, probably 50 to 100 m^2 of surface area available for gas exchange by diffusion.

A dense fibrous network of collagen and elastin in the interstitium that surrounds the airways provides the architectural struts needed to support the large gas exchange surface. Fibers originated in the visceral pleura radiate inward to invest each lobe and subdivide the lobes into lung segments. Elastic fibers contribute to the fibrous elements in which a contractile property is added to the structure. Alterations in this fiber network such as overgrowth (*pulmonary fibrosis*) or destruction (*emphysema*) can cause severe dysfunction.

The acinus (figure 3.1 right) is the basic gas exchange unit of the lung and consists of those structures distal to the terminal bronchiole. The acinus is a descending series of branches and includes two to five orders of respiratory bronchioles, two to five orders of alveolar ducts, one to three alveolar sacs, and alveoli.

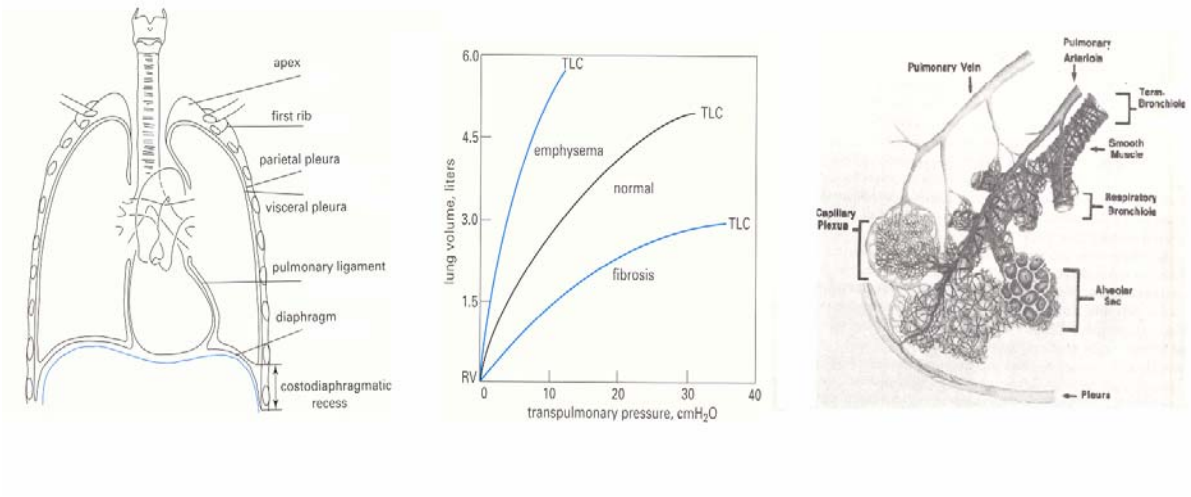

Figure 3.1. Lung structure and function.
Left: repiratory muscles. Centre: Compliance curves comparing normal/abnormal states. Right: Acinus

3.3. LUNG DISEASES[1-16,1-17]

3.3.1. OVERALL VIEW

In taking medical history of a patient with pulmonary problems, the physician should note the primary symptoms of pulmonary disease and also any symptoms arising from other organ system that may relate to pulmonary disorders. These primary symptoms are cough, expectoration, dyspnea, wheezing and chest pain. Key points in the patient's past medical history include smoking, allergies, and occupational exposure to potential disease-producing environments. The presence of a particular syndrome constitutes a functional diagnosis. This will reflect, and in some instances indicate precisely, the underlying changes in structure that are present in CT scans.

Diseases of the airways can be localised, e.g. a tumour or polyp, or generalised. Generalised conditions of the airways can present with cough and expectoration of phlegm, with wheeze which may be episodic, or with breathlessness on exertion. The classical causes are then chronic bronchitis, asthma and emphysema. However the conditions may coexist and all can be associated with airflow limitation so the distinction between them is not always clear cut. This has led to use of the terms CNLD (chronic non-specific lung disease), COLD (chronic obstruction lung disease), COPD (chronic obstructive pulmonary disease) and CAO (chronic airways obstruction). CNLD describes chronic airflow limitation associated with chronic bronchitis, emphysema or asthma and carries the implication that one or more host factors may contribute to the flow limitation. COLD describes airflow limitation which is progressive, mainly irreversible and yet is clearly not due to bronchiolitis or asthma of defined aetiology. COPD is nearly interchangeable with COLD but can also describe emphysema without airflow limitation. CAO describes airflow limitation associated with partial bronchial obstruction and should strictly exclude flow limitation due to a reduced elastic recoil pressure as occurs in emphysema. However, the terms are of limited usefulness because they are seldom used in a precise sense.

In the following section, we will make a division of diseases into those that present an 'above normal' parenchymal density, and those that present a 'below normal' parenchymal density, because this is the splitting that best suits the considerations made in this project.

3.3.2. CLASSIFICATION

The main disorder under study presenting **abnormally low greyscale values** in the lung tissue is *emphysema*. Emphysema (figure 3.2, top left) is a condition of the lung characterised by an increase beyond the normal in the size of air spaces distal to the terminal bronchioles, and destruction of some of the air sacs and the walls of the airways. The expansion commonly affects the second order of respiratory bronchioles. That due to destruction of the walls of the airways is known as centriacinar (or centrilobular) emphysema. When it is due to dilatation which is secondary to the accumulation of inert dust in the lung it is called focal emphysema. Compensatory and senile types of emphysema also occur. Panacinar emphysema refers to the condition in which there is destruction and expansion of more than one order of airway within the acinus, i.e. that portion of the lung which is distal to each terminal bronchiole. The condition can then arise in the air sacs and alveolar ducts. Bullous emphysema may be said to be present when, in the inflated lung, the diameter of one or more of the emphysematous spaces exceeds 1 cm. Emphysema and chronic bronchitis together comprise COPD. The tobacco smoke is the main cause of this disorder. In some occasions, a little fibrosis or structural narrowing of small airways can appear together with emphysema.

The clinical features of emphysema are relatively independent of the underlying pathology, except in the case of large bullae when the symptoms and signs of a space-occupying lesion can be superimposed. The clinical history is of progressively increasing breathlessness, whilst cough and expectoration can be mild and of late onset. Wheeze is common and on clinical examination the chest is held in an inflated position. The chest radiograph typically shows over-inflation and loss of lung tissue.

At an early stage of the disease, the diagnosis can be made by computer-assisted tomography (CT), as showed in figure 3.2. This can identify both bullae and areas of reduced density where lung tissue has been destroyed. The areas of low density can be highlighted in the CT display.

The physiological features of diffuse emphysema are secondary to the loss of lung tissue. This affects both the elastic recoil of the lung and the surface area which is available for the exchange of gas. The elastic recoil is reduced to its greatest extent in panacinar emphysema whilst the loss of surface occurs equally in centriacinar and panacinar emphysema.

The diagnosis of emphysema is suspected on the basis of the clinical features and the chest radiograph. The diagnosis is confirmed by the findings on computer-assisted tomography and the assessment of lung function.

The second group of this classification comprises those disorders presenting **abnormally high greyscale values** in the parenchyma. In this category, the *interstitial lung diseases* play an important role. The quantity of interstitial tissue is increased in many disorders of the lung. Usually reticulum is laid down the alveoli and there is proliferation of cubical type II cells in the walls of the alveoli, sometimes with desqueamation into the lumen, hyperplasia of the bronchiolar epithelium and arteritis of the small pulmonary arteries. These disorders are associated with diffuse nodular and/or irregular opacities on the chest radiograph. They are sometimes accompanied by dry cough and vague chest pains, and usually by breathlessness which is due to characteristic changes in lung function. In many of the conditions the lesions are initially reversible, either by removal of the cause or in response to treatment by steroid or immunosuppressant drugs. Subsequently they can progress via interalveolar and peribronchial fibrosis to diffuse interstitial fibrosis. This may be further complicated by infection, by diffuse dilatation of airways, by emphysematous distortion of parts of the lung and by extensive destruction of the pulmonary vascular bed leading pulmonary hypertension and right heart failure.

The conditions have been classified by Turner Warwick under four headings:

- Widespread granulomas. These include *sarcodiosis*, *beryllium* disease, *extrinsic allergic alveolitis* caused by various organic dusts, *tuberculosis* and other granulomas.

- Interstitial exudates. These are non-inflammatory when due to uraemia or to a raised left atrial pressure as in *mitral stenosis*, or to failure of the left ventricle. Inflammatory exudates occur with infections, *cryptogenic fibrosing alveolitis* and systemic connective tissue disease.

- Disorders caused by inhaled inorganic particulates. These conditions are mainly of occupational origin.

- Tumours and congenital dysplasias.

The lung function of patients with interstitial lung disease usually combines the features of a restrictive ventilatory defect and a defect of gas transfer. However, one or other may predominate depending on the lung pathology. In the late stages the typical features may be obscured by airflow obstruction.

In the interstitial lung diseases the lung volumes can be reduced by an increase in volume of interstitial tissue, replacement of ventilated lung units by fibrous tissue and diffuse interstitial fibrosis which reduces the distensibility of those units which remain. Hence the static lung compliance is reduced and the recoil pressure at total lung capacity is increased (figure 3.1 centre).

In *cryptogenic fibrosing alveolitis*, also known as *diffuse interstitial pulmonary fibrosis* (figure 3.2 top right), the principal feature is thickening of the interstitium of the alveolar wall by collagen. These changes may be dispersed irregularly within the lung. In some patients, a cellular exudates consisting of macrophages and other mononuclear cells is seen within the alveoli in the early stages of the disease. This is called “desquamation”. Eventually, the alveolar architecture is destroyed, and the scarring results in multiple air-filled cystic spaces formed by dilated terminal and respiratory bronchioles. Collagen is scar tissue and is well seen in a scar on the arm,

for example. This can be the end result of a variety of insults including trauma, infection, or a burn. In the same way, interstitial pulmonary fibrosis is the end result of many forms of injury, and it is often impossible to determine what the injurious agent was.

Diffuse pleural fibrosis is another disorder that shows high density in the lung tissue, presenting a thickening of the pleura. This condition usually affects both the parietal and the visceral pleura. It can be primary as following exposure to asbestos or secondary to pleural effusion or haemothorax. The presence of fibrous tissue reduces the lung compliance and increases the recoil pressure.

Finally, we will briefly review *bronchiectasis* and *lung cancer*, also included into this second group of diseases.

The term *bronchiectasis* (figure 3.2 bottom left) means simply dilation of the bronchi, but in general usage also implies the infectious destruction of bronchial walls. The principal clinical feature of bronchiectasis is a chronic, loose cough, usually productive of large amounts of mucopurulent, often foul-smelling sputum. In advanced cases, the sputum settles out into three layers: cloudy mucus on top, clear saliva in the middle, and cloudy purulent material on the bottom. The diagnosis of bronchiectasis can often be made from the history alone. Advanced bronchiectasis can be diagnosed sometimes from the standard postero-anterior chest film. CT scans with thin columnation are proven to be quite useful as well. However, it can be confirmed only after a bronchography is performed.

The symptoms for *bronchogenic carcinoma* (figure 3.2 bottom right) at the time of radiologic diagnosis may be absent. However they usually show one of the following symptom complexes. Cough is the commonest symptom. Hemoptysis, which is the coughing up of blood from the respiratory tract, is a particularly suggestive sign and should never be ignored, particularly in patients over 35 years old who have a history of smoking. Vague nonpleuretic chest pain added to a worsening cough and hemoptysis is a common triad seen in these diseases. Dyspnea may occur, secondary either to an obstruction of an airway or to lymphangitic spread of the tumour.

In the majority of cases of lung cancer, the chest radiograph or CT scan is abnormal at the time of the diagnosis. A common radiographic abnormality is the solitary pulmonary nodule, by definition less than 3 cms in diameter and completely surrounded by lung tissue. Larger lesions are termed masses. In many cases the radiographic appearance of a lesion allows reliable determination of benign nature. Dense, central, concentric, or "popcorn" calcification seen on plain radiographs or CT is a reliable sign that a nodule is benign. Interpretation of calcification on CT scans may be confusing, because lesions containing only a fleck of calcification may represent scar carcinomas growing near an old calcified lesion.

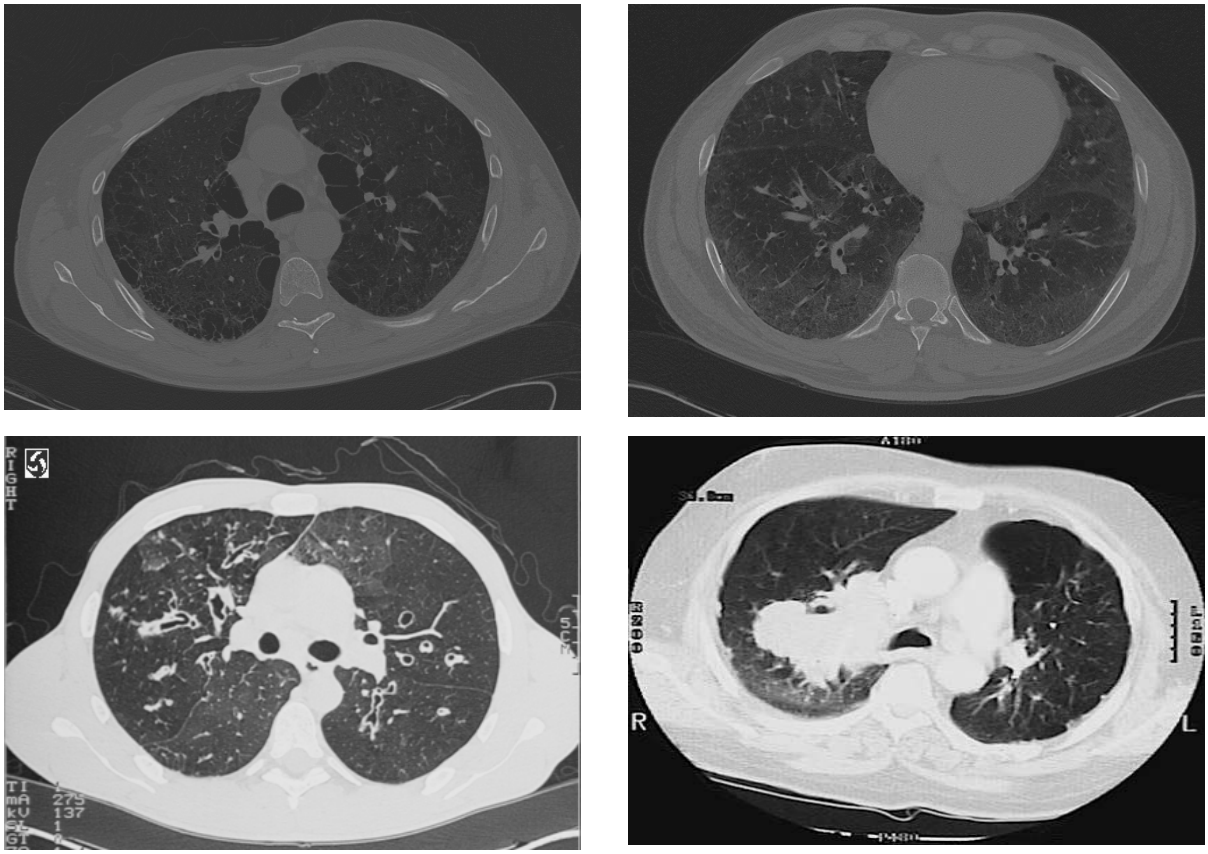

Figure 3.2. CT scans of unhealthy lungs. Top left: emphysema. Top right: fibrosing alveolitis.
Bottom left: Bronchiectasis. Bottom right: Bronchogenic carcinoma

3.4. COMPUTED TOMOGRAPHY (CT)[1-18]

3.4.1. CT AND MAGNETIC RESONANCE IMAGING (MRI)

The application of digital imaging in the field of medicine has been a major innovation. There are many advantages of creating digital chest radiographs, such as the ability to use a computer to manipulate contrast levels and to process the image, the possibility of a greater range of contrast, the control of the image density, and the facility to store the information in high-density form on digital archiving devices.

We will briefly compare here two of the most important modalities of digital radiography, CT and magnetic resonance imaging (MRI). Both of them have inherently good contrast, but the spatial resolution is nor as good as with conventional radiography. This is the main problem with these techniques, however digital radiography is being developed to improve this.

CT has become the screening procedure of choice in evaluating many abdominal and spinal diseases. Improvements in scanning technology have resulted in better spatial resolution and faster data acquisition. This allowed a broader use of CT to areas such as the evaluation of bronchiectasis and other lung diseases such as emphysema. Furthermore, CT is no longer used exclusively for detection and diagnosis, but to guide biopsy and treatment.

As CT technology maturated, MRI has been introduced into clinical practice. The initial application of MRI, like CT, focused on the brain because of its size and lack of respiratory or other physiologic motion. Development of surface coil technology, cardiac triggering, various motion suppression techniques, and fast imaging sequences have made it possible to use MRI for other parts of the body.

The advantages of MRI are that it produces clear anatomical display in any plane, with no radiation risk to the patient, and that it has a soft tissue contrast sensitivity and discrimination unrivalled by any other imaging technique. The underlying technology

is particularly complex and a working knowledge of it is essential to good practice. The radiologist, therefore, needs the support of the technologists.

As stated above, CT has been the initial imaging procedure for evaluating many clinical problems, as well as a non-invasive, efficacious means for assessing and monitoring patients with a wide variety of cancers and other chronic illness. CT has become readily accessible and a vital part of quality medical cares. As the numbers of CT examinations have substantially increased, it has become even more important to conduct CT in a way that maximizes diagnostic information and minimizes risks and costs.

CT is an x-ray-generated imaging study that requires roughly the same photon density as conventional radiography, but is designed to focus x-rays on a limited cross-sectional tissue plane and to utilize those x-rays more efficiently. The efficiency of this method results in excellent contrast sensitivity because of its reduction in scatter, removal of superimposed information, sophisticated detection systems, and sensitive display techniques.

On the other hand, MRI is technique capable of producing thin tomographic sections that, in contrast with CT, requires ionising radiation. It is based on the interaction between radio waves and atomic nuclei in the presence of a strong magnetic field. Whereas the pixel intensity in CT reflects electron density, in MRI it reflects the density of mobile nuclei modified by their magnetic relaxation times, T1 and T2. Because the hydrogen atom, which consists of a single proton in its nucleus, is the most abundant element in the body and because it has a strong magnetic moment, it is the technique used most commonly for “in vivo” imaging. There is a large number of available operator-controlled parameters and it is a real challenge to the practicing radiologist, because even large lesions can go undetected if inappropriate techniques are used. But proper selection of these factors will result in images of high quality that will show appropriate tissues to best advantage with a reasonable expenditure of time.

Despite the fact that very fine soft tissue details can be more readily and clearly seen with MRI, and that in some situations these soft tissues may be obscured by

nearby bone structures in a CT, the point of using Computed Tomography for images in the chest and lungs, is that with MRI, the air-tissue interfaces perturb magnetic fields, and therefore peripheral pulmonary features that may be relevant are not well detected due to the low density of hydrogen atoms in the inflated lungs. Hence, this is the technique used throughout this project to display the images of the lungs.

3.4.2. GENERAL PROCEDURE

CT scannings depend on the same basic physical principles as conventional radiography, namely, the absorption of x-rays by the atoms of the tissues. The difference is that by using multiple projections and computer calculations of radiographic density, it is possible to record finer differences in absorption that can be achieved with conventional films. Also, these differences can be displayed in sectional format without blurring.

The basic components of a CT machine are:

- The x-ray tube, which is similar to the x-ray tube of a conventional machine

- An array of electronic x-ray detectors, placed opposite the tube, is housed in a scanning gantry. In stationary-rotate systems, there are stationary detectors arranged in a ring around the patient, and the x-ray tube rotates within this ring, emitting a fan-shaped beam that always covers the whole width of the body part to be examined. In the rotate-rotate system, both the tube and x-ray detector array rotate synchronously around the patient, the number of detectors being enough to cover the fan-shaped beam, but no more.

- Control devices to rotate the tube (and the detectors where appropriate) around the patient in a short enough time to ensure an image in one breath hold.

- A computer to reconstruct the image from x-rays received by the detectors.

- Some monitors to display the images.

The image is composed of a matrix of picture elements (pixels), the diameter of which determines the resolution of the image. Most machines operate with a fixed number of pixels in the matrix. Thus, the size of each pixel varies according to the diameter of the circle to be scanned. The narrower the scan circle, the smaller the area represented by the pixel and the higher the resolution. By the selection of specific areas (so-called "targeting"), the optimal resolution of the image can be displayed, making available to the operator information that is in the raw data but not displayed when the whole-body section is viewed at one time.

The height of the pixel is determined by the thickness of the section and is chosen by the operator. In chest work, the usual routine is to place the patient supine and obtain contiguous 8-10 mm thick sections from the extreme lung bases to the apices. Thinner sections may be chosen if the lesion being investigated is very small or if partial volume artefacts may be influencing the interpretation of the image. Thus, each pixel has a definite volume. For this reason, it is frequently referred to as a volume element (voxel). The average radiographic density of each voxel is calculated by the computer, and the resulting image consists of a representation of the average density of each of the voxels in the section.

The units have been arbitrarily chosen so that zero is water density, -1000 is air density and +1000 is solid bone. These units have been named Hounsfield units (HU). The range of densities to be displayed on the monitors is selected by the operator. This is necessary because neither the display nor the human eye can appreciate more than approximately 22 shades of grey. Two variables are employed: window width and window centre. The window width is the number of HUs to be displayed. Any densities greater than the upper limit of the window are displayed as white, and any below the limit of the window as black. Between these limits, the densities are displayed in shadows of grey. For lung images, a very wide window of 1000 HU or more and a centre of –400 to –600 HUs are used.

Currently, the instrument in use at the Bristol Royal Infirmary[1-19] (B.R.I.) is the Siemens Somaton Plus 4, a whole-body spiral scanner, whose largest field of view is 500 mm. It has a choice of slice thickness from 1 to 10 mm, determined by

collimation of the x-ray beam. Generally, 2 to 8 mm thickness is used. Various corrections can be applied to minimize any artefacts and image blur from involuntary movement of the patient, such as heartbeat and pulse. Computer storage is sufficient for 3200 images of compressed data in 512x512 format. The endoscopically accessible central airways, which include trachea, carina, lobular and segmented airways, are all routinely visualized using CT scans with 2 mm collimation. This means that one data-set could contain as many as 80 image slices, often in pairs, one "hard" focus and one "soft" focus for each "slice", as we will see in more detail in Chapter 7 when we work with whole data-sets of lung images.

3.4.3. BENEFITS VS. RISKS

Benefits:

- Unlike other imaging methods, CT scanning offers detailed views of many types of tissue, including the lung, bones, soft tissues, and blood vessels.

- CT examinations are fast and simple. Especially in trauma cases, they can reveal internal injuries and bleeding quickly enough to help save lives.

- CT scanning can identify both normal and abnormal structures, making it a useful tool to guide radiotherapy, needle biopsies, and other minimally invasive procedures.

- CT scanning is painless, non-invasive, and accurate.

- Diagnosis made with the assistance of CT can eliminate the need for invasive exploratory surgery and surgical biopsy.

- CT has been shown to be a cost-effective imaging tool for a wide range of clinical problems.

Risks:

- CT does involve exposure to radiation in the form of x-rays, but the benefit of an accurate diagnosis far outweighs the risk. The effective radiation dose from this procedure is about 10 mSv, which is about the same as the average person receives from background radiation in three years.

- Special care is taken during x-ray examinations to ensure maximum safety for the patient by shielding the abdomen and pelvis with a lead apron, with the exception of those examinations in which the abdomen and pelvis are being imaged. Women should always inform their doctor or x-ray technologist if there is any possibility that they are pregnant.

- Nursing mothers should wait for 24 hours after contrast injection before resuming breast-feeding.

- The risk of serious allergic reaction to iodine-containing contrast material are rare, and radiology departments are well equipped to deal with them.

4. DEVELOPMENT OF THE DIAGNOSTIC TOOLS

4.1. PREVIOUS CONSIDERATIONS. CHOSEN LANGUAGE AND ENVIRONMENT.

As we stated in the former Chapter, we will work with data containing lung images and we will make an attempt to extract measures from them in order to detect the presence or absence of aberrations and disorders in the lung tissue. But previously to this, the images must be appropriately processed and treated to obtain a high degree of robustness in our program. Regarding this matter, many researches have been developed in the University's Department of Engineering Mathematics, detailing the image processing from the very first scanning stage to the normalization and segmentation of the extracted slices, which is our starting line. We briefly expose these previous steps below.

First of all, the scanner providing the primary data at the B.R.I. sends information to a host computer[1-19], T.O.S.C.A., from which it can be downloaded in a compressed 3D image format. Lung scans are usually high resolution, with a 2 mm collimation of the x-ray beam. A software package to unstack these slices for analysis with integrity was developed. It was a 16-bit image display and analysis program, optimized for displaying CT data. After this, the normalization and noise suppression problems were treated, and ultimately a consistent segmentation method was proposed, removing information from all the body regions except from the lungs (figure 4.1).

At this point, the aim of this project is to use these segmented images to develop a learning machine that is able to make an automatic detection of diseases.

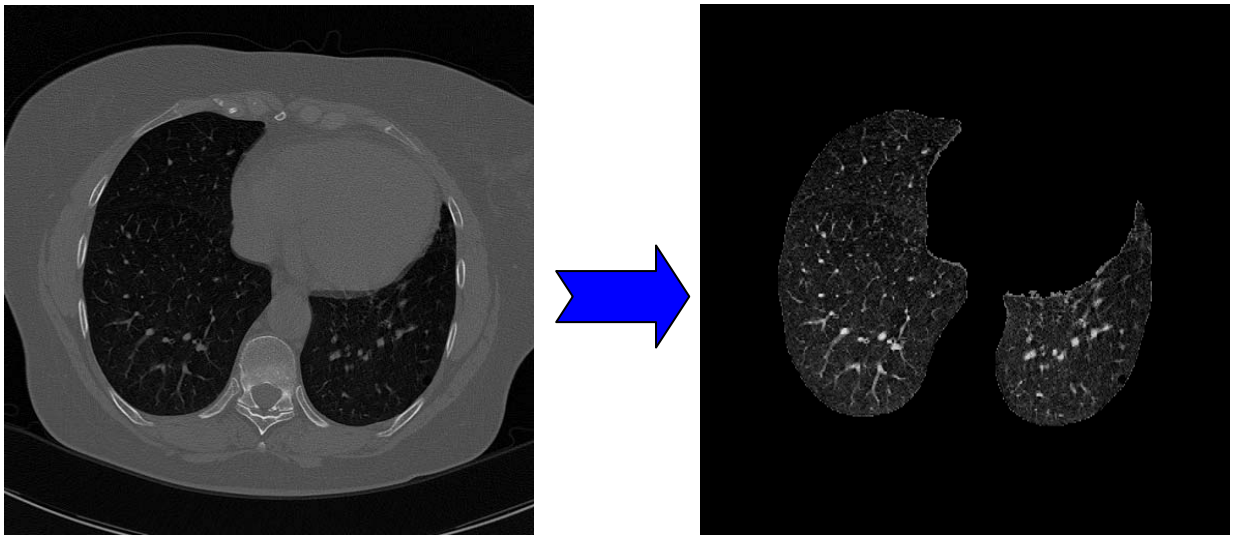

Figure 4.1. Segmentation process

The language that we will utilize is MATLAB v. 6.0., which is a high-performance language for technical computing that integrates computation, visualization, and programming in an easy-to-use environment, and it takes considerably less time to write a program in MATLAB than in a scalar non-interactive language such as C or FORTRAN.

In this environment we can also use a family of application-specific solutions called toolboxes, as well as a large number of mathematical libraries and functions for visualization, which will make our work much easier.

4.2. DEVELOPED METHODOLOGY AND PROGRAM STRUCTURE

The developed methodology to tackle the current subject comprises two different modules:

1. The first module performed texture analysis on the segmented lung images. The image is divided into square regions, and texture measures are extracted from these regions.

2. The second module was the definition of a learning machine based on artificial intelligence techniques (Neural Networks) that merged the discovered textural parameters into a diagnosis regarding the presence or absence of a pulmonary problem.

Image noise and limitations of the human visual system can sometimes hinder radiologists, particularly in low contrast detection problems. In addition, humans often miss a diagnosis because the image feature was simply overlooked.

Computers approach image analysis and feature extraction differently than humans. Computers consistently apply, without bias or distraction, the rules they are programmed with for feature extraction. The different approach taken by computers is often complementary to human observers. Therefore, researchers have recognized that human interpretation of CT images may be improved by adding the second opinion introduced by an expert system[1-20].

Artificial Neural Networks, on the other hand, have gained acceptance as an alternative to traditional statistical modelling[1-21]. They are useful for performing nonlinear statistical analysis without having to define a formal statistical model and, even though they often lack explanatory power and are developed using an empirical training procedure, they have proven to be particularly suitable for prediction and classification problems in medical decision making.

There are several ways to apply computers in the diagnostic interpretation of medical images. Specifically, in our system, we base it on image features extracted by a human observer. These features must be integrated then into a diagnosis, using standard diagnosis criteria.

The criteria used by us, as mentioned in previous Chapters, consists of considering as unhealthy any aberration in the grey level scale in certain areas of the lungs, i.e., abnormally low density levels and abnormally high density levels, without regard to the kind of disease in most of the cases, and leaving this task to a posterior examination by expert physicians.

Initially, the criteria are applied directly by a human observer, and an example of the result of this diagnosis is shown in figure 4.2., where we can see two pictures from two different patients. The first of them shows widespread emphysema with scarring in both lungs, which means a lost of structure and areas with clearly low grey levels. The second one, on the other hand, shows greyscale values significantly higher than normal, suggesting fibrosis.

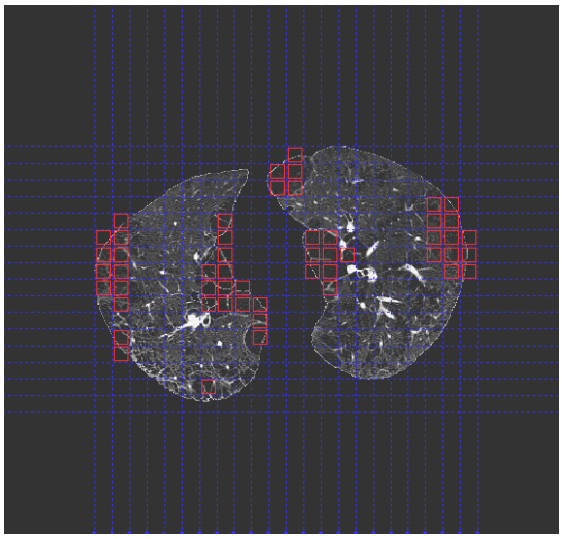

4.2.a. Emphysematous lungs

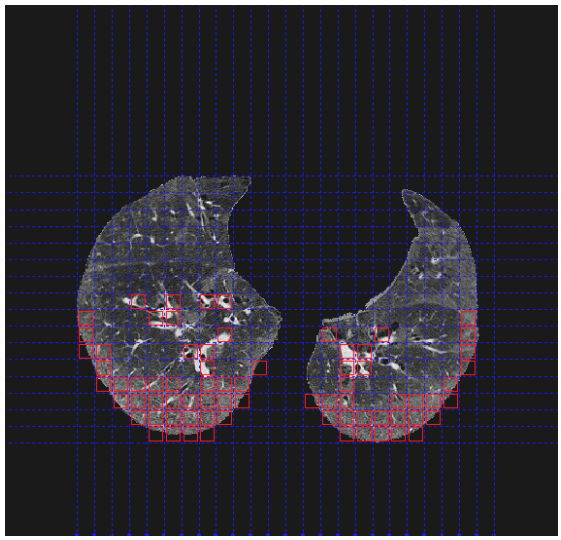

4.2.b. Fibrous lungs

Figure 4.2. Diagnosis made by a human observer. Some brightness corrections have been made to the pictures in order to see them clearer.

The human observer marks the square regions that he thinks are unhealthy, following a visual criterion. Using an expert system, this criterion is transformed into the digital environment, and applied in a consistent and unbiased way. Such a system is typically limited by imperfections and bias in the criteria it uses so we use an artificial neural network to reduce the bias.

An expert system using neural networks learns and develops the decision criteria based on the examples presented to it, in our case these examples consist of square regions as seen in the previous figures. As more representative and unbiased examples are collected and used to train the neural network, it develops a model of the diagnostic process that it can then generalize to unseen cases.

The challenge then becomes collecting the cases, selecting good features to use in the construction of the model and incorporate as much a priori knowledge about the problem as possible.

In next section the different steps to develop the method, as well as the interface of the program used to do the process will be described in more detail.

4.3. 2D INTERFACE

For the current project, we have developed two different user interfaces, one for 2D process and diagnosis, and the other one, which will be introduced in Chapter 7, for a 3D approach of the problem.

Therefore, the initial investigations will focus on isolated single slices. The stored patient CT scans of this study have a resolution of 8 bit grey scale precision, i.e., 256 greyscale levels, and we have developed a graphical interface to aid in the display and process of such images, as shown in the figure below.

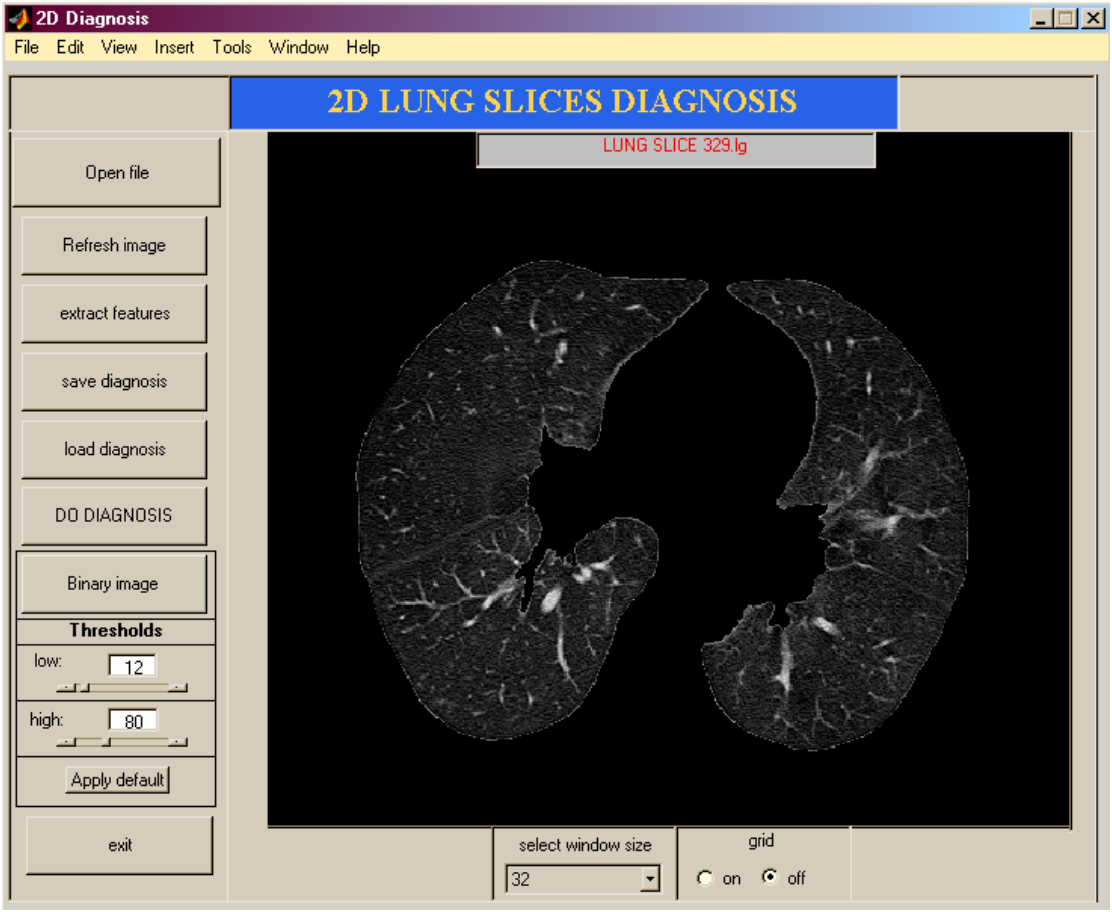

Figure 4.3. 2D Interface

To run the program, it is only necessary to type “interface” in the prompt of MATLAB, without any additional parameters.

Next, the different elements of this graphical tool are described, paying special attention to the *extract features* and *DO DIAGNOSIS* buttons, that comprise the main tasks developed.

- **Open file:**

  Before the pre-process and segmentation of the images, the used files had extension “*.pgm”, and after the segmentation, the information was stored numerically in a file with extension “*.lg”.

  Having the information this way in a file allows us to work in MATLAB with these values, reading them directly from the file and storing them in internal variables. In this type of files, we have the values sorted in four different columns, as follows:

Column1	Column2	Column3	Column4
File name	Row	Column	Intensity

  As we said previously, our work starts after the segmentation of the lungs is done, hence “*.lg” is the extension of the files we will work with, as we can see in figure 4.3., where slice “329.lg” is displayed.

  Two considerations are relevant here. Before opening a file, we have to check if it is the correct kind of file, i.e., “*.lg”, otherwise the following error message is displayed (figure 4.4):

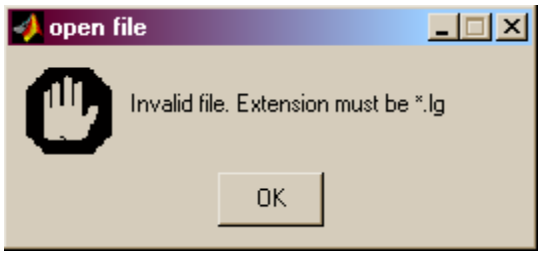

Figure 4.4. Open file error

On the other hand the size of the square window's side we will use to mark the unhealthy areas, and the option of grid on/off are checked and used for the opened file, although this can also be done at any time later on.

- **Refresh image:**

This will remove all the marks from the current image, and besides it will update the parameters for a new window size, and for displaying the binary image, if this option is selected (we will see this below). Therefore, this must be used when we want to make a new process or diagnosis of the image.

- **Extract features:**

The feature extraction and all the texture features used in this project will be fully described in next Chapter. However, we will introduce here some concepts that will give us a general idea of the process.

Before carrying out the extraction of the textural measures, a human observer has to deal with the problem of doing manually some previous diagnoses, based on visual perception. These diagnoses will be used in posterior steps to train the neural network. This step, hence, is critical for our application.

We have used throughout this research 16x16 pixel regions, although a new study could be easily done with the other available window sizes, 8x8 and

32x32. The way to select the areas we think are unhealthy is simply by clicking on them. A double click would remove the mark, leaving it as healthy. After completing the visual diagnosis, it can be saved to a file, so that it would be available for a posterior inspection. This allows the physicians to have available a data base for different patients and for different CT scans from the same person.

Once we have selected the unhealthy areas, we are ready to proceed with the features extraction. When we press this button, a new dialog window appears (figure 4.5), displaying a list of 25 features. This set of features has been selected from all those reviewed in Chapter 2.

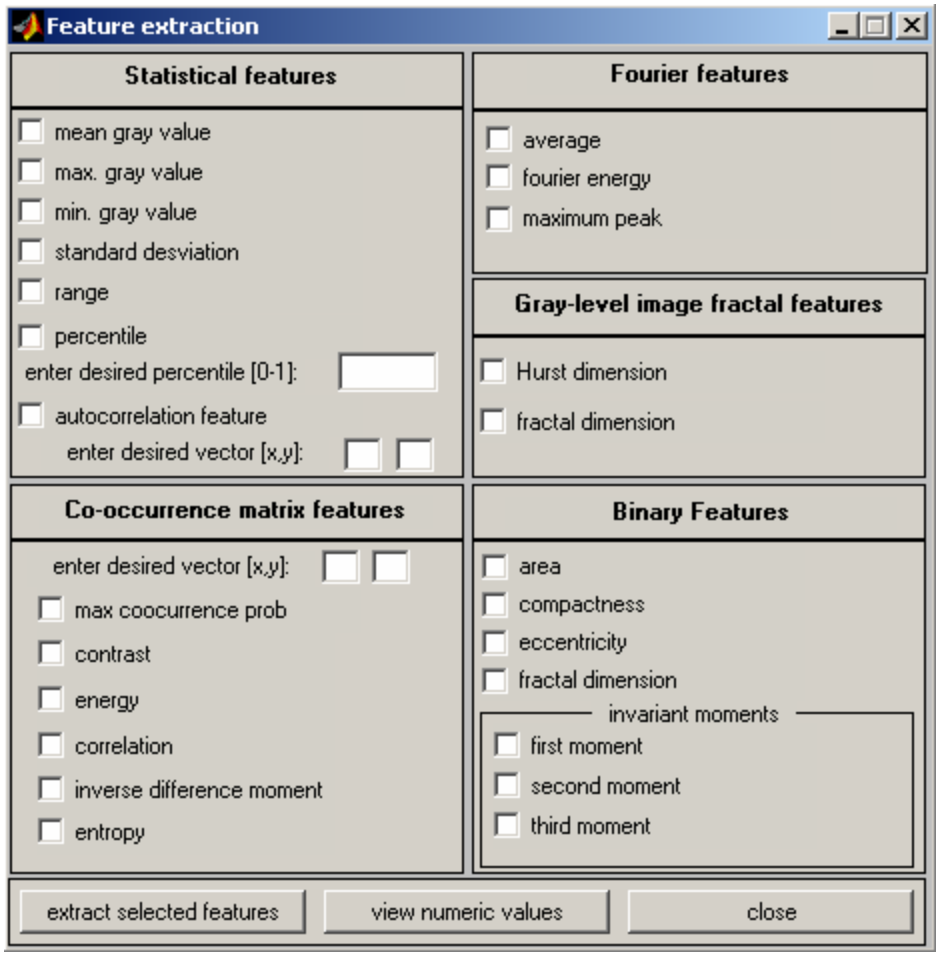

Figure 4.5. Features extraction

We have divided them into five groups:

- **Statistical features:** mean grey value, maximum and minimum grey value, standard deviation, range, percentile, and autocorrelation.

- **Co-occurrence matrix features:** maximum co-occurrence probability, contrast, energy, inverse difference, and entropy.

- **Fourier features:** average, energy, and maximum peak.

- **Grey-level image fractal features: F**ractal dimension and Hurst exponent.

- **Binary features:** area, compactness, eccentricity, binary fractal dimension, and first, second and third invariant moments.

Their definitions and characteristics will be reviewed in Chapter 5. The question now is to understand the process of selecting a subset of features from these initial 25 ones.

Feature selection can be empirical (pick features that have a high correlation with the known diagnosis when considered alone) or analytical (pick features based on some selection algorithm). We will start using an empirical approach, and when we introduce the neural networks, we will describe an analytical algorithm.

For empirical observations, we will assume that the observer's decision can be modelled by a random variable that fits the binormal model (figure 4.6).

The probability density functions of this random variable under the two hypotheses, i.e., normal tissue and abnormal tissue are assumed to be normally distributed. Since the separation between both populations is not perfect, these

distributions will overlap and the decision threshold, that will eventually be fixed by the neural network, will always involve some compromise between false positive (test calls an actually negative case positive) and false negative (test calls an actually positive case negative) decision.

The underlying binormal assumption is generally preferred because many experimentally determined curves are binormal.

Therefore, we have to look for those features that present a better discrimination between both populations of normal and abnormal regions, and combine them in order to discover which one of the many possible combinations of measures has the best performance for differentiating both classes.

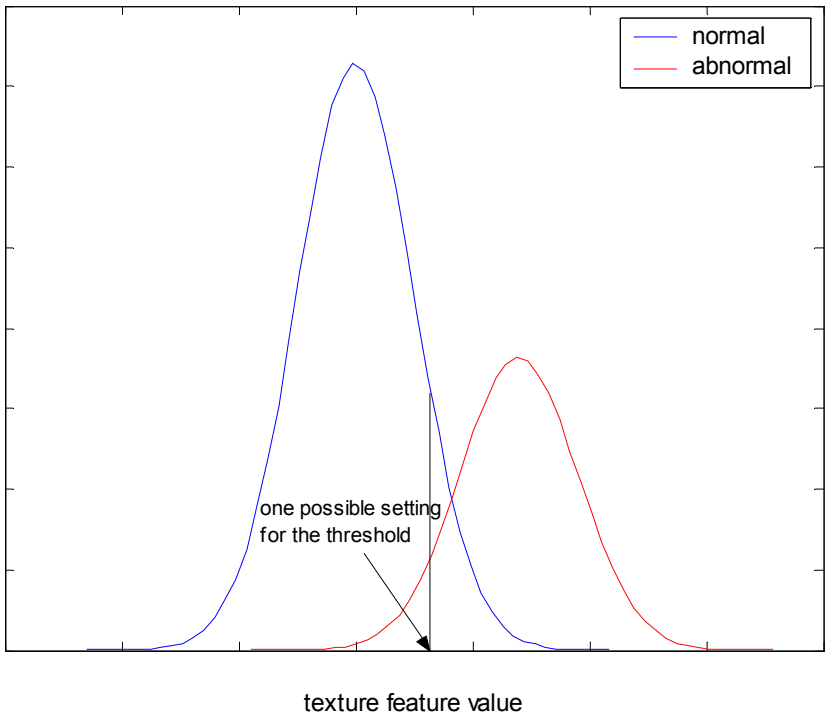

Figure 4.6. The binormal model for empirical analysis.

Now we will show an initial example to demonstrate that the previous considerations are quite close to reality. Using a real lung image, such as that from figure 4.2.a, which presents some very dark areas, we can see if both classes actually fit the binormal model.

To accomplish this, firstly the manual diagnosis was made, and then the "extract features" function was used. Two features were selected: mean grey value, and standard deviation.

For each feature, a normalized histogram and a box plot with lines at the lower quartile, median, and upper quartile values, have been calculated (figure 4.7). In the box plot, apart from the mentioned lines, we also have the whiskers which are lines that extend from each end of the box to show the extent of the rest of the data. The outliers are data with values beyond the ends of the whiskers. If there is no data outside the whisker, a dot is placed at the bottom whisker.

Therefore, the box plot gives an idea of the dispersion of the data inside the same class. We can also make a comparison between both classes, normal and abnormal, and see if they completely overlap or if, on the contrary, they are separable. None of the features that we will use are perfect, hence the classes will not be separable for any of them. On the other hand, the histograms are another way of showing the distributions, and we have normalized them to compare both populations of data.

In figure 4.7, we can see that the diagnostic potential of the mean value is relatively good, if we compare it with the standard deviation feature. In the former, the populations overlap, but we still have some discriminatory power. Besides, the histograms show that the binormal assumption adopted for modelling this problem is close to reality in this case. On the contrary, the standard deviation presents a very poor reliability, because normal and abnormal classes completely overlap, until the point that the median value is almost the same for both of them.

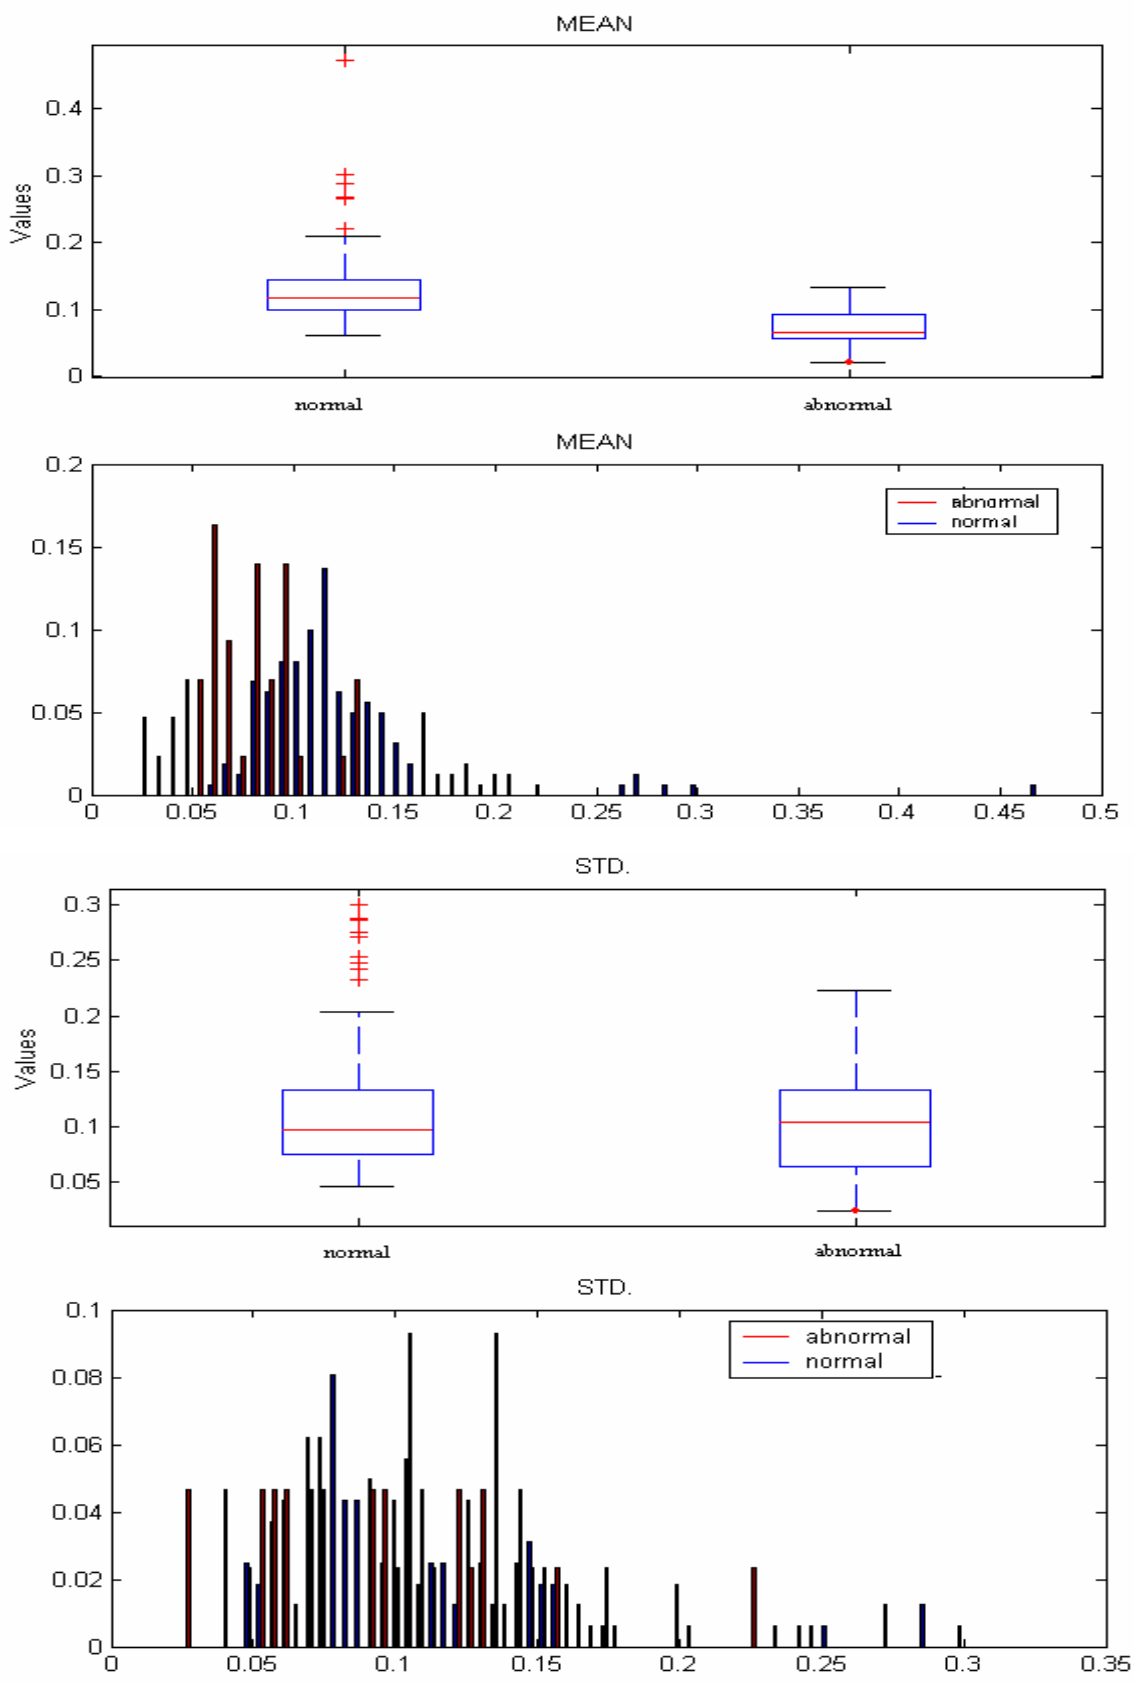

Figure 4.7. Mean and Standard deviation features. Comparison between both classes.
Top of each picture: Box plot of the distributions. Bottom: Normalized Histograms.

Once some features have been extracted, the numeric values are stored in internal variables, and we also have the option of storing them in a file in order to have access to them at any time. These files have extension "*.res". An extract of the file with the numeric results of the previous example is presented in table 4.1.

The classes are labelled as "0", if abnormal, and "1" if normal. We also have a mark of "1" if the regions belong to the contour of the lung, or "0" if they belong to the inner parts. This last distinction is due to the fact that some features can be very good for regions belonging to the centre of the lungs, and very bad for the peripheric areas, since in these zones there is a lack of structure and therefore a lack of information. One clear example of this can be any feature related to the co-occurrence matrix, because for these measures a square matrix with full information is needed to have an accurate result. On the other hand, there are some features, such as the mean grey value, that are not affected by the fact of belonging to the centre or to the contour of the lungs.

CLASS	**MEAN**	**STD**	**EDGES**
1	0.1149	0.0959	1
1	0.1340	0.1261	1
0	0.0618	0.1307	1
0	0.0246	0.0287	0
1	0.0949	0.0714	0
1	0.1127	0.0785	0
1	0.1189	0.1265	1
0	0.0540	0.1036	0

Table 4.1. Numeric values of the selected features, extracted from file.

After the feature extraction, an optimal subset of a reduced number of features must be selected. This subset, if our method if successful, must be the one that better captures the texture information of the images, so that it contributes to the neural network's ability to diagnose diseases and aberrations in the lung parenchyma.

A last consideration must be mentioned here. As we stated before, some features can come up with different results depending on the position of the regions, i.e., centres or edges. Some different results can appear as well if we study the same feature for several diseases, because the conditions and structure of the disorders can change considerably. Therefore, the best subset of features for emphysema does not necessarily have to be the optimal one for fibrosis or bronchiectasis, although it also depends on the type of features that we work with.

- **Save diagnosis:**

After the human observer marks manually the unhealthy areas, or the diagnosis is carried out automatically by our tool, we will want to store this diagnosis into a file. Doing this, the physicians will save time, since they will not have to wait for all the process to be accomplished again each time that they want to consult the CT scans from some patient.

The extension used for this type of files is "*.dgn". If diagnoses using different window sizes are done, they do not need to be stored in three different files. They can be saved in the same file, and we can also overwrite the diagnosis made for a particular window size at any moment, without affecting the diagnoses that belong to the others.

Therefore, a manageable data base with lung scans and their diagnoses can be easily organized in order to save time for the expert physicians.

- **Load diagnosis:**

This allows us to recover the diagnoses that have been previously done for the current image.

The diagnosis is read from a file, and immediately after the unhealthy regions are marked with a red square if the window size is 16 or 32, or with a red X in the corner of the region if the window size is 8, for a clearer visualization of the image.

If the diagnosis that we are trying to load is empty for the current window size, the following error message is displayed:

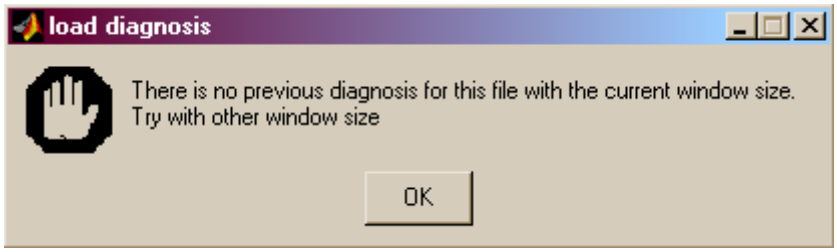

Figure 4.8. Load diagnosis error. Invalid window size.

On the other hand, if the diagnosis file does not belong to the current image, this error message is displayed instead:

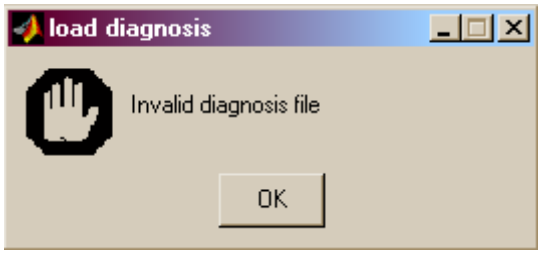

Figure 4.9. Load diagnosis error. Invalid file.

- **DO DIAGNOSIS:**

This is the last step of the whole process, where features extraction and artificial neural networks are combined and merged into the same diagnostic tool.

After the final selection of a subset of representative features, which are the ones that best define the textural properties of the lungs images, these measures are used to train an artificial neural network for the automatic detection of pulmonary diseases.

The neural network used in this project and all its characteristics are explained in Chapter 6. Basically, the process consists of training and testing the net with different subsets of features, the so-called feature vectors, until the performance obtained can not be beaten, and hence the selection process stops at that point. We also have to consider that the size of the subset must not be too large, in order not to complicate the neural network too much. The artificial neural networks are used in the model building process because they are useful for capturing complex hidden relationships between many input variables.

To accomplish the diagnosis of the current lung image, firstly the image is divided into square regions, depending upon the window side, which in our study consists of 16x16 pixels regions. Then, the feature vector is calculated for each area. We will see later in next Chapters that different feature vectors are used for the edges and for the centres, and hence different neural networks are needed as well. At the next stage, these feature vectors are simulated using the neural networks, and then we come up with the final diagnostic results, detecting presence or absence of disease.

Finally, the unhealthy areas are marked, and the diagnosis is showed in the image. Then, we can save it to a file and process a new slice.

- **Binary image / grey level image:**

Some of the features that we work with in this project, as stated in a previous section, are binary features. This group of binary measures, that will be fully reviewed in next Chapter, include *binary area, compactness, eccentricity, binary fractal dimension, and first, second and third invariant moments*.

In order to utilize these features, first we have to convert the image under study into a binary image. As an example of this, we have used the images from figure 4.2, which present abnormally low and high density levels respectively.

The greyscale range that we are using is 0 (black) – 255 (white), and in the conversion, all the values below a certain low threshold and above another high threshold are turned to white, while the rest of the values are turned to black.

The choice of the thresholds is an important and delicate question, because this will have an influence upon the quality of the features. We have chosen a low threshold of "12" for the abnormally low greyscale values, and a high threshold of "80" for the abnormally high greyscale values. The reasons of this are explained in Chapter 5.

In figure 4.10, we have the binary versions of those images from figure 4.2. For the first image, which corresponds to an emphysematous lung with large dark areas, we have used a low threshold of "12", as mentioned before, and a high threshold of "255", because we are only interested in this case in detecting the presence of low-density areas. For the second image, which corresponds to a fibrous lung with large bright areas, the high threshold was fixed to "80", and the low threshold to "0", for similar reasons than the case before.

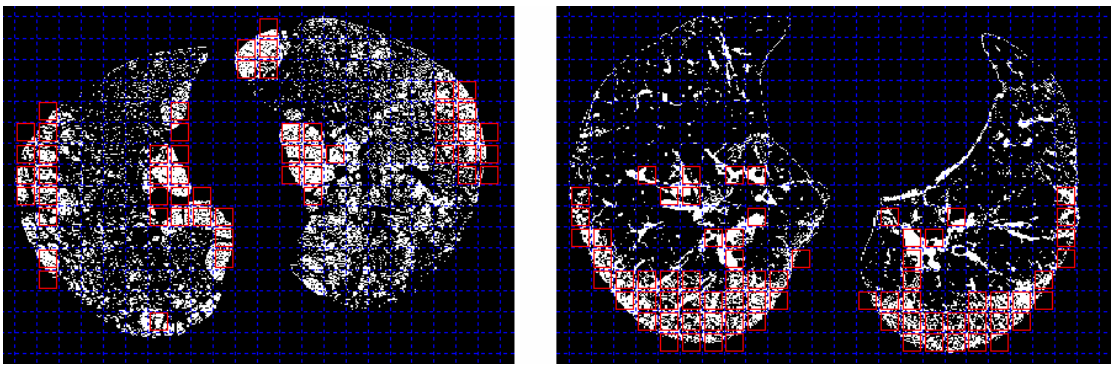

4.10.a. Emphysematous lung 4.10.b. Fibrous lung

Figure 4.10. Binary images, corresponding to figure 4.2.

In the figure above, the unhealthy regions have been also marked, to demonstrate visually that the thresholds actually fit with the diagnoses carried out by a human observer. Therefore, we can also use these binary images as a means to help in the reduction of mistakes when a visual and manual diagnosis by humans is accomplished.

Figure 4.11 shows the graphic tools available in our interface to do the conversion. There is a toggle button that displays the binary image when pressed and the greyscale image otherwise. The thresholds can be set directly by hand or using sliders. The default values for the thresholds are “12” and “80”.

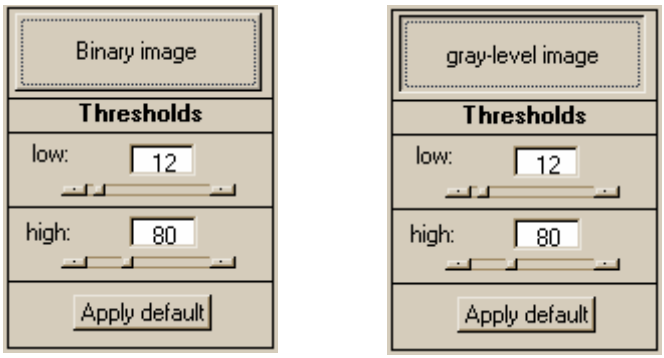

Figure 4.11. Binary conversion tools.

- **Exit:**

Terminate the program. The graphical user interface is closed. As an alternative to this, select Close from the File menu or click the close box in the figure.

- **Select window size:**

The options are these three: 8, 16, and 32. Throughout this project, a window size of 16 has been used, therefore for the obtention of useful results, this is the size that must be utilized when an automatic diagnosis is required.

The same process, i.e., selecting new features and training again some neural networks, should be repeated for the other sizes if we are determined to use them.

- **Grid:**

This turns the current image's grid lines on and off. For a better visualization, it is highly recommended to turn them on.

4.4. SUMMARY AND CONCLUSIONS

An automatic diagnostic tool approach is presented using texture analysis and artificial neural networks to assist physicians in the diagnosis of density disorders in the pulmonary parenchyma.

An optimal subset of textural features is needed to achieve a high performance, being the ideal goal of this system to accomplish a performance similar to physicians for disease detection. In Chapters 5 and 6, we will try to fulfil this aim.

A neural network, basing its diagnosis on these image features extracted from training cases during its development, will be utilized as the chosen artificial intelligence technique to merge the texture measures into a diagnosis. In Chapter 6, the neural networks are explained in detail.

To summarize, the development of the building model process consists of the following steps:

1. Investigate representative cases of study.

2. Determine a criterion for the diagnosis of these cases.

3. Extract several features from image regions.

4. Process the features.

5. Merge the extracted features into a diagnosis using one or more artificial neural networks.

6. Test developed system using unseen images.

Finally, the file extensions used in this project are summarized below:

Extension	**Information contained**
*.lg	lung region data points in row, column, and intensity groups
*.dgn	diagnosis
*.res	numeric values from the features extraction

Table 4.2. File extensions

5. FEATURE SELECTION

5.1. INTRODUCTION

Texture analysis is an active research field and a large number of schemes for texture feature extraction have been proposed[1-22]. In Chapter 2, a review of the most popular techniques and models for texture analysis, as well as its main problems and some applications, was carried out.

We have developed a feature extraction method that consists of splitting the whole lung image in 16x16 pixels regions, and extracting a texture feature vector for each region in each image. We are looking for a method that describes textures in a form suitable for pattern recognition, and as a result of the description, each region's texture is represented by a feature vector of properties. The next step would be to estimate the abnormality of each region based on these feature vectors, and to find a decision rule assigning a texture to some specific class, i.e., normal or abnormal (see Chapter 6).

The aim of this Chapter is to present a set of features, selected from all those described in Chapter 2, and to study which of them have a significant diagnostic potential, and can be reliably utilized to discriminate between healthy tissue and unhealthy tissue.

The measures initially chosen for this investigation can be named as "classical features", in the sense that they can be found in many text books about image processing and analysis[1-23, 1-24]. The main criteria utilized for the selection of the textural parameters were:

- High popularity in the literature, which is a means to check their respective performances in previous investigations, and therefore to justify their use in ours.

- Ease of implementation and use of the parameters in the design of the architecture. This will allow us to reduce the number of unclear assumptions that can cause unknown effects on the final program and the user interface design.

- Simplicity and efficiency. Within this project the procedure for recognizing the problematic regions and distinguishing them from the non problematic ones must be as simple as possible, because very complex features can increase considerably the computational cost. However, some complex features can be very efficient for discriminatory tasks, so we must have a balance between efficiency and computational cost as well.

Once this first set of features has been chosen, the next question in our texture-based tissue characterization method is to select the optimal subset of features from this initial set that best performs the detection of diseases.

In the previous Chapter, we saw that the feature selection can be accomplished in an empirical or in an analytical way. Here we will introduce an empirical approach, picking features without using any selection algorithm, in which a qualitative and visual comparison among the measures will be developed. We will not come up with a quantitative study until Chapter 6, where the neural networks will be utilized to check the performance of the textural parameters.

Each feature has been tested in different lung images with abnormally low, abnormally high and normal greyscale levels, and their respective histograms and density functions have been plotted. The binormal hypothesis described in former sections (figure 4.6) will be further confirmed for some of the features and, on the contrary, rejected for others.

Due to the fact that the differentiation between both populations of texture patterns is not fully complete, some overlapping will appear in the distributions, hence in the final results some false negative and some false positive cases will be present. The neural networks will be the tools used to optimize the task of minimizing these errors.

Initially, we could expect that those features that are actually very clearly related to human visual perception, such as the mean grey value, will come up with a relatively good differentiation between healthy and unhealthy areas. Notwithstanding, there are other texture properties connected to the spatial distribution of the grey values, such as smoothness, coarseness, compactness or regularity, that we are not able to predict with the naked eye how good they are to characterize disease conditions.

To summarize, in the next sections of this Chapter we will deal with the following points:

- Review of the textural features used in the set, which we will divide in five groups as follows: *statistical features, co-occurrence matrix features, Fourier features, grey-level image fractal features*, and finally a new group that was not mentioned in Chapter 2, but that is interesting for our research yielding good results, composed by *binary features*. We have eventually come up with a final list of 25 textural parameters, detailing their main characteristics and equations.

- Initial trials to evaluate the discriminatory potential of each feature, using two sets of samples from lung tissue regions that belong to two different patients. One of them presents abnormally low parenchymal density, which corresponds to a clear condition of emphysema. The other one presents abnormally high parenchymal density, which in this case was diagnosed as fibrosing alveolitis.

We will check if the features that are liable to detect the underlying structural properties that characterize the emphysematous conditions are also able to detect other diseases, such as fibrosis. In this case, the simplicity of our system would be increased, and the expense of time in the feature extraction stage of the process reduced.

The tools used for this purpose are the normalized histograms of the populations of both classes, normal and abnormal, and a box-shaped density function, which allows us to visualize the distributions in a clearer way. This will be described later on.

- Finally, a further word about features in the edges is necessary. In the regions that belong to the contour, there is a lack of structure that can bias the results for some features, and therefore a separate study is suggested.

5.2. SELECTION OF A SET OF FEATURES TO EVALUATE[1-23, 1-24]

The feature computation for this investigation involves extracting textural measures based on:

5.2.1. STATISTICAL FEATURES

In the next equations, *I(i,j)* represents one pixel of the image, and *N* is the window size, which in our research has a value of 16, representing 16x16 pixel regions.

- **Mean Grey Value:**

  An indicator of the average tone of the region.

$$Mean = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} I(i,j) \tag{5.1}$$

- **Maximum and Minimum Grey Value:**

  An indicator of the highest and lowest tones of the region.

$$Max = Max\{I(i,j), i = 1,\ldots,N, j = 1,\ldots,N\} \tag{5.2}$$

$$Min = Min\{I(i,j), i = 1,\ldots,N, j = 1,\ldots,N\} \tag{5.3}$$

- **Standard Deviation:**

  An indicator of how much variation exists in the image with respect to the average tone.

$$Std = \sqrt{\frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} (I(i,j) - mean)^2} \tag{5.4}$$

- **Range:**

  An indicator of how much absolute variation exists in the region. It depends only on the extreme tones in the area.

$$Range = Max - Min \tag{5.5}$$

- **Percentile:**

  The percentile of a distribution of values is a number x_p such that, after the data has been ordered from smallest to largest, a percentage p of the population values are less than or equal to x_p. For example, the 25th percentile, also referred to as the .25 quantile or lower quartile, of a variable is a value (x_p) such that 25% (p) of the values, in this case grey levels, of the variable fall below that value.

  Similarly, the 75th percentile, also referred to as the .75 quantile or upper quartile, is a value such that 75% of the values of the variable fall below that value and is calculated accordingly. The range of valid values for the percentiles, hence, is [0,1]. The lower and upper quartiles are the percentiles most frequently used. For our study, the 0.75 quantile was utilized, yielding satisfactory results, as we will see in next sections. The percentiles are calculated following these steps:

  1. Rank the data from lowest to highest

  2. Compute the location L from equation 5.6, where n is the number of data points and k is the percentile.

$$L = \frac{k}{100} n \tag{5.6}$$

  3. The L^{th} score counting from the lowest is the k^{th} percentile. If L is not a whole number, then round up and down, and average them out.

- **Autocorrelation:**

Measuring spatial frequencies is the basis of a large group of texture recognition methods. Textural character is in direct relation to the spatial size of the texture primitives; coarse textures are built from larger primitives, and fine textures from smaller primitives. Fine textures are characterized by higher spatial frequencies, coarse textures by lower spatial frequencies.

One of the many related spatial frequency methods evaluates the *autocorrelation function of a texture*. In an autocorrelation model, a single pixel is considered a texture primitive, and primitive tone property is the grey level. Texture spatial organization is described by the correlation coefficient that evaluates linear spatial relationships between primitives. If the texture primitives are relatively large, the autocorrelation function value decreases slowly (coarse texture) with increasing distance, while it decreases rapidly (fine texture) if texture consists of small primitives. If primitives are placed periodically in a texture, the autocorrelation increases and decreases periodically with distance.

The normalized autocorrelation function of texture is described by the equation below:

$$Cff(p,q) = \frac{MN}{(M-p)(N-q)} \frac{\sum_{i=1}^{M-p} \sum_{j=1}^{N-q} I(i,j) I(i+p, j+q)}{\sum_{i=1}^{M} \sum_{j=1}^{N} I^2(i,j)} \tag{5.7}$$

In equation 5.7, *p,q* is the position difference in the *i,j* direction, and *M,N* are the image dimension, in our case $M=N=16$. Normally, the values used for parameters *q* and *p* are (0,1), (1,1), and (1,0). After several trials, the pair (0,1) was chosen for a further study here.

5.2.2. CO-OCCURRENCE MATRIX FEATURES

This method of texture description is based on the repeated occurrence of some grey level configuration in the texture. This configuration varies rapidly with distance in fine textures, and slowly in coarse textures[1-25, 1-7]. Let the analysed part of a textured image a rectangular window *N*x*N* (in our case, 16x16). An occurrence of some grey level configuration may be described by a matrix of relative frequencies $P_d(i,j)$ for a displacement vector $d=(dx,dy)$ as follows. The entry (i,j) of P_d is the number of occurrences of the pair of grey levels i and j which are a distance d apart. Formally, it is given as

$$P_d(i,j) = \left|\{((r,s),(t,v)) : I(r,s) = i, I(t,v) = j\}\right| \tag{5.8}$$

where $(r,s), (t,v) \in NxN, (t,v)=(r+dx,s+dy)$, and $|.|$ represents the cardinality of a set of data.

As an example, consider the following 4x4 image containing 3 different grey values:

$$\begin{matrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 2 & 2 \\ 0 & 0 & 2 & 2 \end{matrix}$$

The 3x3 grey level co-occurrence matrix for this image for a displacement vector of $d=(1,0)$ is given as follows:

$$P_d = \begin{pmatrix} 4 & 0 & 2 \\ 2 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

Here the entry (0,0) of P_d is 4 because there are four pixel pairs of (0,0) that are offset by (1,0) amount. Examples of P_d for other displacement vectors are given below (table 5.1):

d=(0,1)	$P_d = \begin{pmatrix} 4 & 2 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}$
d=(1,1)	$P_d = \begin{pmatrix} 2 & 0 & 2 \\ 2 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}$

Table 5.1. Co-occurrence matrix examples.

Notice that the co-occurrence matrix so defined is not symmetric. But a symmetric co-occurrence matrix can be computed by the formula $P=P_d+P_{-d}$.

The co-occurrence method describes second-order image statistics and works well for a large variety of textures[1-26]. Good properties of the co-occurrence method are the description of spatial relations between tonal pixels, and invariance to monotonic grey level transformations. On the other hand, it does not consider primitive shapes, and therefore cannot be recommended if the texture consists of large primitives. Memory requirements represent another big disadvantage, although this is definitely not as limiting as it was few years ago. The number of grey levels of the images we are working with, though, has been reduced from 256 to 32, which decreases the co-occurrence matrix sizes, without experimenting a considerable loss of grey level accuracy in practice.

Although co-occurrence matrices give very good results in discrimination between textures, the method is computationally expensive. A fast algorithm for co-occurrence matrix computation is given in[1-27].

The following algorithm calculates the co-occurrence matrix P_d from the image *I(i,j)* and the displacement vector *d=(dx,dy):*

1. Assign $P_d(i,j)=0$ for all $i,j \in [0,L]$, where L is the maximum brightness

2. For all pixels *(r,s)* in the image, determine *(t,v)* which has the relation *d* with the pixel *(r,s)*, i.e., *(t,v)=(r+dx,s+dy),* and increment P_d thus,

$$P_d(I(r,s),I(t,v)) = P_d(I(r,s),I(t,v)) + 1 \tag{5.9}$$

There are some interesting properties in this matrix. Values of the elements at the diagonal of the co-occurrence matrix $P_d(k,k)$ are equal to the area of the regions in the image with the brightness k. Thus the diagonal elements correspond to the histogram. The values of the elements off the diagonal $P_d(k,j)$ are equal to the length of the border dividing regions with brightnesses k and j, $k \neq j$. For instance, in an image with low contrast the elements of the co-occurrence matrix that are far from the diagonal are equal to zero or are very small, and the texture is coarse. For high contrast images the opposite is true.

A set of features proposed by Haralick[1-7] include:

1. **Energy:**

   It is also called Uniformity or Angular Second Moment, and represents an image homogeneity measure: the more homogeneous the image, the lower the value. Its value is lowest when the elements of the co-occurrence matrix are all equal.

$$Energy = \sum_{i,j} P_d^2(i,j) \tag{5.10}$$

2. **Entropy:**

   It is a measure of randomness, achieving its highest value when all elements of P_d are equal.

$$Entropy = \sum_{i,j} P_d(i,j) \log P_d(i,j) \tag{5.11}$$

3. **Maximum probability:**

   This property gives an indication of the strongest response to P_d.

$$Max.prob. = \max_{i,j} P_d(i,j) \tag{5.12}$$

4. **Contrast:**

This descriptor is a measure of local image variations. It has a relatively low value when the high values of P_d are near the main diagonal since the differences *(i-j)* are smaller there.

$$Contrast = \sum_{i,j} |i-j|^2 P_d(i,j) \tag{5.13}$$

5. **Inverse Difference moment:**

It has the opposite effect of the previous characteristic. Its value is high when the high values of P_d are far from the diagonal.

$$Idm = \sum_{i,j;i\neq j} \frac{P_d(i,j)}{|i-j|^2} \tag{5.14}$$

6. **Correlation:**

It is a measure of image linearity. Linear directional structures in the direction *d* result in large correlation values in this direction.

$$Correlation = \frac{\sum_{i,j}[(ij)P_d(i,j)] - m_X m_Y}{s_X s_Y} \tag{5.15}$$

where m_x, m_y are means and σ_x, σ_y standard deviations:

$$m_X = \sum_i i \sum_j P_d(i,j) \qquad \sigma_X = \sum_i (i-m_X)^2 \sum_j P_d(i,j)$$

$$m_Y = \sum_j j \sum_i P_d(i,j) \qquad \sigma_Y = \sum_{ji} (j-m_Y)^2 \sum_i P_d(i,j)$$

The most usual values for the displacement vector are (0,1), (1,0), and (1,1). Coinciding with the autocorrelation feature, *d=(0,1)* will be the selected one.

5.2.3. FOURIER FEATURES

The Fourier transform describes spatial frequencies extremely well[1-28]. It is ideally suited for describing the directionality of periodic or almost periodic two-dimensional patterns in an image. These global texture patterns, while being easily distinguishable as concentrations of high-energy bursts in the spectrum, are generally quite difficult to detect using spatial methods because of the local nature of these techniques. In opposition to the co-occurrence matrices, the classification measures from the Fourier spectrum of image segments is a different statistical technique that uses the absolute differences between pairs of grey levels in an image region.

Average values of energy in specific wedges and rings of the Fourier spectrum can be used as textural description features. Features evaluated from rings reflect coarseness of the textures; high energy in large radius rings is characteristic of fine textures (high frequencies), while high energy in small radii is characteristic of coarse textures (low spatial frequencies). Apart from these measures extracted from rings and wedges, there is another approach for obtaining Fourier features related to the amplitudes of the spatial frequencies, that is the one we will use here.

This classification measures from the Fourier spectrum of image regions require the calculation of the fast Fourier transform (FFT) for each area and the definition of features in terms of the amplitudes of the spatial frequencies. The discrete Fourier transform *F(n,m)* of a digitized image segment *I(j,k)* of size *NxN* is defined by

$$F(n,m)=\frac{1}{N^2}\sum_{j=0}^{N-1}\sum_{k=0}^{N-1}I(j,k)\exp\left[-\frac{2\pi i}{N}(jn+km)\right] \tag{5.16}$$

where *(i=√-1), (n,m)* are the discrete spatial frequencies, and *(j,k)* are the pixel positions.

The set of features based on the power spectrum consists of three statistical measures. If $|F(n,m)|$ is the matrix containing the amplitudes of the spectrum and N^2 is the number of frequency components then these measures, proposed by Augusteijn et al.[1-29] are given by:

$$Max.peak = \max[|F(n,m)| : (n,m)^{1} \quad (0,0)] \tag{5.17}$$

$$Average = \frac{\sum_{n,m} |F(n,m)|}{N^2} \tag{5.18}$$

$$Energy = \frac{\sum_{n,m} |F(n,m)|^2}{\max(|F(n,m)|^2) N^2} \tag{5.19}$$

5.2.4. GREY-LEVEL IMAGE FRACTAL FEATURES[2-6]

A detailed description of the relevance of applying fractal measures as a possible means to detect the presence of diseases in the pulmonary parenchyma was outlined in Chapter 2. Fractal geometry is becoming increasingly more important in the study of image characteristics, since it allows quantification of structure or pattern across many spatial or temporal scales and therefore could be useful in many biomedical applications[1-30].

The fractal dimension is a measure of randomness, as well as of complexity. In the human body, we have some complex structures, such as the brain and the bronchial tubes, whose fractal dimension is liable to be calculated, and to be used as a tool for the quantification of structure changes.

There are numerous methods available to estimate parameters from images of fractal surfaces. The images' fractal dimension may be measured, for example, by use of the local second-order statistics (interpixel differences change with distance) or the Fourier power spectrum[1-31]. In our study two techniques to estimate the fractal dimension have been developed.

The first one, explained below, comprises a study based on greyscale images of the lungs, yielding a range of values between 2 and 3, which is what we should expect because the fractal geometry can be considered as an extension of Euclidean geometry, and allows non-integer dimensions. Perfectly flat planes have Euclidean dimension of 2, but real surfaces have a dimension greater than 2, which is our case.

The second one, on the other hand, is based on a binary version of the greyscale images, which is a general and powerful technique if the conversion is carried out conveniently and will be discussed in next section.

Many surveys comparing methods for measuring fractal dimensions have been conducted[1-32]. The one chosen here uses the difference between pixels at different scales, coming up with a measure linearly related to the fractal dimension, called the Hurst exponent (H), as follows:

1. For k=1,…, N , a measure of the interpixel grey level difference is obtained:

$$GD(k) = \frac{\sum_{i=1}^{N}\sum_{j=1}^{M-k-1} |I(i,j) - I(i,j+k)| + |I(j,i) - I(j+k,i)|}{2N(N-k-1)} \tag{5.20}$$

   where N is the window size, and k is the increasing distance.

2. In practice, H can be estimated from the slope of the log-log plot:

$$\log(GD(k)) = \log T + H \log(k) \tag{5.21}$$

   where T is a constant, and H the Hurst exponent

3. The fractal dimension (FD) can finally be obtained by:

$$FD = 3 - H \tag{5.22}$$

As an example, in figure 5.1 a zoom into a 16x16 region extracted from the lung of figure 4.2.a and the log-log plot are shown, and its fractal dimension has been calculated following this procedure.

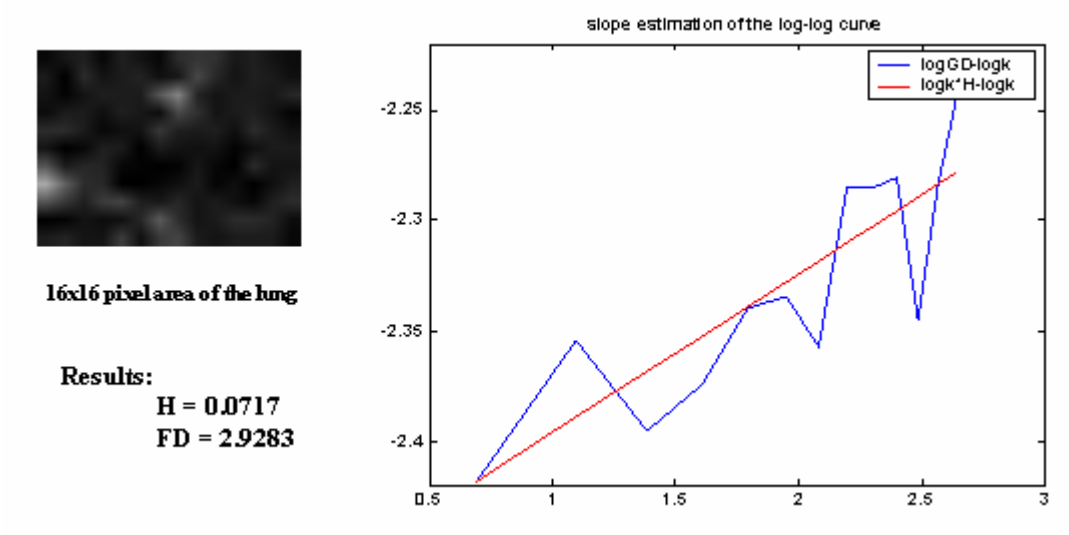

Figure 5.1. Grey-level differences method for measuring the fractal dimension

5.2.5. BINARY FEATURES[2-7]

Binary images are images that have been quantised to two values, usually denoted 0 and 1, but often with pixel values 0 and 255, representing black and white. They are typically obtained by thresholding a grey level image. However, choosing a threshold can be difficult, and is even considered by some[1-33] to be a "black art". Most approaches make use of the histogram of the number of times each grey level occurs in the image.

In our case, two different thresholds have to be selected, one for differentiating abnormally dark areas from normal ones, and the other for differentiating abnormally bright areas. The histograms for abnormal and normal regions are not perfectly bimodal and will overlap, making it more difficult to find good values for the thresholds since there will be no clear way of choosing them.

An empirical approach has been utilized here to find the thresholds that are better for the conversion from a greyscale image to a binary image. Repeated observations of different regions considered unhealthy and healthy using several images were carried out.

In next figure, we present some samples of these observations, extracted again from figures 4.2.a. and 4.2.b:

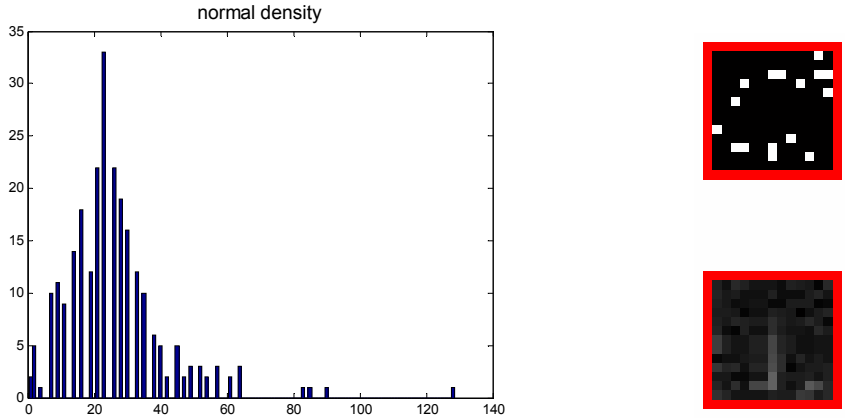

5.2.a. NORMAL DENSITY AREA

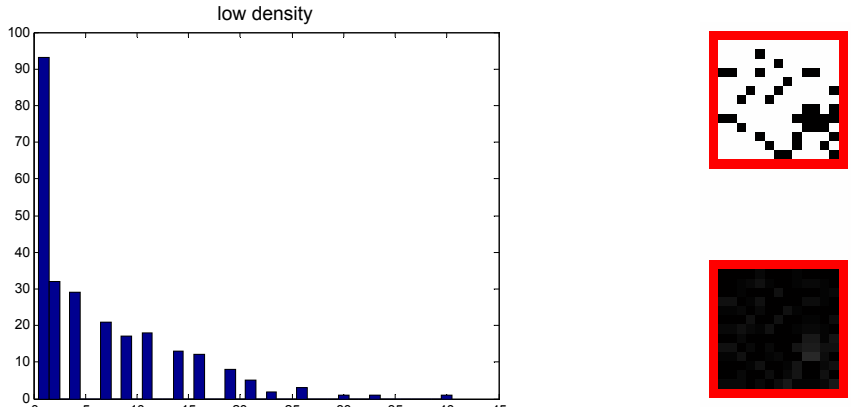

5.2.b. LOW DENSITY AREA

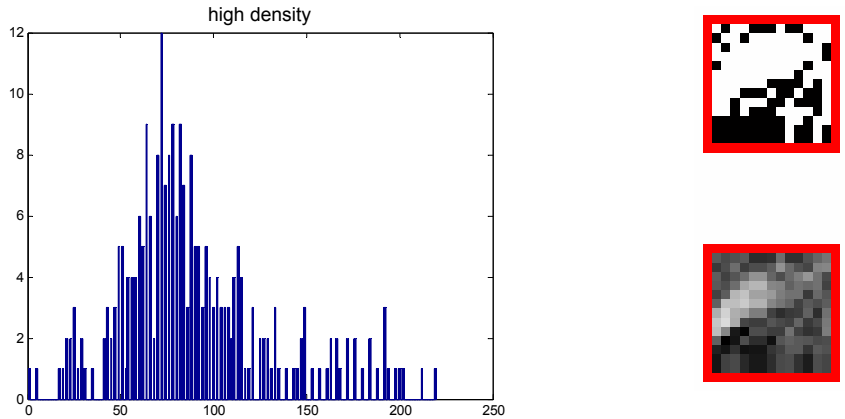

5.2.c. HIGH DENSITY AREA

Figure 5.2. Threshold selection Process

The histograms above show the grey level distributions in these three cases: low, normal and high density areas. The window size, as usual, is 16, and we have done a zoom in on the regions to see them clearer. Along with each region, we also present the binary image as the result of applying the thresholds.

We have to distinguish thus both cases to understand the results. For the high threshold, pixels with a grey level above the threshold are set to 1 (equivalently 255, white), whilst the rest are set to 0, black. On the other hand, for the low threshold, pixels with a grey level below it are set to 1, and the rest to 0. In figure 5.2.a, most of the pixels are black, which is an indication of normality; however in figures 5.2.b and 5.2.c, most of them are white, indicating abnormality.

There is another consideration that we can get out of the previous histograms, which is that the low threshold appears to be much more sensitive to changes than the high threshold. This is due to the fact that the histograms corresponding to low density are more concentrated in the lowest positions of the grey level scale, while the other histograms are more disperse and we can have a larger range of selection for the threshold. Furthermore, there are sometimes healthy and unhealthy parts in the same region, which makes it even more difficult to obtain accurate results. After many trials and observations, we have come to the result that the optimal range of values for the low threshold is [10-15], and [75-85] for the high threshold (the whole range of grey levels is [0-255]), and they have eventually been set to the values “12” and “80” respectively.

Once the conversion from a greyscale image to a binary image has been explained, we can carry out the analysis of some binary measures. The parameters that are used to describe them are various statistical measures, which may be divided into two distinct classes: geometrical descriptors and topological descriptors.

We will narrow the study to the fractal dimension of the binary image, and to some geometrical descriptors, which are: area, compactness, eccentricity, and the first, second and third invariant moments.

- **Area:**

It is the simplest and most natural property of a region, and it is given by the number of pixels of which the region consists, i.e., the number of pixels that have been set to 1. A normalized version is presented below:

$$Area = \frac{\sum_{n=0}^{N-1}\sum_{m=0}^{N-1} b(n,m)}{N^2} \tag{5.23}$$

where *b(n,m)* represents a pixel of the binary image, and N^2 is the number of pixels in the region.

- **Compactness:**

Compactness is a popular shape description characteristic invariant with respect to scaling and rotation.

$$Compactness = \frac{(4N)^2}{Area} \tag{5.24}$$

where *4N* is the region border length.

- **Eccentricity:**

This descriptor is sometimes called elongation, and it is invariant with respect to scaling, rotation and translation. This is the ratio of the maximum length of line or chord that spans the region to the minimum length chord. We can also define this in terms of moments as follows:

$$Eccentricity = \frac{\mu_{20} + \mu_{02} + \sqrt{(\mu_{20} - \mu_{02})^2 + 4\mu_{11}^2}}{\mu_{20} + \mu_{02} - \sqrt{(\mu_{20} - \mu_{02})^2 + 4\mu_{11}^2}} \tag{5.25}$$

where μ_{20},μ_{02}, and μ_{11} are invariant moments, as explained below.

- **Moments:**

Moments can be used for grey level region description and also for binary description[1-34]. A moment of order *(p+q)* is dependent on scaling, translation, rotation and is given by:

$$m_{pq} = \sum_{n,m} n^p m^q b(n,m) \tag{5.26}$$

where *(n,m)* are the pixel coordinates.

Translation invariance can be achieved if we use the central moments:

$$\mu_{pq} = \sum_{n,m} (n - x_c)^p (m - y_c)^q b(n,m) \tag{5.27}$$

where x_x and y_c are the coordinates of the region's centre of gravity (centroid) which can be obtained using the following relationships

$$x_c = \frac{m_{10}}{m_{00}}; y_c = \frac{m_{01}}{m_{00}} \tag{5.28}$$

where m_{00} represents the binary area.

The central moments of order 3 can be expressed as follows:

$$\begin{aligned}
&\mu_{00}=m_{00}, && \mu_{11}=m_{11}-y_c m_{10} \\
&\mu_{10}=0, && \mu_{30}=m_{30}-3x_c m_{20}+2m_{10}x_c^2 \\
&\mu_{01}=0, && \mu_{12}=m_{12}-2y_c m_{11}-x_c m_{02}+2y_c^2 m_{10} \\
&\mu_{20}=m_{20}-x_c m_{10}, && \mu_{21}= m_{21}-2x_c m_{11}-y_c m_{20}+2x_c^2 m_{01} \\
&\mu_{02}= m_{02}-y_c m_{01}, && \mu_{03}= m_{03}-3y_c m_{02}+2y_c^2 m_{01}
\end{aligned} \tag{5.29}$$

Scale invariant features can also be found in the normalized central moments h_{pq}, defined as:

$$h_{pq} = \frac{\mu_{pq}}{\mu_{00}^{\gamma}}; \gamma = \frac{p+q}{2} + 1, p+q = 2,3,... \quad (5.30)$$

From the second and third moments, a set of seven invariant moments can be derived. The first, second and third invariant moments are given by

$$f_1 = h_{20} + h_{02} \quad (5.31)$$

$$f_2 = (h_{20} - h_{02})^2 + 4h_{11}^2 \quad (5.32)$$

$$f_3 = (h_{30} - 3h_{12})^2 + (3h_{21} - h_{03})^2 \quad (5.33)$$

- **Binary Fractal Dimension:**

This is a different approach to study the fractal dimension of an image. In fact, since we will work with binary images obtained from the application of certain brightness thresholds to the grey scale image, we cannot expect that the resulting pattern is a real fractal object. However, a "fractal dimension" can be calculated and we can work with it as a measure of information and complexity of the binary structure.

Some surveys have been developed to demonstrate the reliability of using fractal measures derived from binary images. Pentland[1-31] demonstrated that fractal based segmentation converting a grey image to binary image according to local fractality yielded higher classification accuracy than other methods such as correlation statistics, co-occurrence statistics, and texture energy statistics.

The program used to find the fractal dimension (courtesy of Dr. Jonathan Rossiter), utilizes the box-counting method described in Chapter 2, yielding high values when we test it on abnormal areas (figure 5.2.a and b), and low values for regions considered as normal(figure 5.2.a). The range of possible

values is [0-2]. If the region contains only zeros, then the fractal dimension is 2, and if it contains only ones, it yields a result of 0.

A large number of texture features have been proposed. But these features are not independent as pointed out by Tomita and Tsuji[1-35]. The relationship between the various statistical texture measures and the input image is summarized in the next figure:

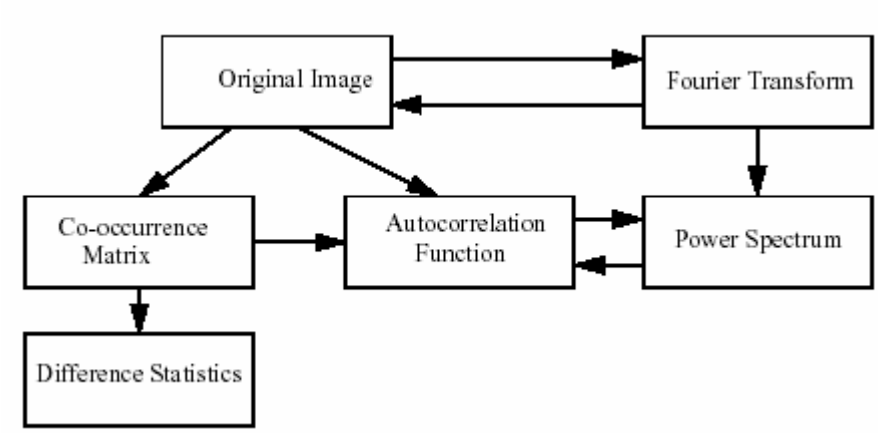

Figure 5.3. The interrelation between the various second-order statistics and the input image

5.3. INITIAL TRIALS.

The next step in our research has to be the evaluation of each feature as a method of differentiating normal lung tissue regions from abnormal ones. A comparison amongst all the features belonging to the first set has to be carried out in order to obtain a final reduced subset with the ones that yield the best performances.

To make the comparison between both classes of regions, healthy and unhealthy, we have plotted histograms and box-shaped density functions. The histograms have been normalized in order to make a clearer comparison, making the sum of all the grey level occurrences of each class equal to one. This is though a qualitative and empirical approach, and the final quantitative method for the current project will be based on neural networks.

However, there are certain tools in MATLAB that we could use to make a first attempt to evaluate the performance of each feature in a numerical way. This function is *anova1*, an extended version of the *boxplot* function that was initially described in Chapter 4. It performs a comparison between both sets of data (healthy/unhealthy) under the null hypothesis that all the samples in both groups are drawn from the same population or from different populations with the same mean, and returns a p-value. If this value is near zero, this casts doubt on the null hypothesis and suggests that both means are significantly different. Therefore, the lower the p-value, the higher the discriminatory power of the feature under study.

The test makes the following assumptions about both populations:

- They have equal variance
- They are normally distributed
- Both observations are mutually independent

The test is known to be robust to modest violations of the first two assumptions, and we will illustrate these first trials with these results.

A further word about the features in the edges is needed before starting off the following process. There are some features that yield poor and unreliable results if they are calculated in regions that include pixels belonging to the outer parts of the lungs. The reasons of this are explained in next section, but here we can see an example in figure 5.4, where the histograms for the energy of the co-occurrence matrix corresponding to the lungs from image 4.2.a are plotted. We can see that the range of values for the edges concentrates in [0-0.12], while the range for the centres is [0-0.35], plus the distributions clearly vary. For this reason, if we mix in the same study edges and centres, the results will appear biased. Hence, we will only work here with the inner regions, which comprise the most representative information about the state of the lungs, and the study of the boundaries will be done separately in Chapter 6

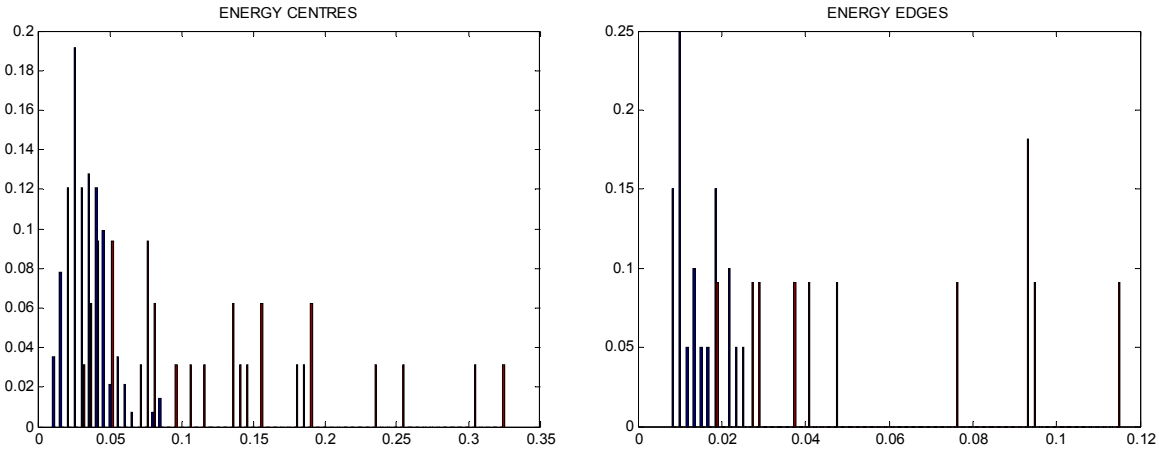

Figure 5.4. Co-occurrence matrix energy in edges and centres. In red: Unhealthy class; in blue: healthy class.

A study was carried out using the two main sample images that we have seen so far, which are those from figure 4.2, because the diseases appear very distinctly in them. The first one presents an emphysematous condition, and the second one is a lung with fibrosis. In the next Chapters more images will be used and more tests developed.

5.3.1. ABNORMALLY LOW PARENCHYMAL DENSITY: *EMPHYSEMA*

The areas marked with red squares in figure 5.5 present lower than normal greyscale values. The diagnosis is the presence of widespread emphysema with scarring in both lungs. We can see that most of the bullae are clustered around the inner borders of the lungs.

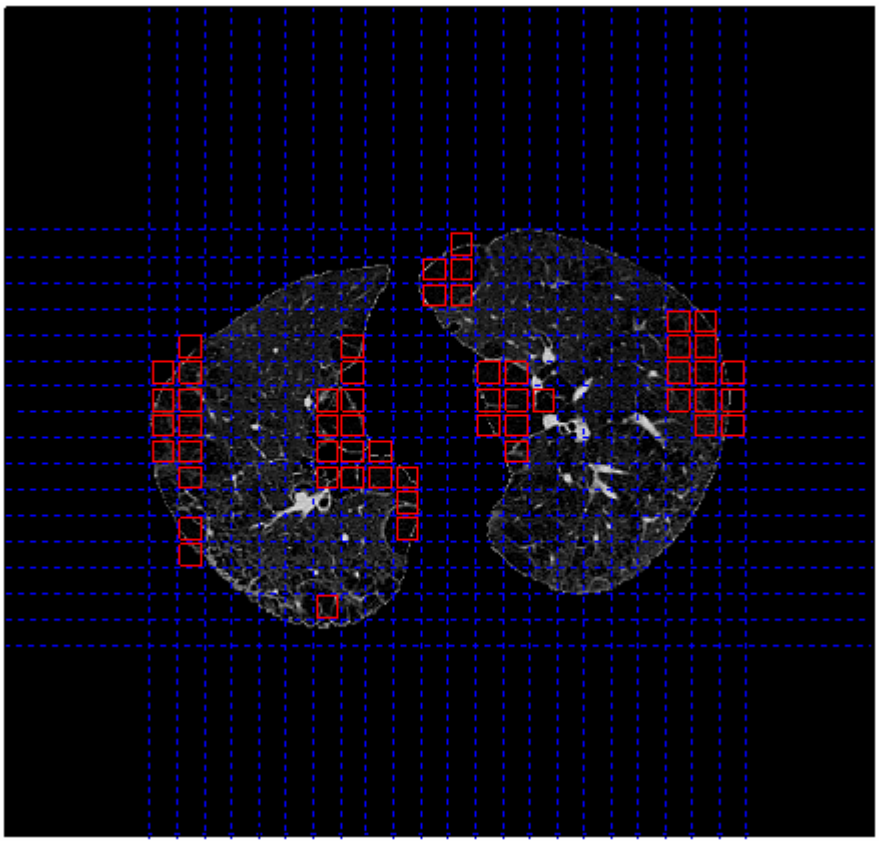

Figure 5.5. Widespread emphysema

We will proceed to evaluate each feature, using the MATLAB function *anova1*. Table 5.1 comprises the results obtained for the *p-values* in order of better performance. As we can see, it is not a very accurate means of coming up with an optimal selection, but it gives a general idea about which features will be finally included in the optimal subset. Actually, at the end of Chapter 6, the six selected features are: mean, percentile, autocorrelation, Fourier max. peak, binary fractal

dimension, and binary area, which present here a *p-value* of 0, so the results below are not very accurate but they are robust.

Order	feature	p-value
1	Mean	0
2	Percentile	0
3	Autocorrelation	0
4	Fourier max. peak	0
5	Binary fractal dimension	0
6	Binary area	0
7	Fourier energy	0
8	Max. co-occurrence prob.	0
9	Entropy	0
10	Compactness	0
11	Contrast	0
12	First moment	0
13	Eccentricity	0.0023
14	Max. grey value	0.0063
15	Range	0.0064
16	Hurst exponent	0.0143
17	Grey-level fractal dimension	0.0143
18	Third moment	0.0229
19	Correlation	0.0268
20	Standard Deviation	0.0568
21	Second moment	0.0648
22	Inverse diff. moment	0.0714
23	Fourier average	0.5151
24	Min. grey value	0.6352
25	Co-occurrence matrix energy	0.7944

Table 5.1. Features evaluation for abnormally low tissue density with *anova1*

The histograms and *boxplots* of four of the best features, percentile, binary area, binary fractal dimension and Fourier maximum peak, are shown in the figure below. For the binary features, the thresholds utilized are 12 for the low threshold and 255 for the high threshold. The class unhealthy is represented in red in the histograms, and the class healthy in blue.

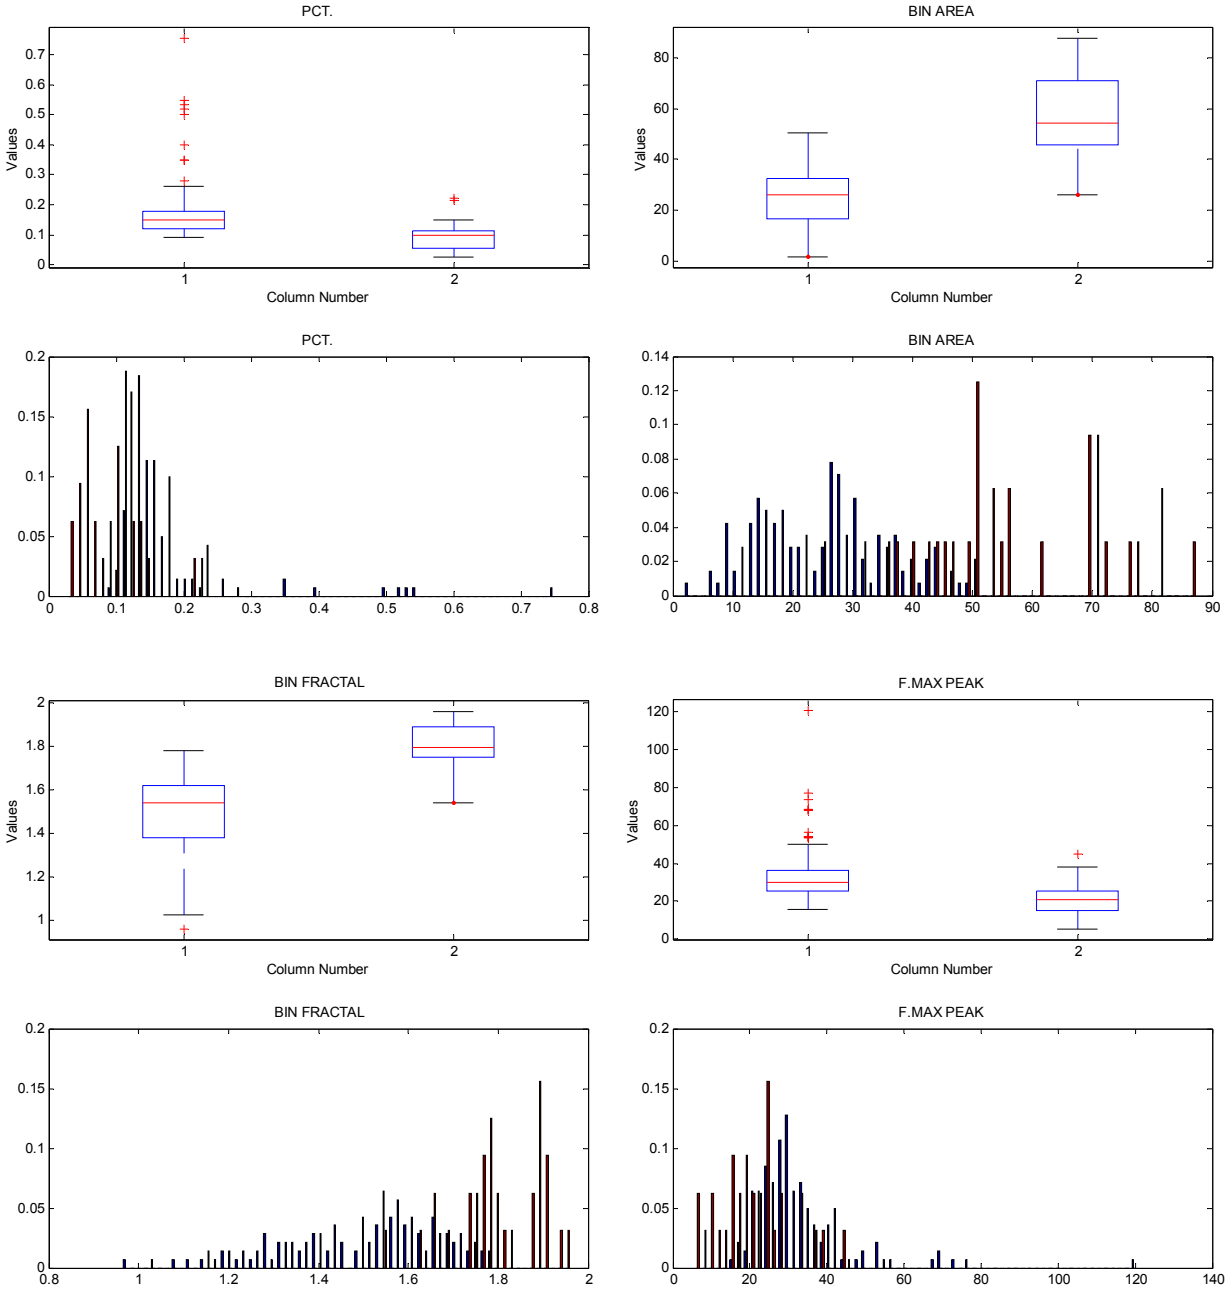

Figure 5.6. Histograms and boxplots.

5.3.2. ABNORMALLY HIGH PARENCHYMAL DENSITY: *FIBROSIS*

We will follow the previous procedure for the lungs of figure 5.7, which present higher than normal greyscale values, corresponding to fibrosing alveolitis, especially relevant in the inferior part of both lungs.

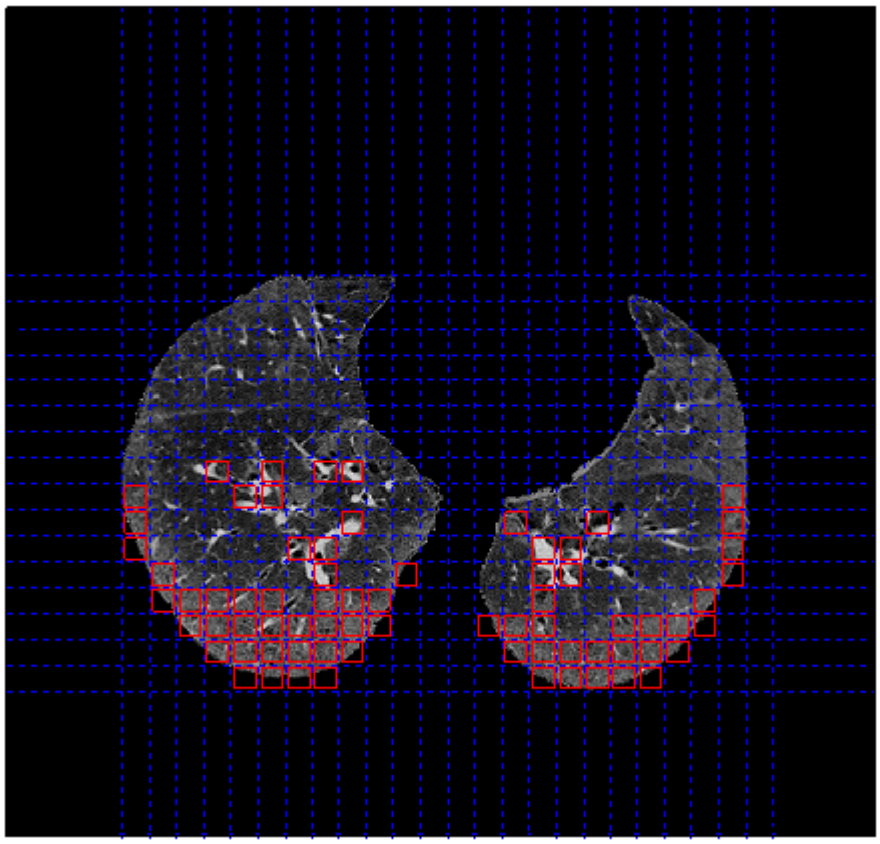

Figure 5.7. Fibrosing alveolitis

Table 5.2 comprises the results obtained for the *p-values* in order of better performance. The thresholds for the binary features, in this case, are 0 and 80. In Chapter 6 we will discover that the optimal subset of features coincides with the optimal subset for detecting emphysema. Their *p-value* is 0 here, which proves again the robustness of this method. In figure 5.8, the histograms and *boxplots* for the percentile, binary area, binary fractal dimension and Fourier maximum peak are calculated.

Order	feature	p-value
1	Mean	0
2	Percentile	0
3	Autocorrelation	0
4	Fourier max. peak	0
5	Binary fractal dimension	0
6	Binary area	0
7	Co-occurrence matrix energy	0
8	Max. co-occurrence prob.	0
9	Entropy	0
10	Compactness	0
11	Contrast	0
12	Standard deviation	0
13	Correlation	0
14	Max. grey value	0
15	Min. grey value	0
16	Inverse diff. moment	0
17	Fourier average	0
18	First moment	0.0001
19	Range	0.0011
20	Eccentricity	0.0233
21	Fourier energy	0.0341
22	Second moment	0.1790
23	Third moment	0.1965
24	Hurst exponent	0.7507
25	Grey-level fractal dimension	0.7507

Table 5.2. Features evaluation for abnormally high tissue density with *anova1*

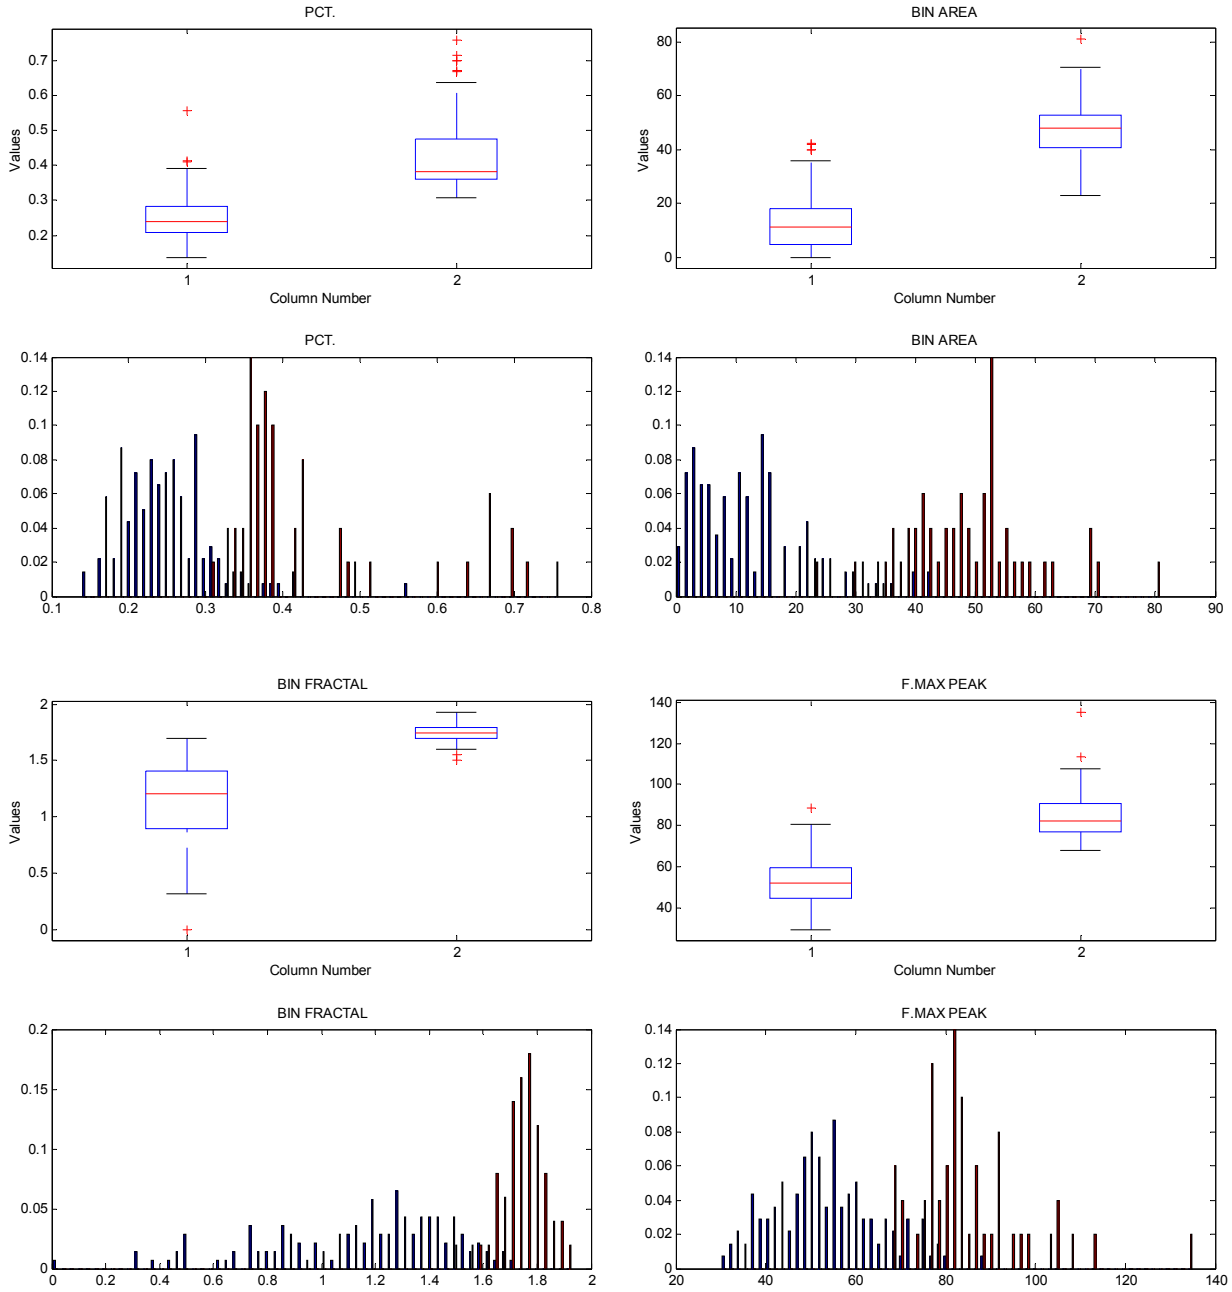

Figure 5.8. Histograms and boxplots.

5.4. A SPECIAL CASE: FEATURES IN EDGES

We have already mentioned the problem of this matter. We can divide the set of features into two groups. The first group is composed by the features that yield the same result regardless of the position of the region under study, which can thus belong to the edge or to the inner parts of the lung. This group comprises the following 8 features: mean, maximum and minimum grey value, range, standard deviation, percentile, binary area and compactness. The second group is composed by the rest of the features, and the reason for their bad behaviour in the contours is that they need a complete square-shaped structure due to their own definition. For example, for the co-occurrence matrices, we need to use a rectangular region in order to achieve some results, and the same thing happens with the rest of the features. This is why if edges and centres were treated together without any differences, the results would appear biased.

Some solutions have been proposed to this problem. Firstly, all the pixels that do not belong to the lung, i.e., the image background, were marked with a label (a number out of the greyscale range, such as -2). These labels were replaced with random noise inside the greyscale range. The idea was to push all the peripheral regions in the same way with noise, so that certain comparisons could be carried out amongst these areas in order to differentiate the abnormalities in them. This approximation only has a sense with the second group of features that we mentioned before, and would not be applied to the first group of features. Another way we used to deal with the problem will be treated in next Chapter, where we propose to train a different neural network specifically only with edges. Notwithstanding, we have included those regions with at least an 85% of relevant information as belonging to the inner parts of the lungs and therefore they have been treated accordingly. On the other hand, since the pixel areas we are working with is 16x16 which is not very large, we have only considered regions that contain at least a 40% of information, discarding those that cannot really contribute to the diagnosis due to their lack of structure.

5.5. SUMMARY AND CONCLUSIONS

In this Chapter we have selected 25 texture features to investigate their discriminatory potential for detecting lung diseases. Some initial trials have been carried out utilizing sample images with abnormally low and abnormally high tissue density, presenting clear symptoms of emphysema and fibrosing alveolitis respectively.

We have obtained some qualitative results in the observation of the histograms and density function, and a first approach of a quantitative method for measuring the performance of the features has been proposed, showing little accuracy but high robustness.

We can finally conclude that the treatment of the features in edges is a difficult problem that will impoverish the performance of our diagnostic tool.

6. NEURAL NETWORKS. THE LEARNING PROCESS

6.1. INTRODUCTION[1-36, 1-37, 1-38, 1-39]

Artificial intelligence has become an important element for investigation in tasks such as pattern recognition, computer vision and medical applications. There are many AI techniques (artificial neural networks, genetic algorithms, case-based reasoning systems), the most popular and historically successful of them in medical decision making, has been the artificial neural networks[1-40]. They have already outperformed conventional techniques on a number of problems, so it is realistic to expect that neural computing could be the dominant approach to computing in this century.

They are useful to carry out complex analyses using an empirical training procedure without having to describe a formal model, and are able to capture the underlying relationships between the textural features extracted from the images, which will be the input variables of our neural network.

This Chapter presents an overall view of the neural networks architectures, narrowing the focus to the *cascade-forward* architecture and also the *standard back-error propagation* or *BP algorithm*, which is a supervised learning algorithm for multi-layered networks. The MATLAB neural network toolbox provides the necessary functions to develop these strategies, and will be reviewed to fully understand the procedure.

The sequential forward and the sequential backward algorithms for feature selection are explained, and applied to obtain the two optimal subsets of features corresponding to the two main diseases under study, emphysema and fibrosis. The

'leave-one-out cross-validation' method for the training/validation process is carried out along with these sequential feature selection algorithms, and we will eventually come up with the resultant performances.

A differentiation between both diseases and also between edges and centres is necessary; hence four cascade-forward neural networks have been trained to yield satisfactory results.

6.1.1. NEURAL NETWORKS ARCHITECTURES

The inspiration for neural network theory comes from biological neural systems, such as the human brain and their basic component: the neuron. This field has a history of six decades, since McCulloch and Pitts outlined in 1943[1-41] the first model of an artificial neuron, but has seen solid application only in the past twenty years, when several people in the mid-1980s found a learning algorithm, called the *back-propagation algorithm*, that could adjust the weights in multi-layer nets, and the field is still developing rapidly.

A neuron with a single scalar input and no bias appears on figure 6.1, and another one with a scalar bias on its right. The scalar input p is transmitted through a connection that multiplies its strength by the scalar weight w, to form the product wp, again a scalar. Here the weighted input wp is the only argument of the transfer function T, which produces the scalar output o. The neuron on the right has a scalar bias, b, that is added to the product wp, and the resulting net input is $wp+b$. The transfer function is typically a step function, a linear function or a sigmoid function, and produces the output $o=T(wp+b)$. The weight and the bias are parameters liable to be adjusted so that the network exhibits some desired or interesting behaviour. Thus, we can train the network to do a particular job by adjusting these parameters, or perhaps the network itself will adjust these parameters to achieve some desired end.

More complex structures can be derived from this first model. A neuron can have an input vector $\boldsymbol{p}=[p1, p2, ..., pR]$ instead of a scalar input, and the weights $w1, w2, ..., wR$ correspond to a single row matrix $\boldsymbol{W}$. The dot product $\boldsymbol{Wp}$ plus the scalar bias

b form the net input *n= w1p1+w2p2+...+w3p3+b*, which is the argument of the transfer function *T*.

Two or more of the neurons can be combined in a layer, and a particular network could contain one or more such layers. First consider a single layer of *S* neurons with *R* elements in input vector. In this network, each element of the input vector ***p*** is connected to each neuron input through the weight matrix ***W***. The *ith* neuron has a summer that gathers its weighted inputs and bias to form its own scalar output *ni*. The various *ni* taken together form an *S*-element net input vector ***n***. Finally, the neuron layer outputs form a column vector ***o***.

The input vector elements enter the network through the weight matrix ***W***. The row indices on the elements of matrix ***W*** indicate the destination neuron of the weight, and the column indices indicate which source is the input for that weight. Thus, the indices $w_{i,j}$ say that the strength of the signal from the *jth* input element to the *ith* neuron is $w_{i,j}$.

$$W = \begin{bmatrix} w_{11} & w_{12} & \dots & w_{1R} \\ w_{21} & w_{22} & \dots & w_{2R} \\ \\ w_{S1} & w_{S2} & & w_{SR} \end{bmatrix}$$

Finally, several layers can be combined in a more complex architecture, presenting basically an input layer, an output layer and one or more hidden layers.

We will eventually come up with neural networks that have one single hidden layer and only one output neuron as the architecture that better suits our problem We can see its general scheme in figure 6.1.

Each layer has a weight matrix ***W***, a bias vector ***b***, and an output vector ***o***. We have appended in the figure the number of the layer as a superscript to the variable of interest.

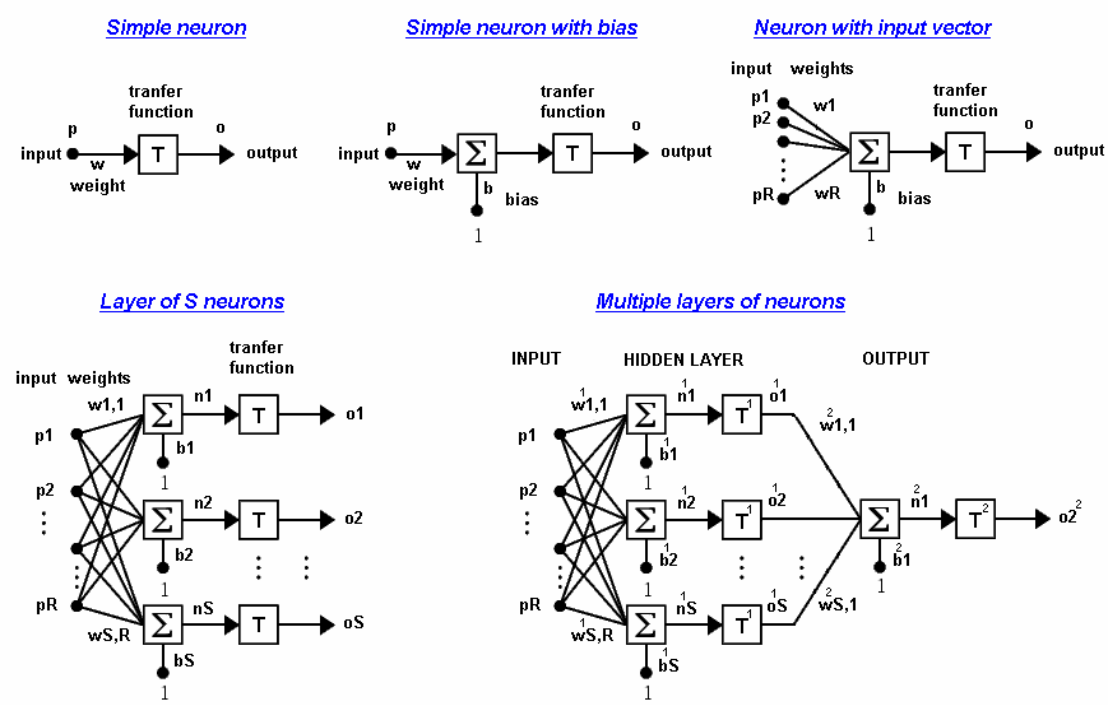

Figure 6.1. Neural Network Architectures

6.1.2. CASCADE-FORWARD NETWORKS

One basic split is between *recurrent* (*feed-back*) and *feed-forward* or *cascade-forward architectures*. A recurrent processing architecture includes direct or indirect loops of connections; the current state of each single layer of nodes is fed through an updating function of the nodes to produce a new state, which is in turn feedback into the nodes to produce a new state again, etc. The alternative type of architecture is one that have asymmetric connection matrices $w_{i,j}$ in which information is presented and proceeded uni-directionally from input to output. This type of nets resembles conventional methods of statistical pattern recognition much more closely. The interconnections between nodes in adjacent layers start with random weights that are changed during an iterative process of neural net learning which may be unsupervised or supervised, as we will describe below.

We will use cascade-forward networks, belonging to the latter type, which can be subdivided into two main classes: *single-layered* and *multi-layered networks*. Single-layered structures consist only of an input layer and an output layer, and cannot

perform certain classification tasks or certain logic operations, such as the *XOR* problem, because they can solely solve linearly separable problems. As we discussed in previous Chapters, we do not have separable classes for the features, so we need to complicate the architecture by adding more layers until the results are satisfactory. In practice, one single hidden layer should suffice for most of the datasets, and can compute most of the functions. Problems that require 2 hidden layers are obscure and involve datasets generated by discontinuous functions[1-42]. The optimal number of nodes for an errorless classification is not known beforehand; hence several trials are necessary to find it.

Cascade-forward network learning is based on the *back-propagation algorithm*[1-43], that calculates a set of weights $w_{i,j}$ from a training set of examples, where the weight $w_{i,j}$ represents an interconnection between node i and j in the next layer. In next section, this algorithm is further detailed.

6.1.3. LEARNING PROCESS. SUPERVISED LEARNING & BACK-PROPAGATION ALGORITHM

The property that is of primary significance for a neural network is the ability of the network to learn from its environment, and to improve its performance through learning. The improvement in performance takes place over time in accordance with some prescribed measure. A neural network learns about its environment through and interactive process of adjustments applied to its weights and bias levels. Ideally, the network becomes more knowledgeable about its environment after each iteration of the learning process.

The current learning process takes place under the tutelage of a teacher, and it is called *supervised learning*. Figure 6.2 shows a block diagram that illustrates this form of learning. The teacher has knowledge about the environment, with that knowledge being represented by a set of *input-output* examples. In our case, the teacher is a human observer making decisions about the state of the lungs, and the input-output pairs correspond to *feature_vector-diagnosis* examples. The feature vectors comprise the information about the environment; in this case the environment consists of 16x16

pixel areas of lung images, and the features vectors are integrated by several measures that capture the textural information of these lung segments. The targets or desired responses are divided into two classes of diagnosis, healthy and unhealthy, which correspond to a binary classification.

Let us suppose now that the teacher and the neural network are both exposed to a training vector drawn from the environment. By virtue of the built-in knowledge, the teacher is able to provide the neural network with a desired response for that training vector. Indeed, the desired response represents the optimum action to be performed by the neural network. The net parameters are adjusted under the combined influence of the training vector and the error signal. The error signal is defined as the difference between the desired response and the actual response of the network. This adjustment is carried out iteratively in a step-by-step fashion with the aim of eventually making the neural network emulate the teacher; the simulation is presumed to be optimal in some statistical sense. In this way knowledge of the environment available to the teacher is transferred to the neural network through training as fully as possible. When this condition is reached, we then dispense with the teacher and let the net deal with the environment completely by itself.

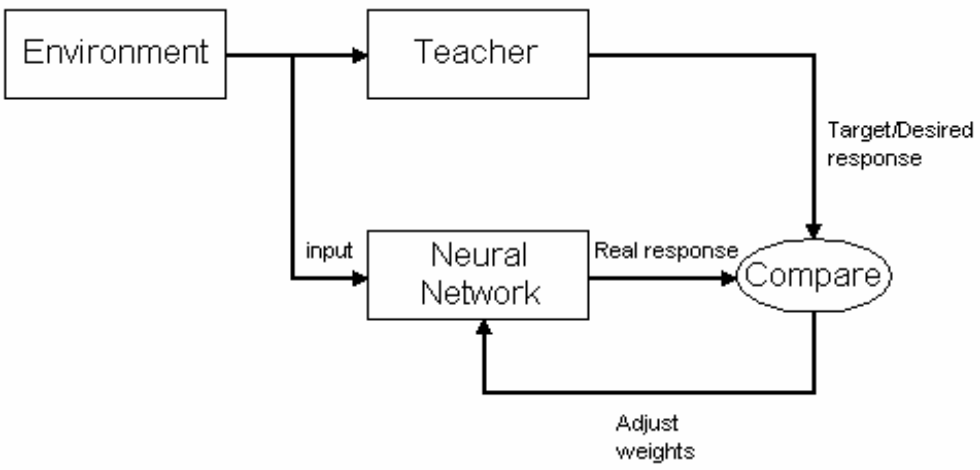

Figure 6.2. Supervised training.

There are several supervised learning methods, such as quickdrop, conjugate gradient methods, quasi-Newton methods, etc. We will use a gradient descent algorithm for multi-layered networks, the *standard back-error propagation* or *BP* algorithm.

Back-propagation works by feeding the inputs for a case through an initially random network and comparing the output to the desired output for that case. It then back-propagates the error through the network, incrementally adjusting the weights along each connection to minimize the mean squared error between the actual network output and the desired output. This has been shown to achieve a gradient descent of the error surface[1-44]. Two distinct passes of computation are distinguished, the first pass is referred to as the forward pass, and the second one is referred to as the backward pass. In the forward pass the synaptic weights remain unaltered throughout the network, and the function signals of the network are computed on a neuron-by-neuron basis. The backward pass, on the other hand, starts at the output layer by passing the error signals layer by layer, computing the local gradient for each neuron. The algorithm is described below.

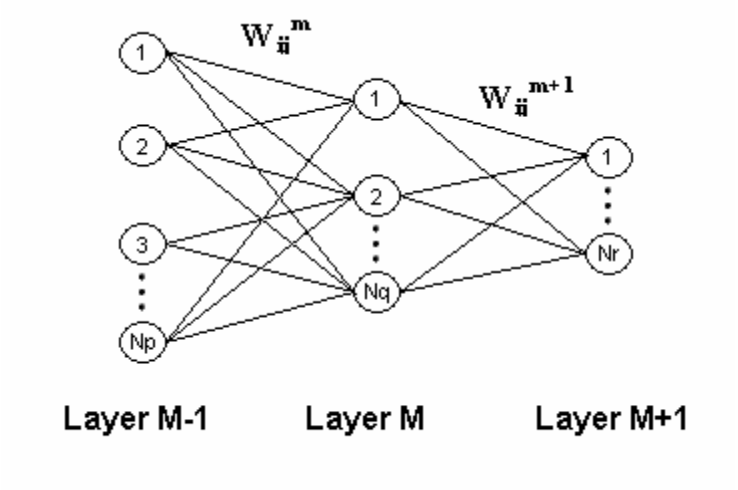

Figure 6.3. Cascade-forward network to illustrate BP algorithm

Let us consider the network illustrated schematically above (we will change slightly the notation used in section 6.1.1. to clarify the algorithm). Suppose we consider a network with an input layer (which we label $m=0$) and q processing layers

(labelled *m=1, 2, ..., q*). o_i^m is the output on layer *m*. Let W_{ij}^m be the weights leading into layer *m* i.e. connecting layers *m-1* and *m*. Then, the algorithm works as follows:

1. Weights on the connections between neurons are randomly initialized. Small values are used to avoid early saturation of the activation function.

2. An input pattern x_k is presented to the input layer (*m=0*) and the output on the final output layer (*m=q*) is found. This output is o_i^q.

3. Compute the error terms δ_i^q on the output layer from

$$\delta_i^q = T'(f_i^q)(y_i - o_i^q) \tag{6.1}$$

   where f_i^q are the inputs for last layer *q*, T' is the derivative of the transfer function, o_i^q the outputs on the ouput layer and y_i is the corresponding target for input x_k.

4. Find the errors δ_j^{m-1} for the preceding layer by recursively using

$$\delta_j^{m-1} = T'(f_j^{m-1}) \sum_i W_{ij}^m \delta_i^m \tag{6.2}$$

5. Find the weight changes for the weights leading into each layer *m*

$$\Delta W_{ij}^m = h \delta_i^m o_j^{m-1} \tag{6.3}$$

   and update the weights according to

$$W_{ij}^{new} = W_{ij}^{old} + \Delta W_{ij} \tag{6.4}$$

6. Go to step (2) and present the next pattern in the pattern set; until the pattern set is exhausted.

7. Stop if the mean squared error performance function (*MSE*) *E* is less than the selected tolerance, or re-present the entire pattern set otherwise.

The constant h in (6.3) defines the learning rate of the network. With a small learning rate, the back-propagation algorithm will minimize the cumulative mean square error between the target outputs and the actual network outputs as training progresses. In most real-world problems, unfortunately, a small learning rate will result in the network settling in a false local error minimum, and some convergence problems can appear. On the contrary, if the value is too large, oscillations can occur. In order to avoid such oscillations and to avoid get caught in a shallow minimum, it is common practice to add a *momemtum term*:

$$\Delta W_{ij}' = -h\frac{\partial E}{\partial W_{ij}} + m_c \Delta W_{ij}^{prev} \tag{6.5}$$

where m_c is a momentum constant that multiplies the previous weight change, controlling how much momentum is used. The other factor is the product of the learning rate and the gradient with respect to the performance *MSE*.

Finally, the back-propagation algorithm can be implemented in two different ways: *batch mode* and *sequential mode*. One complete presentation of the entire training set during the learning process is called an *epoch*. In batch mode, weight updating is performed after the presentation of all the training examples that constitute an epoch. On the contrary, in sequential mode weight updating is performed after the presentation of each training example. We will use batch training, because it is more efficient to present all the input vectors at the same time than presenting the vectors one at a time.

6.1.4. MATLAB FUNCTIONS

MATLAB provides the necessary tools to design many types of neural networks. Here we will review the function that produces trainable multi-layered back-propagation cascade forward networks, ***newcf***, as well as its parameters, which can be adjusted to suit our needs, and that are summarized as follows:

- *Number of layers, and number of neurons in each layer*:

  Unfortunately, there are no reliable rules for predicting how many hidden layers are required for a given problem or how many hidden nodes are required in each layer, so this must be decided by us during the process (we will describe this in next section).

- *Transfer function for each layer:*

  We can choose between a large range of linear and non-linear transfer functions, such as the saturating linear (*satlin*), log-sigmoid (*logsig*), hyperbolic tangent sigmoid (*tansig*), and purely linear (*purelin*). The tansig function has been used in all the hidden neurons, because it has been commonly utilized in back-propagation networks[1-45], and it is a good trade off for neural nets, where speed is important and the exact shape of the transfer function is not. The purelin function is used in the output neuron, because we want to obtain its input value with a linear transformation.

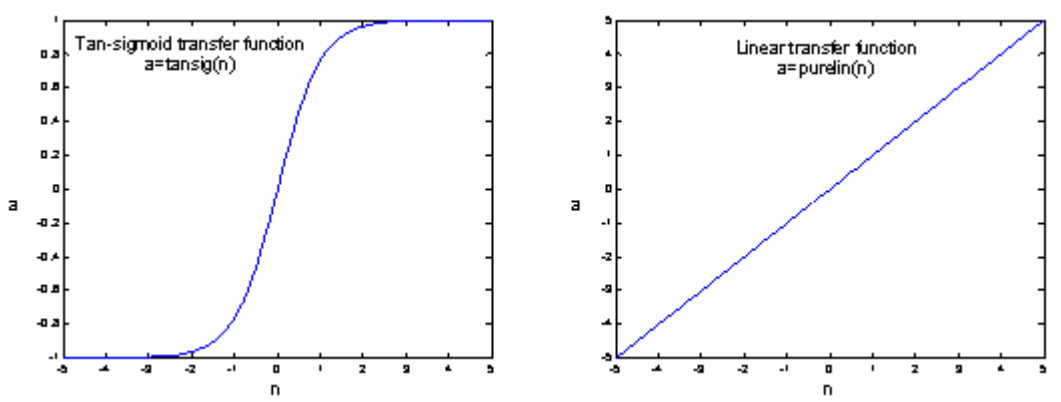

Figure 6.4. Transfer functions

- *Back-propagation network training function:*

  Training functions repeatedly apply a set of input vectors to a network, updating the network each time, until some stopping criteria is met. We will use *trainlm*, which is very fast, and uses the gradient descent back-propagation algorithm described before.

- *Back-propagation weight/bias learning function,:*

  The most specific kind of learning function is a weight and bias learning function. These functions are used to update individual weights and biases during learning. The one we will use is *learngdm*, which uses gradient descent algorithm with momentum, to avoid convergence problems.

- *Performance function:*

  The mean squared error performance function *mse* is utilized. It measures the network's performance according to the mean of squared errors.

Once that the parameters have been explained, the *algorithm* for *newcf* can be described as follows:

Cascade-forward networks consist of *N* layers using the *dotprod* weight function, *netsum* net input function, and the specified transfer functions (*tansig* and *purelin*). The inputs are appropriately normalized and pre-processed so that its mean value is close to zero and the process is accelerated. The first layer has weights coming from the input. Each subsequent layer has weights coming from the input and all previous layers. All layers have biases. The last layer is the network output. Each layer's weights and biases are initialized with *initnw*. Adaption is done with *trains*, which updates weights with the specified learning function (*learngdm*). Training is done with the specified training function (*traingd*). Performance is measured according to the specified performance function (*mse*).

Dotprod applies weights to an input to get weighted input. It returns the dot product of the weight matrix *W* and the input vectors *P*.

Netsum calculates a layer's net input by combining its weighted inputs and biases.

Initnw initializes a layer's weights and biases according to the Nguyen-Widrow algorithm[1-46], which chooses values in order to distribute the active region of each

neuron in the layer approximately evenly across the layer's input space. This presents several advantages over purely random initialization.

6.1.5. MEDICAL APPLICATIONS

Neural networks have been successfully applied to a range of medical problems and entire workshops are devoted to this area. The main applications are:

- Medical imaging, especially ultrasound, magnetic resonance imaging (MRI), Computed Tomography (CT), planar thallium scans and thermal imaging.

- Waveform analysis especially EEG (electro-encephalograms), ECG (electro-cardiograms), etc.

- Detection of cancerous tissue, and of cardiac and pulmonary problems.

We will briefly describe two studies carried out in the latter application mentioned above, to illustrate the techniques:

Ricketts[1-47] applied neural networks to the recognition of single cancer cells. The study used the back-propagation algorithm and the networks was trained on 524 single cells, and tested on 524 previously unseen cells. 80 features were extracted from the cell images. Best performance 96.0% was achieved using a multi-layer network containing 80 inputs, 4 nodes in a hidden layer and a single output node.

In the field of automatic detection of lung diseases, Friman et al.[2-8] made an attempt to detect emphysema in lungs. In the study, 500 CT images were used for training/testing. 14 features were extracted from the images. The results outperformed other statistical methods, the best performance 89.4% was achieved using a multi-layered network containing 8 hidden neurons and one output neuron.

6.2. A SECOND APPROACH ON FEATURES SELECTION

6.2.1. METHODOLOGY[1-48]

In the design of a neural network, there are many free parameters that we have to adjust in order to achieve the best results for our problem, such as the number of inputs, the number of hidden layers, and the number of hidden neurons in each layer.

The number of relevant texture features to obtain an optimum classification is unknown a priori. Many features (25) have been extracted from the images, but some of them do not contribute or even worsen the classification performance, thus they have to be discarded. Many algorithms exist which typically consist of four basic steps[1-49]:

1. A generation procedure to generate the next subset of features X

2. An evaluation criterion J to evaluate the quality of X

3. A stopping criterion for concluding the search. It can either be based on the generation procedure or on the evaluation function

4. A validation procedure for verifying the validity of the selected subset.

The task of feature selection is to reduce the number of extracted features to a set of a few significant features which optimize the classification performance. The best subset

$$X = \{x_i \mid i = 1,..., d; x_i \in Y\} \tag{6.6}$$

is selected from the set

$$Y = \{y_i \mid i = 1,..., D\} \tag{6.7}$$

where D is the number of extracted features, 25 in our case, and $d \leq D$ denotes the size of the feature subset. A feature selection criterion $J(X)$ evaluates a chosen subset X, whereby a higher value of J indicates a better subset. The problem of features selection is to find a subset X from Y such that the number of chosen features is d and J reaches the maximum

$$J(X^{opt}) = \max_{X \subseteq Y, |X|=d} J(X) \tag{6.8}$$

The evaluation criterion J is proposed to be the performance of our neural network. An exhaustive search for feature selection is too time consuming, hence we will use a combination of two suboptimal algorithms, the **Sequential Forward Selection (SFS)** and the **Sequential Backward Selection (SBS)**. In the former, the method starts with an empty set, and one feature among the remaining features is added to the subset with each iteration, so that the subset maximizes the evaluation criterion J. In the later, the algorithm starts with all features selected, and one feature is rejected in each step so that the remaining subset gives the best result. The search concludes when the best *suboptimal* performance is achieved.

To find the best number of hidden layers and neurons, a process of training/testing is used to check the performance of the network when we use different numbers of hidden elements. This has to be applied for each subset of features. To carry out this process, in the back-propagation algorithm we usually have the data used split into a *training set*, which we use to train the network, i.e., find the weights and biases, and a *test set*, which is utilized to evaluate the net performance.

One robust method is the **leave-one-out scheme**, which is a special case of the cross-validation sampling scheme[1-50]. In leave-one-out with n cases, the network is trained n times. Each time the network is trained, $n-1$ cases are used to train the network and it is tested on the one that was left out. This is repeated until each case has been used for testing. In k-fold cross-validation, the data is randomly divided into k subsets. Each subset is used for testing while the remainder are used for training. Training and testing are repeated for each of the k subsets and the testing results are

accumulated and reported. The leave-one-out scheme overcomes the problem of limited data very effectively since it uses almost all available data to develop the network.

Finally, we can see the issue of the **generalisation** ability from two perspectives[1-51]. We can consider that the architecture of the net is fixed, and the issue to be resolved is that of determining the size of the training set needed for a good generalisation, or we can consider that the size of the training set is fixed, and therefore we have to obtain the best architecture for achieving good generalization. We will use the latter approach. For a good generalisation, the ratio of the number of patterns over the number of parameters has been shown to be greater than 3[1-52]. For example, if we had 10 features as inputs, we would need a training set of at least 30 feature vectors, each of them containing its 10 respective features. We will not have problems of poor generalisation in our study, since the number of available patterns exceeds more than three times the number of features utilized in each subset, as we will see in next sections.

6.2.2. TRAINING/TESTING PROCESS FOR EMPHYSEMA

We have based this study on 635 patterns (486 healthy "1", and 149 unhealthy "0"), extracted from 3 different images that present clear emphysematous state. We will use the algorithms previously explained. First, to simplify the process, after studying the histograms for the three images, we have reduced the set of features to 7, because the others present too much overlapping between both classes to expect any results of interest, plus the study would be considerably complicated in computational terms. Therefore, the set of features under study is comprised by: *mean, percentile 0.75, autocorrelation (0,1), Fourier energy, Fourier max. peak, binary area, and binary fractal dimension*. As we can see, the ratio of the number of patterns over the number of inputs is much greater than 3, so we will not have generalisation problems. We have also made a different study for edges and for centres, in order to improve the performance, as we detailed in Chapter 5. Removing the edges, the number of patterns is reduced to 476 (380 healthy and 96 unhealthy), and removing the centres, the number of patterns is reduced to 159 (106 healthy and 53 unhealthy).

The first trials **without edges** using each feature individually yielded very poor results (table 6.1 left). The reason of this is that we have used a much larger number of healthy patterns (380) in comparison to the unhealthy patterns (only 138). Hence, the results are highly biased, and the percentages of correct detection for "0" (unhealthy) are very low. We thought of two possible ways of solving this inconvenient: either we can reduce the number of healthy patterns to 138 or we can increase the number of unhealthy patterns to 499, so that we have the same number of patterns for each class and the bias would disappear. We chose the latter alternative, in order not to lose information, and the results were encouragingly improved (table 6.1 right).

The leave-one-out cross-validation algorithm explained before has been used, and since normally more than 10 runs are needed to obtain accurate results, we have used 15 here.

For individual features, the used network architecture has no hidden layers (since they are not necessary at this stage where there is only one input), and one output with a linear transfer function, where we apply a simple threshold (0.5) to decide if the output belongs to the class healthy "1" (>0.5) or unhealthy "0" (<0.5).

In table 6.1, the results are presented as cross-validation matrices. *T* represents the desired targets or desired response, and *O* the real outputs. For example for the mean grey value, we obtain the following results without bias: 80.95% of true positives (i.e., both the desired and the real response are "0", disease), 79.51% of true negatives (real and desired response equal to "1"), 19.05% false positives (real response disease, desired response healthy), and 20.49% false negatives (real response healthy, desired response disease). We can see the bias in the first table, where the performance 39.13% for true positives is extremely low, while the percentage for true negatives 97.03% is extremely high.

In figure 6.5, it is illustrated that the larger the number of epochs of training, the better the performance. The number of epoch has been fixed to 50, where the performance does not improve any more.

Mean

O\T	0	1
0	39.13	60.87
1	2.97	97.03

O\T	0	1
0	80.95	19.05
1	20.49	79.51

Percentile

O\T	0	1
0	35.51	64.49
1	0.66	99.44

O\T	0	1
0	73.94	26.06
1	20.77	79.23

Autocorrelation

O\T	0	1
0	26.81	73.19
1	4.45	95.55

O\T	0	1
0	68.13	31.87
1	15.04	84.96

F. Energy

O\T	0	1
0	31.88	68.12
1	3.15	96.85

O\T	0	1
0	60.36	39.64
1	15.00	85.00

F. Max. Peak

O\T	0	1
0	32.15	67.85
1	2.78	97.22

O\T	0	1
0	72.79	27.21
1	25.81	74.19

Binary area

O\T	0	1
0	60.87	39.13
1	4.64	95.36

O\T	0	1
0	89.13	10.87
1	11.99	88.01

Binary fractal dimension

O\T	0	1
0	27.54	72.46
1	4.27	95.73

O\T	0	1
0	89.52	10.48
1	12.77	87.23

Table 6.1. Features' performances. Left: Biased results. Right: Unbiased results

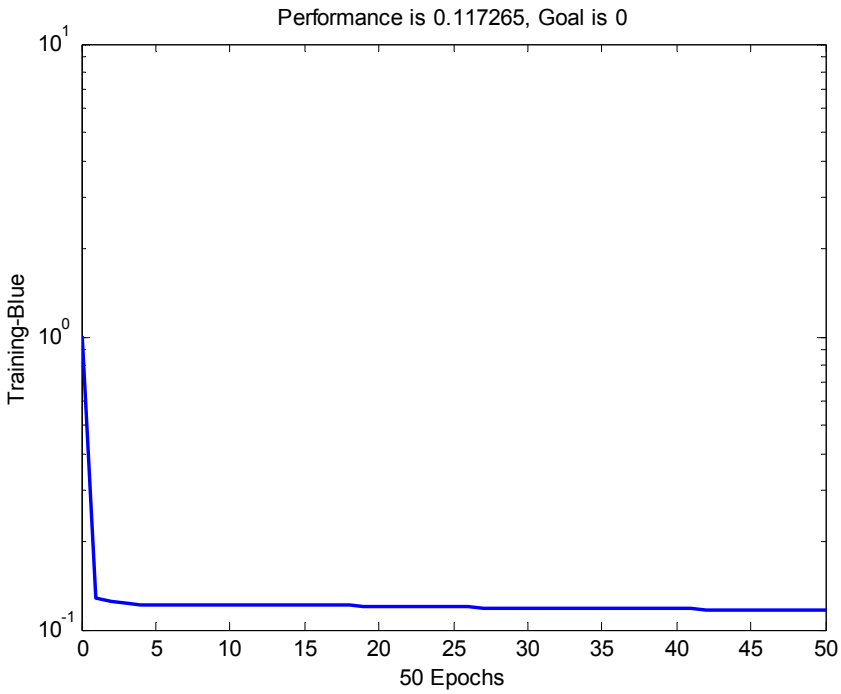

Figure 6.5. Training with 50 epoch at each step of the leave-one-out algorithm

The features can be sorted out by their respective global performances, i.e., the average of the true negatives and true positives, as follows: 1. Binary area (88.57%), 2. Binary fractal dimension (88.37%), 3. Mean (80.26%), 4. Percentile (76.59%), 5. Autocorrelation (76.58%), 6. Max. peak (73.48%), and 7. Energy (72.71%).

Both the Sequential Forward and the Sequential Backward Selection algorithms (SFS and SBS) for discovering the best suboptimal subset of features are utilized. First, we try with the SFS, starting with the best feature, area, and adding features one by one in order of performance, until the results do not improve any more:

- *Area + fractal dimension:*

O\T	**0**	**1**
0	81.00	19.00
1	28.04	71.96

Table 6.2. Binary area and fractal dimension without hidden layer.

The global performance is 76.37%, which is worse than the 88.57% corresponding to the area. This means that we have to add hidden elements. With 1 hidden layer comprising 2 nodes the global performance is 90.31%:

O\T	**0**	**1**
0	91.46	8.54
1	10.83	89.17

Table 6.3. Binary area and fractal dimension with hidden layer of 2 nodes

The results do not improve adding more hidden elements. Throughout all this process, the best results have been proved to appear using only one hidden layer with two nodes. Adding more layers or more nodes was not effective due to the size of the subsets of features, and the complexity of our problem. Hence, in order to abbreviate the procedure, *from now on, at each step one hidden layer with 2 neurons is used*, although many other trials changing the number of layers and nodes have been carried out to back up the results.

- *Area + fractal dimension + mean:*

O\T	**0**	**1**
0	91.84	8.16
1	9.67	90.33

Table 6.4. Binary area, fractal dimension and mean.

The global performance is 91.09%, which outperforms the previous subset.

- *Area + fractal dimension + mean + percentile:*

O\T	**0**	**1**
0	89.97	10.03
1	10.86	89.14

Table 6.5. Binary area, fractal dimension, mean and percentile

The performance is 89.56%, which is worse than the previous one. The algorithm SFS stops here. However, we made other attempts with other subsets of four and five features which turned out to be worse as well.

Now the SBS algorithm will be used. All the features are selected at first, and we will start removing features one by one in order of worse performance, until the performance of the new subset is worse than that of the previous one.

- *Area + fractal dimension + mean + energy + max. peak + percentile + autocorrelation:*

O\T	**0**	**1**
0	91.45	8.55
1	7.73	92.27

Table 6.6. All features, with hidden layer of 3 nodes.

The global performance is 91.87%, but this time, as an exception, it was achieved with 3 hidden nodes instead of 2. The worst feature of this subset proved to be the energy, hence the new subset is:

- *Area + fractal dimension + mean + max peak + percentile + autocorrelation:*

O\T	**0**	**1**
0	92.88	7.12
1	7.35	92.65

Table 6.7. All features except energy.

The global performance, again with 2 hidden nodes, is 92.77%. Removing more features or adding a new hidden layer has worsened this result. Hence, this is our optimal subset of features for centres in emphysema.

For the study **only with** *edges*, the idea was to find a subset of features from the best group of parameters obtained for the centres. As explained in previous sections, those features that require a complete structure in a region will not produce useful

results for edges. Thus in the set of features composed by area, mean, percentile, autocorrelation, fractal dimension, and max. peak, the three latter will not yield good performances, as proved in the table below.

Mean

O\T	0	1
0	76.36	23.64
1	18.15	81.85

Percentile

O\T	0	1
0	65.62	34.38
1	28.03	71.97

Autocorrelation

O\T	0	1
0	51.90	48.10
1	24.82	75.18

F. Max. Peak

O\T	0	1
0	52.54	47.46
1	34.26	65.74

Binary area

O\T	0	1
0	89.86	10.14
1	17.50	82.50

Binary fractal dimension

O\T	0	1
0	75.48	24.52
1	62.41	37.59

Table 6.8. Features' performances in edges.

The best features in order of performance are: 1. Area (86.18%), 2. Mean (79.14%), 3. Percentile (68.82%), 4. Autocorrelation (63.58%), 5. Max. peak (59.18%), and 6. Fractal dimension (56.45%).

Applying the SFS algorithm:

- *Area + mean:*

O\T	**0**	**1**
0	88.88	11.12
1	16.15	83.85

Table 6.9. Binary area and mean for edges with a hidden layer of 2 nodes.

The global performance is 86.36%, which improves the results above, using 1 hidden layer with 2 neurons.

- *Area + mean + percentile:*

O\T	**0**	**1**
0	86.45	13.55
1	13.52	86.48

Table 6.10. Binary area, mean and percentile for edges with a hidden layer of 2 nodes.

The performance continues to improve: 86.46%. Here the SFS algorithm stops, because adding more features worsen the results.

Applying the SBS algorithm:

- *Area + mean + percentile + autocorrelation + max.peak + fractal dimension:*

O\T	**0**	**1**
0	81.41	18.59
1	19.01	80.99

Table 6.11. All features for edges with a hidden layer of 2 nodes.

The performance is 81.03%. Let us remove now the worst feature of the set, which is the fractal dimension.

- *Area + mean + percentile + autocorrelation + max. peak::*

O\T	0	1
0	82.18	17.82
1	17.33	82.67

Table 6.12. All features except fractal dimension for edges with hidden layer of 2 nodes.

The performance improves: 82.42%

- *Area + mean +percentile + autocorrelation:*

O\T	0	1
0	87.01	12.99
1	16.38	83.62

Table 6.13. All features except fractal dimension and max. peak with hidden layer of 2 nodes.

The performance improves again: 85.31%. If we remove the next feature, autocorrelation, we obtain the same result as in the SFS algorithm. Therefore, the best subset of features for edges comprises: binary area, mean and percentile, with a hidden layer of 2 neurons.

After all these procedures, we have come up with two neural networks, one for detecting emphysema in the inner parts of the lungs, and the other in the edges. For the first one, an optimal subset of 6 features extracted from the lung images, composed by binary area, binary fractal dimension, mean, percentile, autocorrelation, and Fourier max. peak, is presented to the net, which has a hidden layer with two neurons and one output neuron, achieving a performance of 92.77%. For the second one, the subset of features comprises mean, percentile and binary area, and the neural net has also one hidden layer with 2 neurons and one output. The performance in this case is lower, 86.46%.

In the figure 6.6, we show the images from the emphysematous lungs that have been used in this section. We present with green X the diagnosis made a priori by the human observer, and with red squares the diagnosis made automatically by the neural

nets. We can visualize clearly that the nets are trying to adjust their diagnosis to that of the observer. Finally, in the last image (bottom left) we can see how the net works with another image of a lung with emphysema that it has not seen before, yielding a good result.

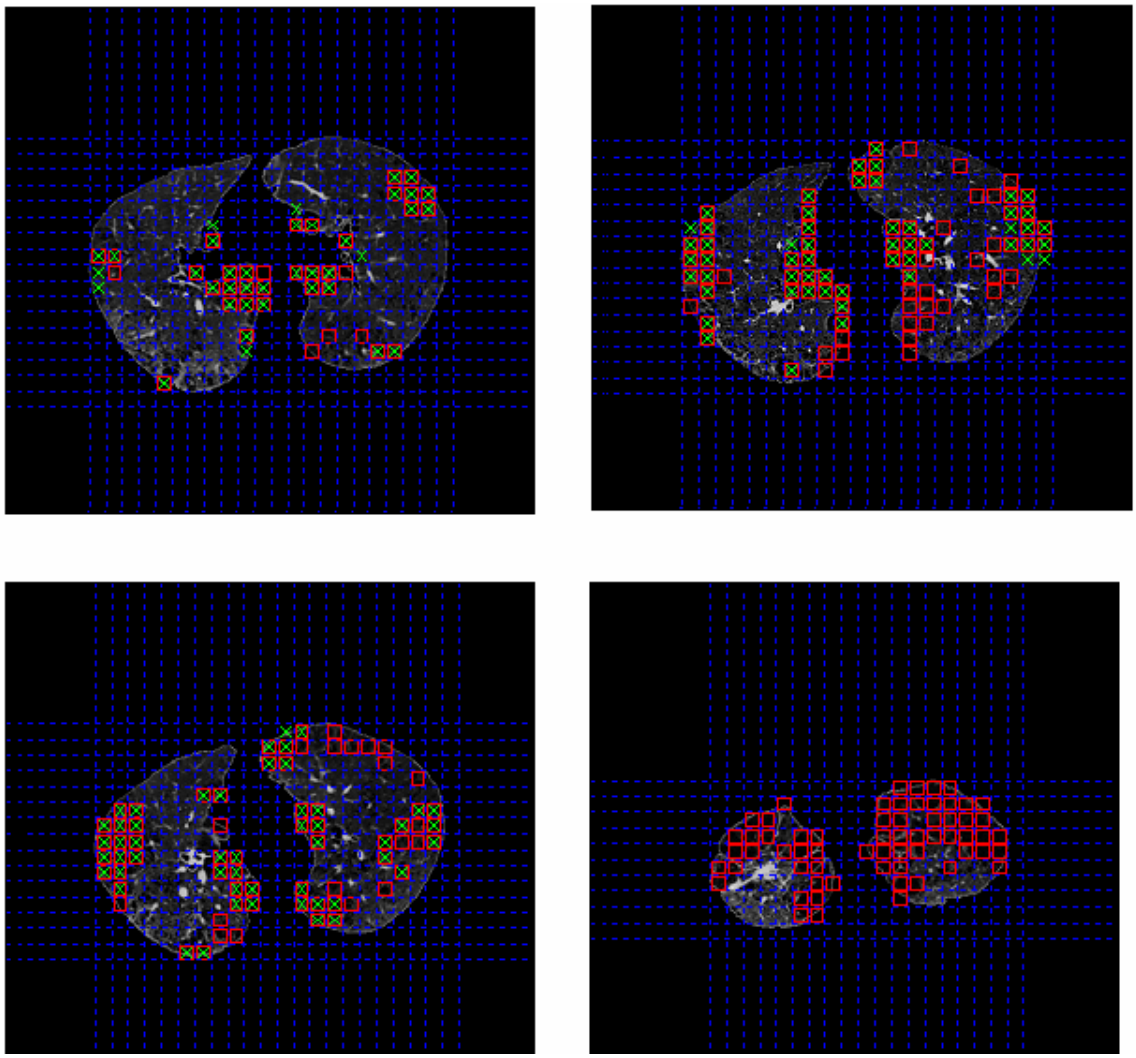

Figure 6.6. Visual results of the process. Diagnosis of emphysema.

6.2.3. TRAINING/TESTING PROCESS FOR FIBROSIS

This procedure is similar to that of the previous section. The number of patterns is 409 (296 healthy "1", and 113 unhealthy "0"), extracted from 2 different images of lungs affected with fibrosis. The initial set of features has been reduced to the same 7 parameters used for emphysema, because they appear to be the best ones for fibrosis as well, although for different ranges of values (for example a very high mean grey level will correspond to a fibrous area, and a very low value to a emphysematous area, so we can see that the same feature can have a good discriminatory power for both diseases).

Again, the ratio of the number of patterns over the number of inputs is much greater than 3, so the generalisation ability is expected to be more than enough. Two separate studies are developed: for centres the number of patterns is reduced to 314 (229 healthy and 85 unhealthy), and for edges this number is reduced to 95 (67 healthy and 28 unhealthy). In order to avoid the bias, the number of unhealthy patterns has to be made equal to the number of healthy patterns.

Table 6.14 summarizes the performances **without edges** obtained when we use each feature individually to train a neural network with only one input, and without hidden layers. The features can be sorted out by their respective global performances as follows: 1. Percentile (94.76%), 2. Binary fractal dimension (93.45%), 3. Mean (93.01%), 4. Max. peak (92.58%), 5. Binary area (92.36%), 6. Autocorrelation (70.31%), and 7. Energy (70.00%). At first sight, we realize that these values are much better than those corresponding to emphysema. The explanation of this has a relation with the study of the histograms for normal areas, and for areas with disease (emphysema and fibrosis) that we developed in Chapter 5. In those graphics it was shown that the overlapping between normal and abnormal regions was more relevant for emphysema than for fibrosis. Hence, the parameters for emphysema cannot discriminate both populations as well as the parameters for fibrosis.

Since all the features appear to have similar performances in a range above 90%, instead of using the SFS and SBS algorithms, we will develop a more exhaustive and

accurate method to find the subset. This consists of studying all the possibilities and performances of groups of 2, 3, 4, 5, 6 and 7 features.

Mean

O\T	0	1
0	96.51	3.49
1	10.48	89.52

Percentile

O\T	0	1
0	98.25	1.75
1	8.73	91.27

Autocorrelation

O\T	0	1
0	73.80	26.20
1	33.19	66.81

F. Max. Peak

O\T	0	1
0	93.01	6.99
1	7.86	92.14

F. Energy

O\T	0	1
0	73.68	26.32
1	33.68	66.32

Binary area

O\T	0	1
0	89.52	10.48
1	4.80	95.20

Binary fractal dimension

O\T	0	1
0	93.89	6.11
1	6.99	93.01

Table 6.14. Features' performances for fibrosis.

Let us apply the selection algorithm:

- *Pairs:*

  Most of the pairs have achieved performances above 93% and 94%. The best result 94.76% has been obtained combining percentile and fractal dimension, and exactly the same value was achieved with the mean and fractal dimension. Here we illustrate it with the cross-matrix correspondent to the percentile and fractal dimension pair.

O\T	**0**	**1**
0	97.38	2.62
1	7.86	92.14

Table 6.15. Percentile and fractal dimension with hidden layer of 2 nodes, for fibrosis.

  The neural network has a hidden layer of 2 neurons and this architecture will be the best for the rest of subsets as well, as discussed in previous section.

- *Groups of 3 features:*

  The best performance, again 94.76% has been achieved by the group formed by percentile, max. peak, and area:

O\T	**0**	**1**
0	96.51	3.49
1	6.99	93.01

Table 6.16. Percentile, max. peak, and area with hidden layer of 2 nodes, for fibrosis.

- *Groups of 4,5,6, and 7:*

  Unexpectedly, the best performance for each one of these groups is fixed to 94.76%. It turns out that adding more features does not improve the results, but they do not worsen either. So at this point we have to choose how many features we want in our subset. Normally, it is better to have the minimum

number of features possible, in order to fasten the process. However, in this case, we will choose the same group of 6 features that we had in emphysema. The reasons are two: first, having exactly the same subset of features for both diseases will reduce the execution time and computational cost of our program; second, when new images are presented to the net, having more features can compensate the result for the performance in case that one feature fail to capture the information of the new image as well as expected. The cross-validation matrix for the subset composed by fractal dimension, area, mean, percentile, max. peak, and autocorrelation is shown in the table below. Tables 6.16 and 6.17 are exactly the same, although they have been calculated in independent processes that could have yielded different results.

O\T	**0**	**1**
0	97.38	2.62
1	7.86	92.14

Table 6.17. Subset of 6 features with hidden layer of 2 nodes, for fibrosis.

Finally, the process has to be repeated for **edges** in the same way. Therefore, we will outline here that the best features are percentile (92.36%), mean (91.3%), and area (89.13%). The best pair is composed by mean and percentile (91.3%), and the optimal subset is formed by these three features, mean, percentile, and binary area (93.48%), as illustrated in the table 6.18.

O\T	**0**	**1**
0	95.65	4.35
1	8.70	91.30

Table 6.18. Subset of 6 features with hidden layer of 2 nodes, for fibrosis.

Therefore, we have designed two neural networks to detect the presence of fibrosis in the centres and edges of the lung, using the same subsets of features that carry out that task for emphysema. Both nets have one single hidden layer, with 2 nodes, and their performances are 94.76% for centres and 93.48% for edges.

The utilized images from the fibrous lungs are displayed below, with both the automatic diagnosis (red squares) and the human observer's diagnosis (green X). The last two images have been used as examples of the behaviour of the nets when new images are presented to them.

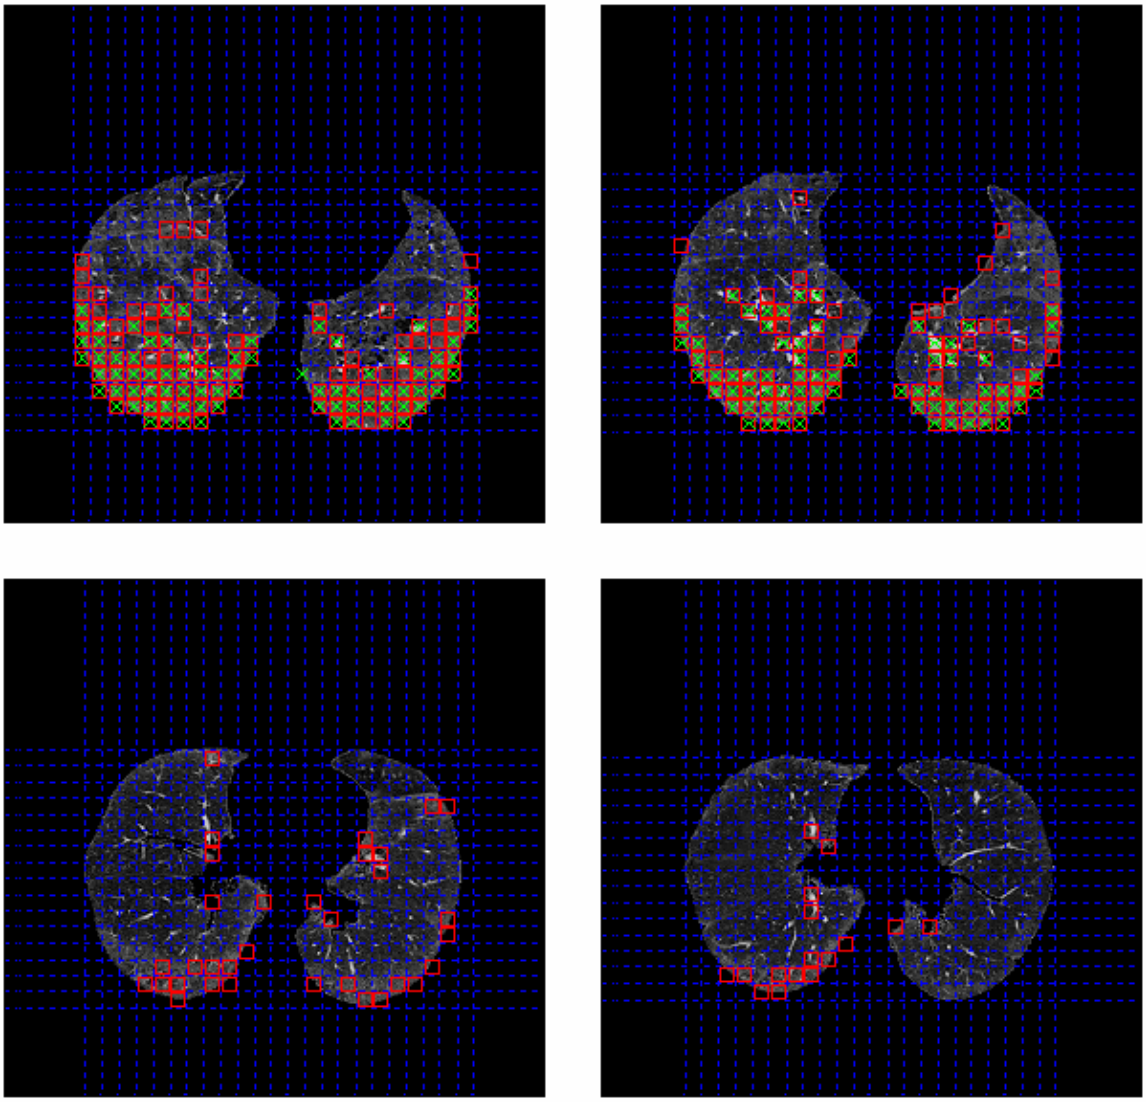

Figure 6.7. Visualization of the process. Diagnosis of fibrosis.

6.2.4. OTHER POINTS OF INTEREST

An interesting consideration is to check the behaviour of the program when an image containing both emphysema and fibrosis is presented to it, which is a common thing because both diseases are often together. The program first extracts from the image the features that will feed the neural networks trained to detect emphysema and

fibrosis. These features are the same for both diseases, so the extraction process will not require a large amount of time. Then, the neural networks will carry out a classification of each region of the lung, marking them as unhealthy if they are detected as abnormal by the neural network trained to predict emphysema, or by the neural network trained to predict fibrosis. Hence, it is a simple **OR** comparison to decide the diagnosis. In figure 6.8, we present again one of the images of figure 6.6, which presented emphysema, but also presents some areas with abnormally high density levels. We can see that the final diagnosis actually covers the detection of both diseases.

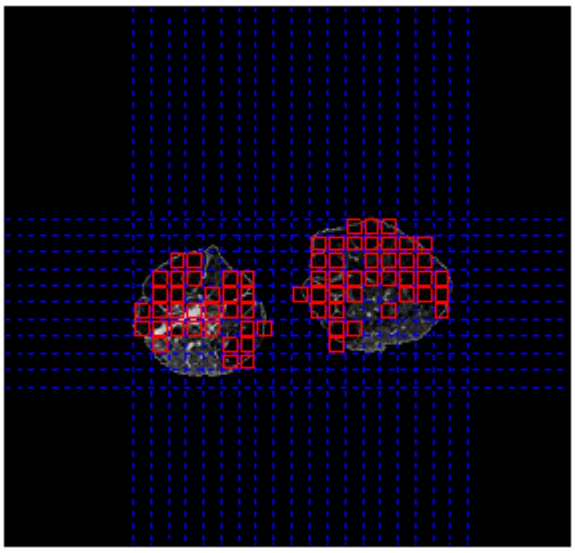

6.8. The final diagnosis for both diseases.

We have also made an attempt to utilize only one net to detect both diseases. After repeating the training/testing process, we came up with the conclusion that the results are worse than in our approach of separating both problems. In figure 6.9, we can check visually that the number of false positives increases considerably, if we compare this image with the diagnosis resultant in figure 6.6, where we have studied the same lung.

Another point of interest is related to the algorithms that have been used to select the subset of features. As we discussed before, there is no reliable method to find the optimal subset, and these alternative methods can only find suboptimal solutions. We present in figure 6.10 the curves *performance/size of the subset* that we have found applying the algorithms, which have different shape and behaviour for each single

case, making it difficult to find the real optimal subset, especially if the initial set of features is very large.

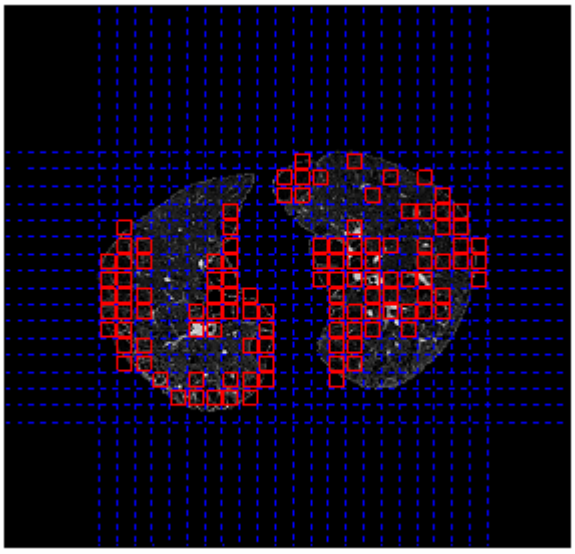

Figure 6.9. One net for both diseases.

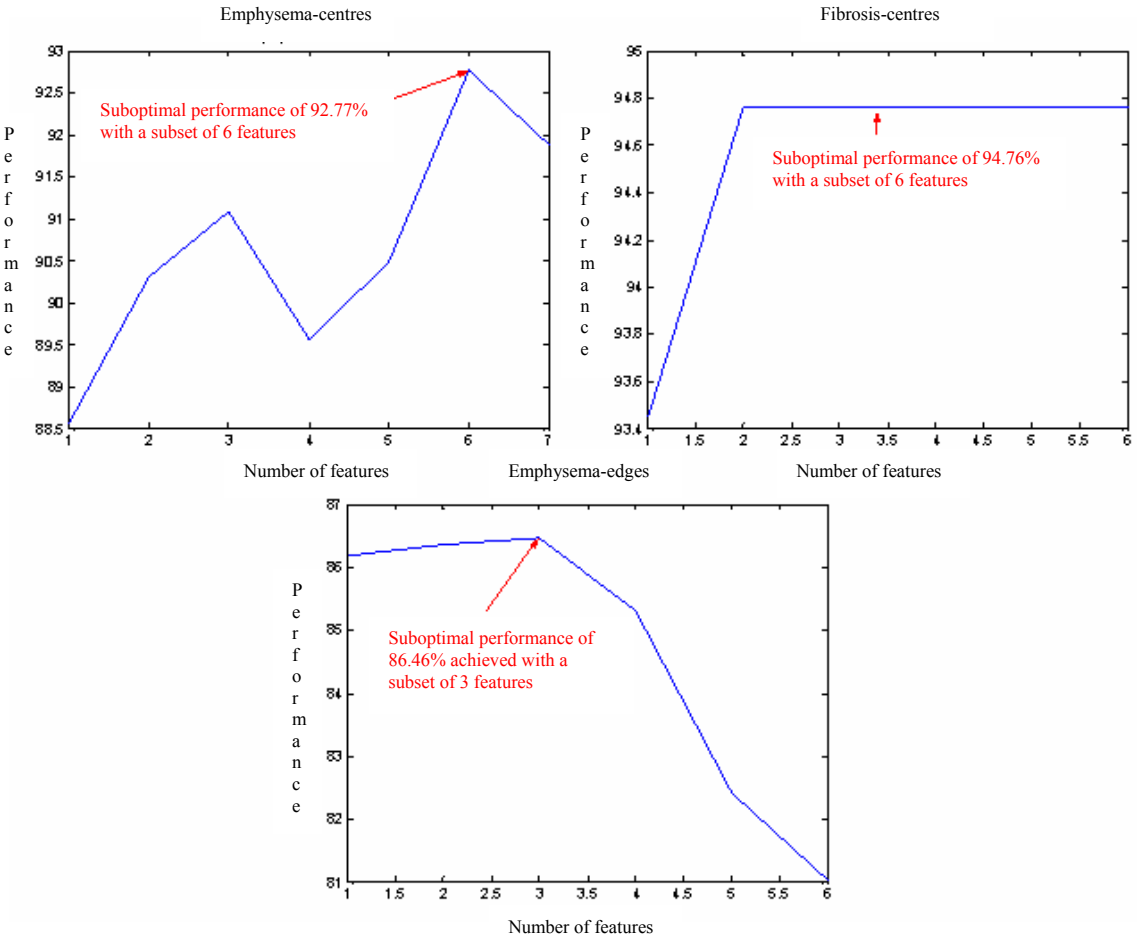

Figure 6.10. Performance/size of the subset curves.

6.3. SUMMARY AND CONCLUSIONS

In this Chapter we have made a review of neural networks, the AI techniques used in this project to accomplish an automatic diagnosis of the state of the lungs. We have started describing the main elements that form the different architectures, paying special attention to the structures utilized, the cascade-forward networks, as well as the learning algorithms, and the tools provided in MATLAB to develop the process.

4 different neural networks have been trained to detect the presence of abnormalities in the lung parenchyma, and stored in the file "neuronal_networks.mat". Two of these nets have been trained to detect the presence of emphysema (one of them in the centres, and the other one in the edges), and the other two to detect the presence of fibrosis (centres and edges again). Applying a series of suboptimal algorithms, we have come to the conclusion that the best subset of features for detecting emphysema and fibrosis in the inner parts are comprised by six features: **mean, percentile, autocorrelation, Fourier energy, binary area, and binary fractal dimension**, yielding performances of **92.77%** and **94.76%** respectively. For edges, the subset includes **mean, percentile, and binary area**, and produces performances of **86.46%** and **93.48%**.

With this study, the part of the project corresponding to the 2D diagnosis of CT lung scans has finished. In next Chapter, a 3D approach is introduced and compared to the results obtain in 2D.

7. A 3D APPROACH

7.1. MOTIVATIONS

Until now, our work has been based only on the information contained in a single slice extracted from CT lung scans. However, one single slice does not represent the real state of the lung. As explained in Chapter 2, the CT scanner used in the B.R.I. can produce data-sets with as many as 80 image slices, which appear normally in pairs, one “hard” or sharp focus, and one “soft” focus for each slice. Hence, when we work with only one image we are using 1/40 of the whole information available.

A 3D investigation will give us a more general idea of the actual extension of the disease. If we take images of several slices and sort them out, we could expect to obtain the state of the lung with more accuracy and reliability than with just one slice, since one slice gives incomplete information, and sometimes this information can be erroneous. For instance, if we consider a single slice, we could want to search for the upper one and the lower one in order to compare them, so that they all contribute to the final diagnosis, i.e., if the central slice has an unhealthy area, and the other two are healthy, then maybe that region is actually healthy, so we have to look for a balance between the three of them to make a final decision. The study will be based on 2D features extracted separately from each slice, although another approach could consist of using volume features of all the 3 images together.

We will discuss the changes made to the 2D user graphical interface, and a comparison between the advantages and disadvantages of a 2D and 3D diagnosis will be carried out.

7.2. 3D INTERFACE

Some changes in the 2D user interface that we described in Chapter 4 are necessary to make a convenient process and treatment of several images. To run the program, it is only necessary to type "interface3d" in the prompt of MATLAB and the figure 7.1 will appear.

It comprises some new basic tools to deal with a group of images, such as *insert/remove slice*, and also the necessary elements to obtain a 3D diagnosis using in our case 3 slices (the current one, the previous one and the next one).

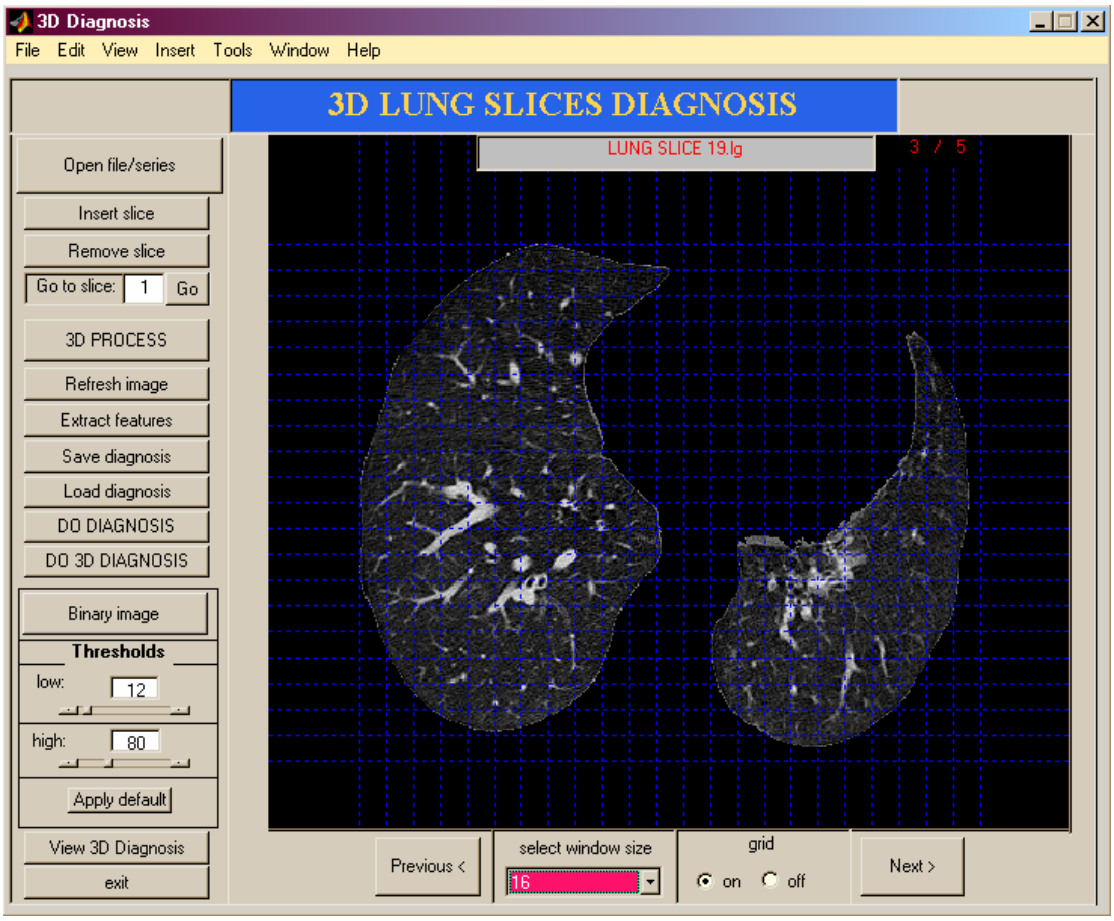

Figure 7.1. 3D Interface

We will review the changes and the new tools:

- **Open file/series:**

  It opens one single image or a part of a whole dataset. The images have extension "*.lg". When one image is selected, the dialog box from figure 7.2 appears. If we are interested in loading more than one image, we have to make sure that they all belong to the same dataset and are in the same directory, plus their names must be consecutive. For example, in figure 7.1, 5 images have been loaded, the first of them was 17.lg, and the rest were 18.lg, 19.lg, 20.lg, and 21.lg. If one of them did not exist in the current directory, the loading process would stop at that point.

  Due to the memory limitations of MATLAB, the process is considerably slow when the number of load slices is above 40. For illustrating this procedure, a large number of slices is not needed, so it will not be a problem.

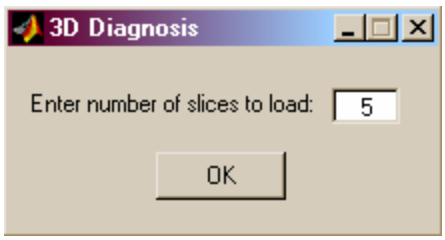

Figure 7.2. Number of slices to load

- **Insert slice/Remove slice/Go to slice:**

  Insert slice adds a new image after the current one, and *remove slice* removes the current image. *Go to slice* displays the selected image. These three functions also refresh the number *current_ image/number_ of total images*, that we can see as 3/5 in figure 7.1.

- **Next/Previous buttons:**

They will allow us to move from one image to the previous or the next one. If we are in the first image and we press the *previous button*, an error message will be displayed. The same happens if we are in the last one and press the *next button* (figure 7.3)

Figure 7.3. Error messages when we try to access to images out of range

- **3D PROCESS:**

This option processes pairs of images, and is basically used to demonstrate the differences between the "sharp" and the "soft" version of the same slice. Figure 7.4 displays the options available for the comparison, which are image subtraction, and the comparison between the diagnoses corresponding to the different images.

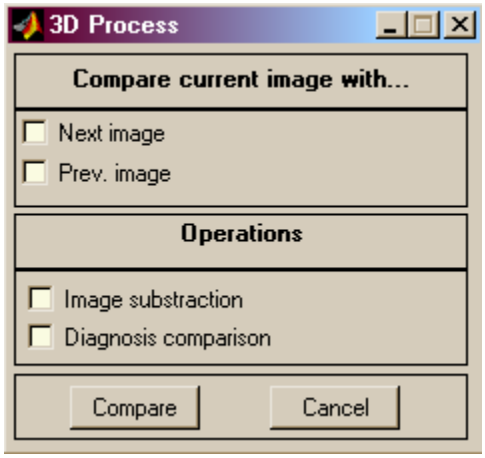

Figure 7.4. 3D Process of pairs of images

An example of this simple process is shown in figure below. The image produced by the subtraction of both images shows a small difference between them, especially relevant in the contour. This is due to the blur in the edges produced by the smoothing process to obtain the soft image.

The diagnosis comparison shows that in the sharper version of the images more abnormalities are detected, because during the smooth filtering some noise is removed but there is also a loss of information and structure. Therefore, from now on, we will only work with the "sharp" version of the images, because they keep the integrity of the information.

As we can see, this process of comparison only has sense when the two images belong to the same pair correspondent to a single slice, because images originated in different positions of the lung will have different shape, size and hence the comparison would yield useless results.

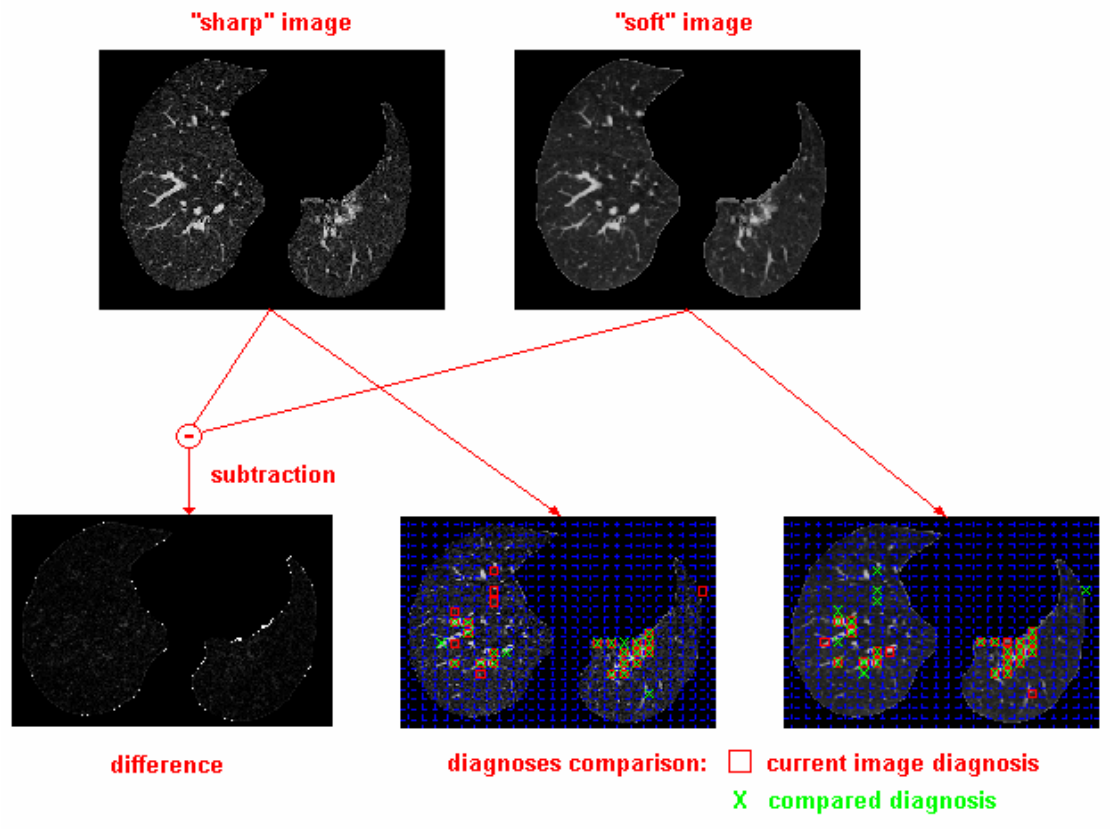

Figure 7.5. Comparison of a pair of images (sharp/soft) corresponding to the same slice.

- **DO 3D DIAGNOSIS:**

  This process will be further reviewed in next section. The base is the same as the diagnosis of one single slice, i.e., it comprises two modules: *features extraction*, and *classification* of these features using some previously trained neural networks.

  The difference is that 3 slices are used in the process, in each of them there has to be an extraction of features, then the classification, and finally a relation between the 3 output values (from the 3 slices) obtained for each region has to be found in order to achieve the best performance in the automatic detection of diseases.

- **View 3D Diagnosis:**

  This displays a visual representation of how the disease extends throughout the whole lung. In order to see the inner parts of the lungs, in the representation, we have only plotted the edges for each slice, and then we have loaded the diagnosis for each slice, and mark the unhealthy areas with a red X.

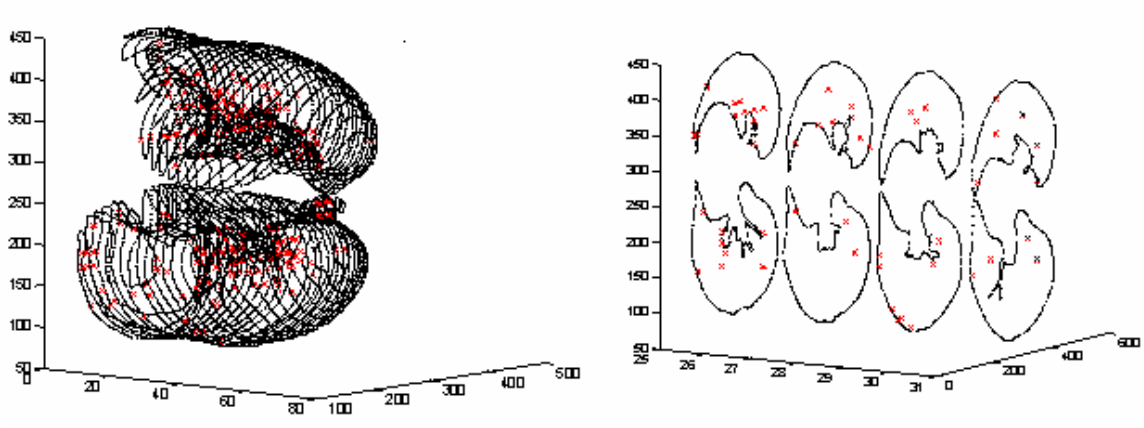

Figure 7.6. 3D representation of the state of the lung. Left: the whole dataset 1. Right: Zoom in on 4 slices.

To make the representation, we have to select the first file with information about the edges, which has the format "*alledges.lg" (for example 1_alledges.lg). After selecting the first edges file, we have to select the first diagnosis file that fit it (1.dgn). If the rest of the files in the dataset are in the same directory and their names are in consecutive order (2_alledges.lg, 3_alledges.lg… 2.dgn, 3.dgn …), then they are all loaded and representations like that of figure 7.6. are plotted.

For this figure, we have used the dataset number 1, that we will discuss later in next section. In fact, as we explained before, only the files correspondent to the "sharp" version of each slice are used, therefore instead of loading all the files consecutively (1_alledges.lg, 2_alledges.lg, 3_alledges.lg…), they are loaded by twos (1_alledges.lg, 3_alledges.lg, 5_alledges.lg…).

7.3. COMPARISON BETWEEN 2D AND 3D DIAGNOSIS. ADVANTAGES AND DISADVANTAGES.

The tools that we will use to develop a 3D Diagnosis are the same than for 2D, only that now a combination of several slices is needed. The same textural features that comprise information about singles slices are utilized, instead of some volumetric features such as the mean grey value applied to a cubic region containing 3 slices.

Hence, we will make an attempt to combine these planar features that belong to different slices in order to achieve an improvement in the final performance of the diagnosis.

The same four neural networks defined in Chapter 6 will receive the features and will classify them to come up with the diagnosis. These neural networks are composed by one hidden layer of two nodes, and one single output. Each net receive the following inputs:

- **Neural network for emphysema in centres:** mean grey value, percentile 0.75, autocorrelation (0,1), Fourier max. peak, binary area, and binary fractal dimension.

- **Neural network for emphysema in edges:** mean grey value, percentile 0.75, and binary area.

- **Neural network for fibrosis in centres:** mean grey value, percentile 0.75, autocorrelation (0,1), Fourier max. peak, binary area, and binary fractal dimension.

- **Neural network for fibrosis in edges:** mean grey value, percentile 0.75, and binary area.

To accomplish the diagnosis of the current lung image, first it is divided into 16x16 pixels regions, and the upper and lower images are also divided in the same way. Narrowing the focus to detect the abnormalities in only one single 16x16 pixels region, the process consists of extracting the feature vector corresponding to that region, and the feature vectors corresponding to the same region in the previous and next images. Then, these three vectors are classified independently using the neural networks, and we come up with 3 different diagnoses. The problem that we attempt to solve is to decide the best way of combining these results to obtain the final diagnosis of the 16x16 region under study.

In figure 7.7 this procedure is illustrated. The three classifications for the selected area and the three slices are named *S1, S2,* and *S3*. The final decision *D* about the state of the region has to be a function *f* of the classification results. Initially, before finding this function, we can expect that the most relevant parameter has to be *S1*, although including *S2* and *S3* in the diagnosis can help to solve certain situations where the normality or abnormality is not clear enough.

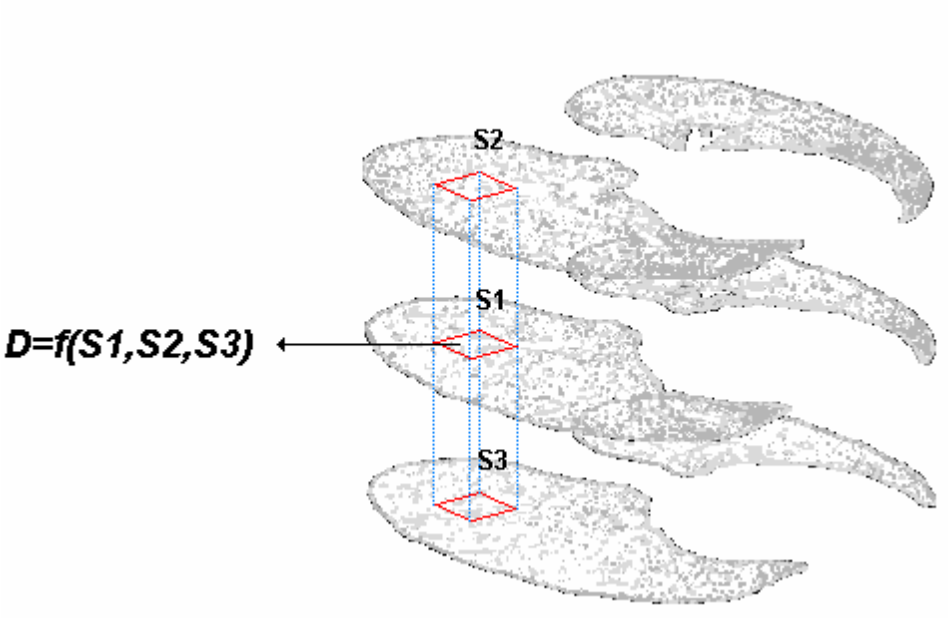

Figure 7.7. Illustration of the 3D Diagnosis process.

There are some previous points that we have to consider. A special case can appear when we study a region whose upper or lower homologous regions do not exist, because they are out of the boundaries of the lung. In this case, they cannot contribute to the final diagnosis, and if neither the upper nor the lower region exists, then the diagnosis will be equivalent to a simple 2D diagnosis of the area of interest.

Other important point is that we have to be careful when we split the images. This is carried out at the first stage when we load some images, and it also depends on the number of images. The splitting is done in a way that a 3D representation of the lung with the loaded images would fit in a cubic volume, as we can see in figure 7.8. This has to be done to work with planar regions that correspond to the same volumetric 16x16x3 cube (16x16 pixels planar area x3 slices). The difference with the 2D interface is that there we squash the single planar image so that it would fit in the smallest possible rectangular region, to have the lower number of 16x16 regions, and hence to obtain the diagnosis in less time.

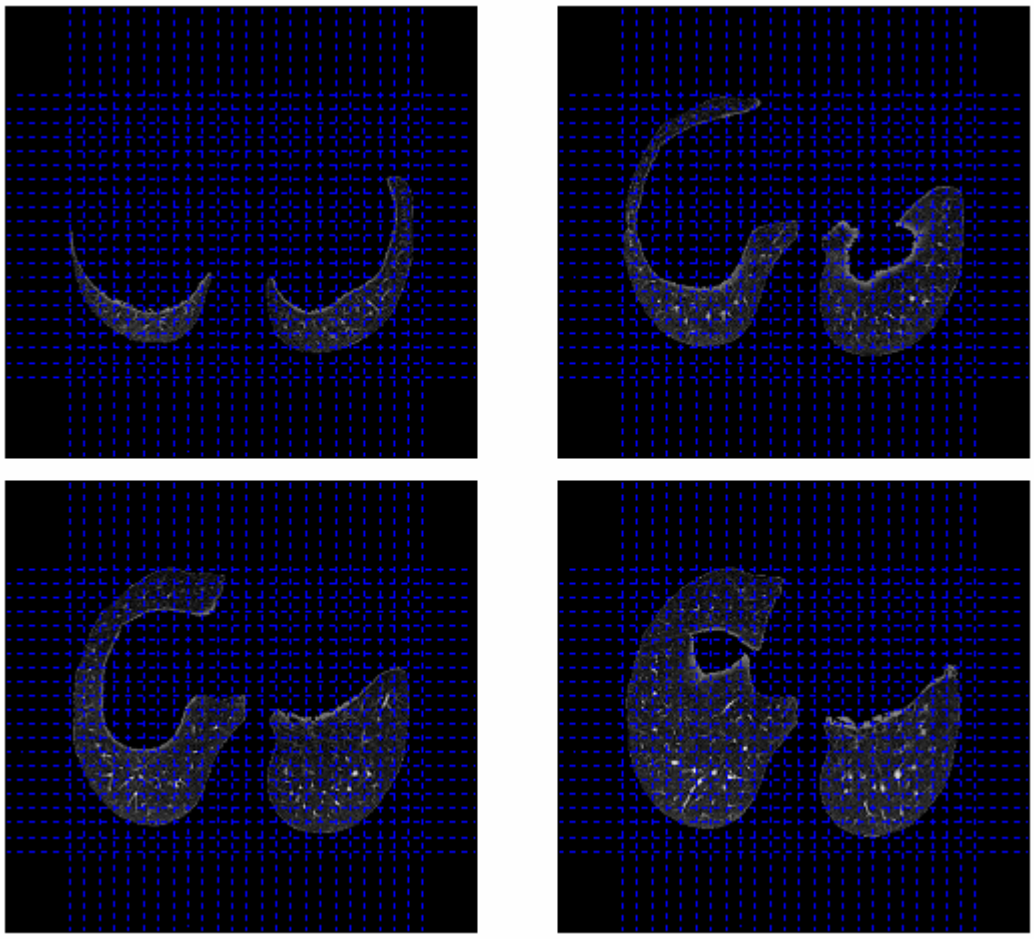

Figure 7.8. Splitting process.

Finally, the equation proposed to obtain the diagnosis combining *S1, S2, S3* is:

$$D = (1 - 2a)S1 + a(S2 + S3) \quad (7.1)$$

where α is the parameter that we have to determine. If for example α=0.33, the three parameters would contribute exactly the same weight to the final value. On the other hand, if α=0, then $D=S1$, which correspond to the 2D Diagnosis.

We can expect that lower values of α give better results, because from one image to the next or previous one, some areas change considerably (figure 7.9), and the results would not be reliable if S2 and S3 had a bigger contribution.

Figure 7.9. Changes in the regions from one image to the upper and lower ones.

In the figure above it is shown that in the current image there is a clear area with abnormally high greyscale levels, while in the lower one it completely disappear and seems to be healthy. In the upper one, however, there is still some disease. Therefore, the selection of α is a difficult and delicate problem.

We have carried out a study on 4 different datasets of lung images to determine the best value of α. We have to say, though, that the results using the slices belonging to the lower and upper part of the lungs have not been considered, because it is not possible to get a certain degree of good performance using a 3D diagnosis, due to the big changes in size and shape from one slice to the next one (we can see this in figure 7.8). However, for the central slices, which are similar in shape and size, we obtain some results of interest.

Next figures show curves that illustrate the percentage of error when we change the parameter α. They correspond to the 4 datasets previously mentioned. 5 central images of each dataset have been utilized to obtain the curves.

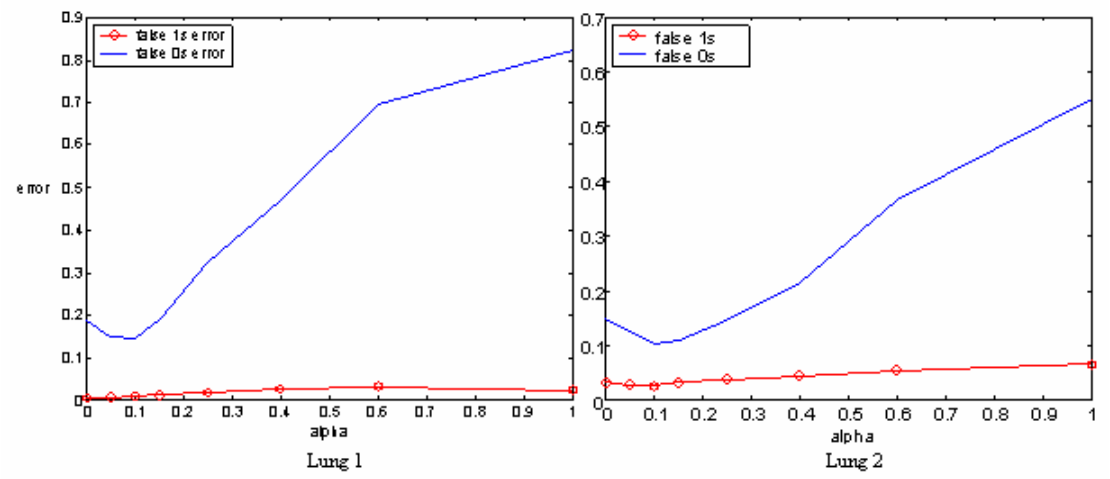

Figure 7.9. Curves of *error-α* for datasets 1 and 2.

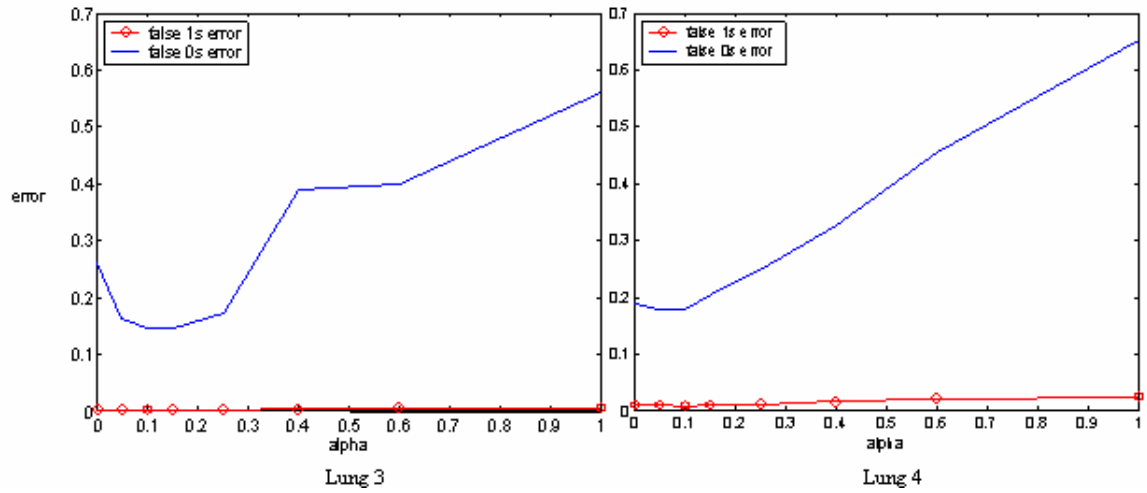

Figure 7.10. Curves of *error-α* for datasets 3 and 4.

For the trials we have used the following values of α: 0, 0.05, 0.1, 0.15, 0.25, 0.4, 0.6, and 1.The graphics present the best performance for values of α in very narrow range around 0.1. Hence, we take α=0.1 as our chosen value. The equation 7.1 finally turns into:

$$D = 0.8S1 + 0.1(S2 + S3) \tag{7.2}$$

The error percentages and the improvements with respect to the 2D diagnosis are presented in the table 7.1

Dataset	2D performance error%	3D performance error%	Improvement %
1	18.45	14.84	3.61
2	14.92	12.68	2.24
3	26	16	10
4	18.84	17.68	1.16

Table 7.1. Comparison 2D-3D diagnoses.

We have displayed below some images belonging to datasets 1 and 2 to see these results visually, showing the diagnosis done by the human observer, and the automatic 2D and 3D diagnoses. Lung 1 presents some areas with abnormally high density tissue, probably corresponding to the initial stage of carcinoma. Lung 2 presents abnormally low greyscale value in some regions, corresponding to emphysema. Lung 3 is almost healthy, presenting only a few isolated areas with high grey levels. Lung 4 presents a fibrous state.

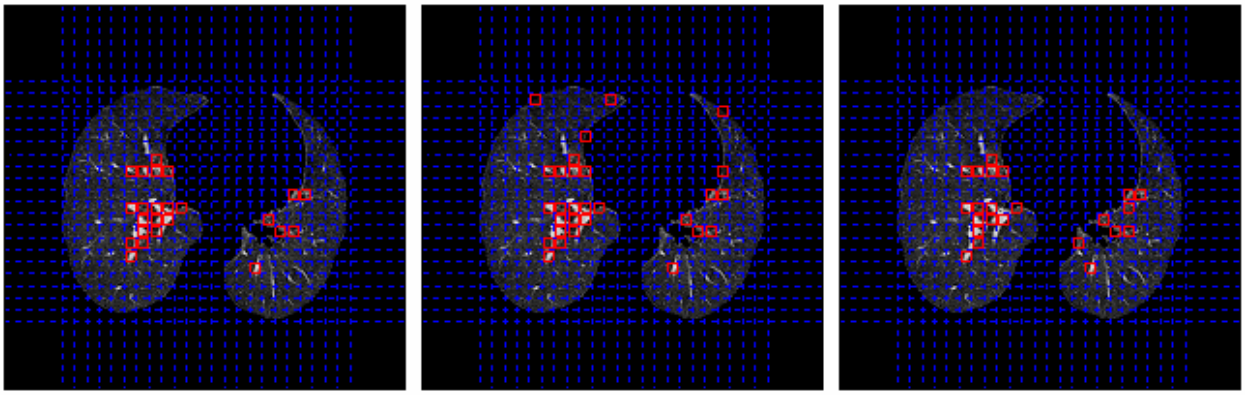

Human observer 2D Diagnosis 3D Diagnosis, α=0.1

Figure 7.11. Diagnosis comparison. Dataset 1.

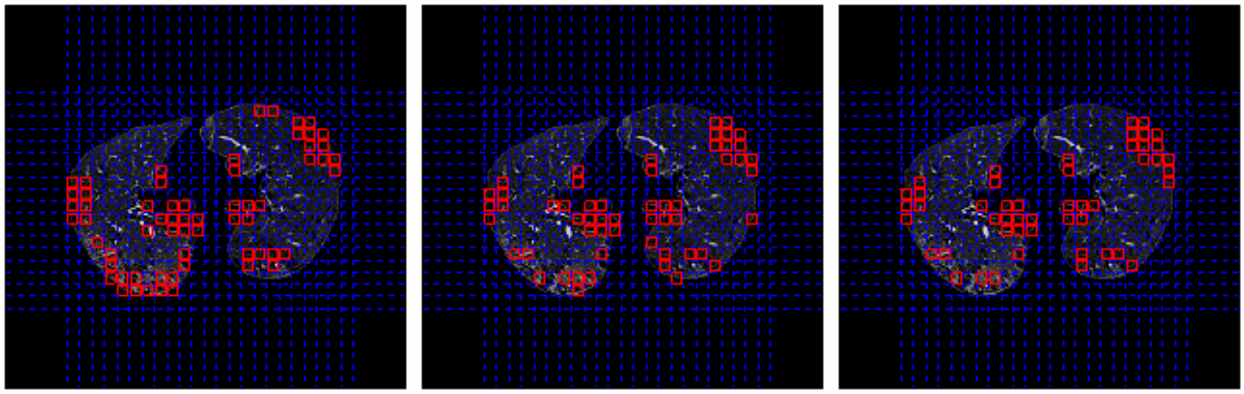

Human observer | 2D Diagnosis | 3D Diagnosis, α=0.1

Figure 7.12. Diagnosis comparison. Dataset 2.

The improvement in the first dataset is clearer than in the second. We can see how those isolated regions which appeared marked as unhealthy in the upper part of the lungs in the 2D diagnosis disappear in the 3D diagnosis, which is nearer to reality. On the other hand, in the second dataset the improvement is difficult to see visually, because it is actually very low, only a 2 or 3%.

7.4. SUMMARY AND CONCLUSIONS

An attempt to design an automatic diagnostic tool using 3D information contained in subsets of 3 lung slices has been carried out in this Chapter. A new graphical user interface has been presented, adding some modifications and new elements to the 2D interface, in order to display and deal with whole datasets of CT lung scans.

A study on 4 different datasets has been developed, in an attempt to find the best relation amongst the parameters obtained after simulations in each one of the three slices. Finally, equation 7.2 illustrates this relation, where we can see that the final diagnosis is mostly based on the information comprised in the slice under study, and the upper and lower slices contribute to the result in a small proportion.

The investigation only yields interesting results when images belonging to the central part of the lungs are considered. A little percentage of improvement compared to the 2D diagnosis has been discovered in this case, although the results are not always clear visually.

8. CONCLUSIONS AND FUTURE WORK

8.1. CONCLUSIONS

Two automatic diagnostic tools for detecting abnormalities in the lung parenchyma, for 2D and 3D, have been developed and their performances investigated. A process of features selection was carried out, yielding a subset of 6 features for the inner parts of the lungs, and of 3 features for the boundaries.

4 different cascade forward networks were trained and tested in lungs with emphysema and fibrosis, yielding a performance above 90% in most of the cases, and showing a slight improvement when a 3D version of the method is applied to slices belonging to the central parts of the lungs.

The computational cost of the procedure was proved to be reasonable in the 2D processes, but the expense of time increased considerably when a 3D diagnosis was carried out.

The results of this research can be included into a more complex computer-aided diagnostic system, adding investigations into new diseases to help the expert physicians in the process.

8.2. FUTURE WORK

There are many open questions for further discussion. First, many more datasets are needed to enhance the robustness of the techniques and the estimation of their performance. Next, a more extensive exploration of new features, such as Markov random fields and especially Gabor filters, could be carried out and might be more successful. Different techniques to improve the accuracy of some of the features already existent in this project, such as the grey-level fractal dimension and the optimal thresholds for binary images can be investigated as well. Another new line of work could be making new features based on the results of the investigations, and also we could start a research using volumetric features in 3D, instead of only 2D measures. However, as we have seen, the feature selection problem is very large and time consuming and will be even more complex with the addition of more cases.

There are also many points of interest regarding the neural networks. There is too much freedom to build this type of structures, and our architecture could be modified in many ways. Adding new hidden elements or changing the number of epochs or the functions that we have used for training, learning, initialization of the net, performance, etc. can lead to an improvement in the results. Although in an initial attempt we were not able to integrate all the neural networks into a single one, it would be a very interesting aspect to consider, because it would reduce the complexity of the system.

9. REFERENCES

BOOKS AND ARTICLES:

[1-1] Schwartz, W.B., **"Medicine and the Computer: The Promise and Problems of Change"**, New Engl. J. Med. 283, 1970

[1-2] Winston, P. W., **"Artificial Intelligence"**, Addison-Wesley, Reading, Mass., 1977

[1-3] Chen, C. H., L. F. Pau, P. S. P. Wang (eds.), **"The Handbook of Pattern Recognition and Computer Vision"** (2nd Edition), pp. 207-248, World Scientific Publishing Co., 1998

[1-4] Vera Valencia, M., **"Análisis, evaluación de parámetros físicos en imágenes de resonancia magnética. Desarrollo de algoritmos de procesado aplicado al diagnóstico de lesiones neurológicas"**, Proyecto Fin de Carrera. Escuela Superior de Ingenieros de la Universidad de Sevilla 2002. Director: Emilio Gómez González..

[1-5] Manjunath, B., **"Texture features for browsing and retrieval of image data"**, IEEE Transaction on pattern analysis and machine intelligence, 8, pp. 837-842, 1996

[1-6] Coggins, J. M., **"A Framework for Texture Analysis Based on Spatial Filtering"**, Ph.D. Thesis, Computer Science Department, Michigan State University, East Lansing, Michigan, 1982

[1-7] Haralick, R.M., **"Statistical and structural approaches to texture"**, Proceedings of the IEEE, V67, No. 5, pp. 786-804, 1979

[1-8] Ahuja, N., and B.J. Schachter, **"Image models"**, Computing Surveys, V13, pp. 373-397, 1981

[1-9] Mandelbrot, B. B., **"The Fractal Geometry of Nature"**, Freeman, San Francisco, 1983

[1-10] Harms, H., U. Gunzer, and H. M. Aus, **"Combined Local Color and Texture Analysis of Stained Cells"**, Computer Vision, Graphics, and Image *Processing*, 33, pp.364-376, 1986

[1-11] Landeweerd, G. H. and E. S. Gelsema, **"The Use of Nuclear Texture Parameters in the Automatic Analysis of Leukocytes"**, Pattern Recognition, 10, pp. 57-61, 1978

[1-12] Lundervold, A., **"Ultrasonic Tissue Characterization - A Pattern Recognition Approach"**, Technical Report, Norwegian Computing Center, Oslo, Norway, 1992

[1-13] Sutton, R. and E. L. Hall, **"Texture Measures for Automatic Classification of Pulmonary Disease"**, IEEE Transactions on Computers, C-21, pp. 667-676, 1972

[1-14] Levitzky, M. G., **"Pulmonary Physiology"**, 5th ed., McGraw-Hill, 1999

[1-15] Mitchell, R.S., T.L. Petty, and M.I. Schwarz, **"Synopsis of clinical pulmonary disease"**, 4th ed., The C.V. Mosby Company, 1989

[1-16] Cotes, J.E., **"Lung Function, Assessment and Application in Medicine"**, 5th ed., Blackwell Scientific Publications, 1993.

[1-17] West, J.B., **"Pulmonary Physiology and Pathophysiology"**, 1st ed. Lippincott Williams and Wilkins, 2001

[1-18] Armstrong, P., A. G. Wilson, P. Dee, **"Imaging of Diseases of the Chest"**, Year Book Medical Publishers, Inc., 1990

[1-19] Horwood, A., **"Computer Diagnosis of Tomographic Pulmonary Images"**, Phd Thesis, Department of Artificial Intelligence, University of Bristol, 2000

[1-20] Chan, H. P., B. Shahiner, M.A. Helvie, N. Petrick, M.A. Roubidoux, TE Wilson, D.D. Adler, C. Paramagul, J.A. Newman, and S. Sanjay-Gopal, **"Improvement of radiologists' characterization of mammographic masses by using computer aided diagnosis: a roc study"**. Radiol, 212(3), pp. 817–827, 1999

[1-21] Tu, J.V., **"Advantages and disadvantages of using artificial neural networks versus logistic regression for predicting medical outcomes"**, J Clin Epidemiol, 1996

[1-22] Reed, R. T., 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

[1-23] Sonka, M., V. Hlavac and R.Boyle, **"Image processing, Analysis and Machine Vision"**, International Thompson Computer Press, 1996

[1-24] Gonzalez and Wintz, **"Digital Image Processing"**, Addison-Wesley Publishing Company, Sec. Ed., 1987

[1-25] Pavlidis, T., **"Algorithms for Graphics and Image Processing"**, Computer Science Press, New York, 1982

[1-26] Gotlieb, C. C., and H.E. Kreyszig, **"Texture descriptors based on co-occurrence matrices"**, Computer Vision, Graphics, and Image Processing, 51, pp. 70-86, 1990

[1-27] Argenti, F., L. Alparone, and G. Benelli, **"Fast algorithms for texture analysis using co-occurrence matrices"**, IEEE Proceedings, Part F: Radar and Signal Processing, 137, pp. 443-448, 1990

[1-28] Shulman, A. R., **"Optical Data Processing"**, John Wiley and Sons, New York, 1970

[1-29] Augusteijn, M. F., Clemens, L. E., and Shaw, K. A., **"Performance evaluation of texture measures for ground cover identification in**

satellite images by means of a neural network classifier", IEEE Trans. Geoscience and Remote Sensing, 33(3), pp. 616-626, 1995

[1-30] Losa G.A., Merlini D., Nonnenmacher T.F., Weibel E.R. (Eds.), **"Fractals in Biology and Medicine"**, Vol II, Birkhauser, Basel, Boston, Berlin, 1995

[1-31] Pentland A.P., **"Fractal-Based Description of Natural Scenes",** IEEE Trans. Patt. Anal. Mach. Intel. PAMI-6(6), pp. 661-674, 1994

[1-32] Barlett M.L., **"Comparison of methods for measuring fractal dimension"**, Austral.Phys.Eng.Sci.Med. 14(3), pp.146-152, 1993

[1-33] Faugeras. O., **"Three-Dimensional Computer Vision: A Geometric Viewpoint"**, M.I.T. Press, 1993

[1-34] Papoulis, A., **"Probability, Random variables and Stochastic Processes"**, McGraw Hill, New York, Third Edition, 1991

[1-35] Tomita, Fumiaki and S. Tsuji, **"Computer Analysis of Visual Textures"**, Kluwer Academic Publishers, Boston, 1990

[1-36] Hertz, J., A. Krogh, and R. G. Palmer, **"Introduction to the Theory of Neural Computation"**, Addison-Wesley Publishing Company, 1991

[1-37 Veelenturf, L. P., **"Analysis and Applications of Artificial Neural Networks"**, Prentice Hall, 1995

[1-38] Haykin, S., **"Neural Networks, A Comprehensive Foundation"**, sec. ed., Prentice Hall, 1999

[1-39] Hagan, M., H. Demuth and M. Beale, **"Neural Network Design"**, The Brooks/Cole Publishing Company, 1996

[1-40] Baxt, W.G., **"Application of artificial neural networks in clinical medicine"**, Lancet, 1995

[1-41] McCullogh, W.S. and Pitts, W., **"A logical calculus of the ideas immanent in neurons activity"**, Bulletin of Mathematical Biophysics, 5, pp. 115-133, 1943

[1-42] Minsky, M. L., and S.A. Papert, **"Perceptrons"**, 2nd Edition, M.I.T., 1988

[1-43] Rumelhart, D., and J. McClelland, **"Parallel Distributed Processing"**, MIT Press, Cambridge, 1986

[1-44] Smith, M., **"Neural networks for statistical modelling"**, Van Nostrand Reinhold, New York, 1993

[1-45] Vogl, T.P., J.K. Mangis, A.K. Risler, W.T. Zink, and D.L. Alkon, **"Accelerating the convergence of the back-propagation method"**, Biological Cybernetics, Vol. 59, pp. 257-263, 1988

[1-46] Nguyen, D. et al. **"Improving the learning speed of 2-layer neural networks by choosing initial values of the adaptive weights"**, International Joint Conference on Neural Networks (conference paper), San Diego, CA, USA, 1990

[1-47] Ricketts, I. W., **"Cervical cell image inspection - a task for artificial neural networks"**, Network: Computation in Neural Systems, Vol. 3, pp. 15-18, 1992

[1-48] Pernkopf, F., and O'Leary, P., **"Feature Selection for Classification Using Genetic Algorithms with a Novel Encoding"**, University of Leoben, Austria, 2001

[1-49] Dash, M., Liu, H., **"Feature selection for Classification"**, Intelligent Data Analysis, 1, pp. 131-156, 1997

[1-50] Gong, G., **"Cross-validation, the jacknife, and the bootstrap: excess error estimation in forward regression logistic regression"**, J Am Statistical Assoc., 81(393), pp. 108–113, 1986

[1-51] Hush, D. R., and B. G. Horne, **"Progress in supervised neural networks: What is new since Lippmann?"**, IEEE Signal Processing Magazine, vol. 10, pp. 8-39, 1993

[1-52] Foley, D.H., **"Consideration of sample and feature size"**, IEEE Transactions on Information Theory, IT-18(5): pp. 618-626, 1972

INTERNET REFERENCES:

[2-1] Szolovits, P., **"Artificial Intelligence and Medicine"**, Westview Press, Boulder, Colorado, 1982

http://medg.lcs.mit.edu/ftp/psz/AIM82/ch1.html

2-2] Chantler, M. J., **"The effect of variation in illuminant direction on texture classification**", Ph.D. thesis, Dept. Computing and Electrical Engineering, Heriot-Watt University, 1994

www.cee.hw.ac.uk/~mjc/texture/mjc-phd/ch4.pdf

[2-3] Douglas, E., G.A. Johnson, G.E. Trahey, **"Computer Aided Diagnosis of Acute Pulmonary Embolism"**, Technical Report, Department of Biomedical Engineering, Duke University, 2000

www.duke.edu/~edf/thesis-final-color.pdf

[2-4] Kelly, P.M., and T.M. Cannon, **"CANDID: Comparison Algorithm for Navigating Digital Image Databases"**, Computer Research and Applications Group, CIC-3- Los Alamos National Laboratory, NM, 87545, 1994

www.c3.lanl.gov/~kelly/ml/pubs/1994_ssdbm/paper.pdf

[2-5] Sonka, M., R., Uppaluri, G. McLennan and E.A. Hoffman, **"AMFM: Adaptive Multiple Feature Method"**, Technical Report, Department of Medicine and Radiology, Universities of Iowa and Arizona, 1998

www.jradiology.com/arts/22.pdf

[2-6] Klonowski, W., **"Signal and Image Analysis using Chaos Theory and Fractal Geometry"**, Lab. of Biosignal Analysis Fundamentals Institute of Biocybernetics and Biomedical Engineering, Polish Academy of Sciences, 2000

http://hrabia.ibib.waw.pl/~lbaf/PDF_Doc/gkpo2000.pdf

[2-7] Fisher, R. B., **"CVonline: The Evolving, Distributed, Non-Proprietary, On-Line Compendium of Computer Vision"**, School of Informatics, University of Edinburgh, 2003

http://www.dai.ed.ac.uk/CVonline/

[2-8] Friman, O., M. Borga, M. Lundberg, U. Tylen, H. Knutsson, **"Recognizing Emphysema-A Neural Network Approach"**, Dept. of Biomedical Engineering, Linkoping University and Dept. of Radiology, Goteborg University, Sweeden ,2002

http://www.imt.liu.se/mi/Publications/Papers/ICPR02emphysema.pdf