Image Normalization, a Basic Requirement for Computer-based Automatic Diagnostic Applications.

Dr Alan C Horwood BA, PhD (A.C.Horwood@Bristol.ac.uk) Tel. +44 (0)117 9743

Professor S J Hogan MA, PhD

Professor P R Goddard BSc, MD

Dr Jonathan Rossiter PhD

Department of Engineering Maths, University of Bristol, Clifton, Bristol BS8 1TR, UK.

May 2001

Purpose: to formulate a method for normalizing computed tomographic (CT) lung image data as a preparation for computer-based automatic, or semi-automatic, diagnostic applications.

Materials and Methods: histograms of greyscale distributions in comparable thoracic image slices from eight CT data-sets showed different modal values for normal, constituent tissues. In a given data-set, the usually consistent modes for muscle tissue, fatty tissue, spinal process and the descending aorta have a close correlation with the brightness increase necessary to bring an anterior 50x50 image region of visually normal parenchyma to a modal greyscale value of 35 – an arbitrarily chosen normal reference value. A straight line equation relates the mode for muscle tissue in a data-set with the required percentage brightness correction. The validity of the processing was tested using the 'information dimension' of noise-reduced pixel patterns, created when standard upper and lower greyscale thresholds are applied to 50x50 regions to confine the values closely around the normalized mode. An empirically based information dimension ≥ 1.85 is taken as an indicator of 'health'.

Results: the criterion information dimension is a useful index of normal lung parenchyma in normalized CT data-sets.

Conclusion: image normalization is a prerequisite for computer-based diagnosis of CT pulmonary images.

Key terms: computed tomography (CT), densitometry, greyscale distribution, statistical mode, information dimension, correlation coefficient.

Image Normalization, a Basic Requirement for Computer-based Automatic Diagnostic Applications.

INTRODUCTION

Reliable software for automatic, objective assessment of, for example, lung parenchyma from x-ray CT would prove a most useful diagnostic tool. A number of projects have been dedicated to finding practical applications towards this end. [1,2,3,4,5,6,7,8,9]

Regular images and standardized processing are a sine qua non for the success of any suggested approach. However, irregularities in image data have been found and recognized as a possible source of error in tissue classification – especially where data has been acquired indirectly (from archives or a series of 'downloads') or from using different scanning protocols and/or scanners. [10,11,12,13]

Hounsfield units and image greyscale values

Hounsfield units provide a scale for measuring local tissue density recorded as the reduction in x-ray intensity per unit thickness of material. The scale is arbitrarily defined such that air, water and fully calcified bone have values of –1000, 0 and +1000 (Hu) respectively. CT images are reconstructed using a greyscale contrast range of 0 for black pixels and 255 for white. The image pixels are points of varying brightness that depend directly on the attenuation coefficients. Studies show a good linear correspondence between Hu and greyscale values – particularly over the range of tissue densities involved in this study. [1,11] A simple heuristic, using the straight line plot obtained by setting –1000 Hu = 0 greyscale and +1000 Hu = 255, produces a possible conversion equation, $y_{(\text{greyscale})} = 127 + 0.1275x_{(\text{Hu})}$. In normalized images, the typical greyscale mode for fat is 115, which gives a corresponding Hu value of –94; for

soft tissue it is 132, which converts to 39Hu; and for bone anything between 140 and 255, which corresponds to Hu range 101 to 1000. The quoted Hu values are, typically, –100 to –80, 22 to 46, and 80 to 1000 respectively. [1]

Computer analysis of a medical image derived from densitometry depends crucially on the recognition of a specific range within a greyscale distribution that is indicative of a particular physical condition. Specially important is the modal value representing 'normal' in a given location within a thoracic region, appropriate for the patient's posture during scanning. CT scans are usually (though not always) taken with the patient in the supine position, which means that, among other things, effects of pressure and gravity on organs such as the lungs must be taken into account. Differences in patients' size and weight affect image quality. Other complicating factors include the use of a contrast medium and variable scanning parameters of range and level. [1,11,14]

In the dissertation, 'Computer Diagnosis of Tomographic Pulmonary Images' [11], the point is made that it is sometimes quite a small deviation from the 'normal' modal greyscale ranges that can signal a physical abnormality. For training the computer, 50x50 image regions are selected (referred to as 'regions of interest' or 'roi's') that are known to contain a significant number of pixels representing lower-than-normal, normal, or higher-than-normal tissue density. From each of these classes, three pixel patterns (plots) are segmented using empirically obtained upper and lower threshold values to contain, respectively, the characteristic density-related greyscale distribution ranges. A histogram of the greyscale distribution in the pixel plot shows its mode(s). However, as already suggested, though one can expect a low-deviation modal greyscale distribution in images of most lesion-free parenchyma – certainly in the lung fields – the actual mode number is not guaranteed to be the same from one data-set to another, even

for similar locations. [11] The term ‘normalization’ refers to an image-processing method aimed at correcting such differences by shifting the relative brightness values up or down before assessing the number of pixels contained within the threshold limits and the information dimension characterizing the resulting pattern. The actual degree of brightness correction needed is determined by reference to a standardized value. How one arrives at this ‘standard’ value is the subject of this paper.

MATERIALS AND METHODS

Fractal estimator

The information dimension (see appendix) is used as the principle estimator for characterizing processed regions from digitized CT lung-slice images. Individual slices, 512x512 format, are extracted from complete data-sets which are derived chiefly, though not exclusively, from spiral scans; (source: Siemens Somaton Plus 4 whole-body spiral scanner). ‘Regions of interest’, 50x50 greyscale pixel plots, are selected by the software to cover, progressively, the entire lung fields. Each region is then processed as three black and white images, using three pre-determined sets of upper and lower threshold limits: values chosen to ‘capture’ greyscale distributions that characterize the three different tissue densities referred to above. These thresholded distributions must be consistent for any subsequent analysis to be reliable.

Information dimension is used as an estimator in preference to capacity dimension based on straightforward box-counting, because it is more sensitive to the pattern of pixel distribution in thresholded images. Other workers researching methods of estimating fractal dimensions of image regions have discovered the limitations of simple box-counting methods, especially where images are somewhat restricted in data.

Impulsive noise and scatter in images – which is not uncommon – present other problems. [8,11,14]

One approach to improving the quantitative characterization of information in images depends on extracting a large number of features – as in the Adaptive Multiple Feature Method developed at the University of Iowa, [6,7] – though most workers have found some application of fractal measure to be a 'best feature'. [10,13] This paper demonstrates a method that relies upon pre-processing of selected regions of interest to normalize pixel brightness to a chosen 'standard' based on 'normal' tissue density. This, allied to image smoothing to eliminate 'noise', allows data to be effectively characterized using the information dimension as the chief discriminating feature.

Differences among data-sets

The images in figure 1 illustrate the irregularities in greyscale distribution for similar tissues in selected slices, one from each of eight data-sets. The thoracic location is much the same in each. Experiments have shown that such irregularities occur between one set and another, rather than between one slice and another in a particular data-set. This observation is based on an analysis of the eight sets, comprising 418 slices. [11] The main exception occurs either where the enhancement by contrast medium happens partway through a scan, or its application is uneven. In the histograms, the spike nearest zero shows the sum of pixels resulting from x-ray attenuation detected beyond the chest wall and from any very low attenuation values not recorded as data – as determined by the CT range and level settings. The other 'modal' distributions occurring progressively towards the higher greyscale values correspond, respectively, to the lung parenchyma, fatty tissue and muscle tissue. Compacted material, such as bone, is represented by the non-modal greyscale distribution beyond that of muscle.

Fig. 1: histograms of image slices from around anatomical level T6 in eight data-sets.

Lung 1, slice 30.

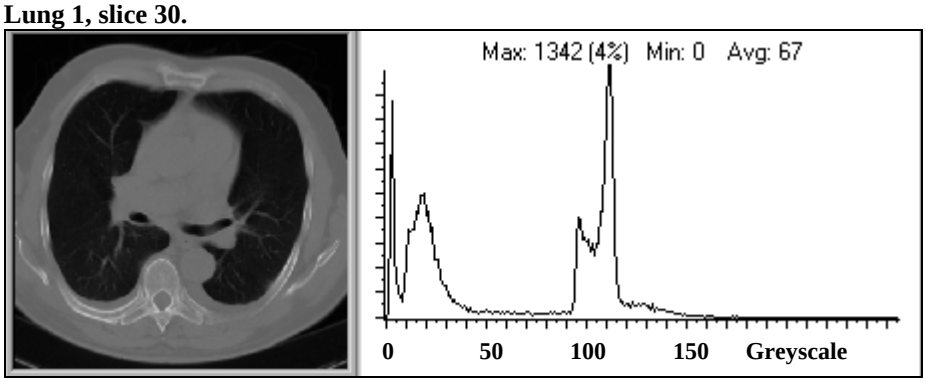

Lung 2, slice 39.

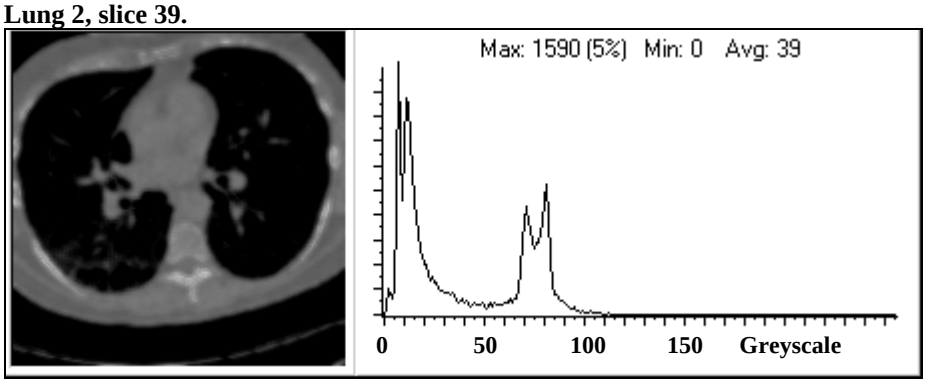

Lung 3, slice 32.

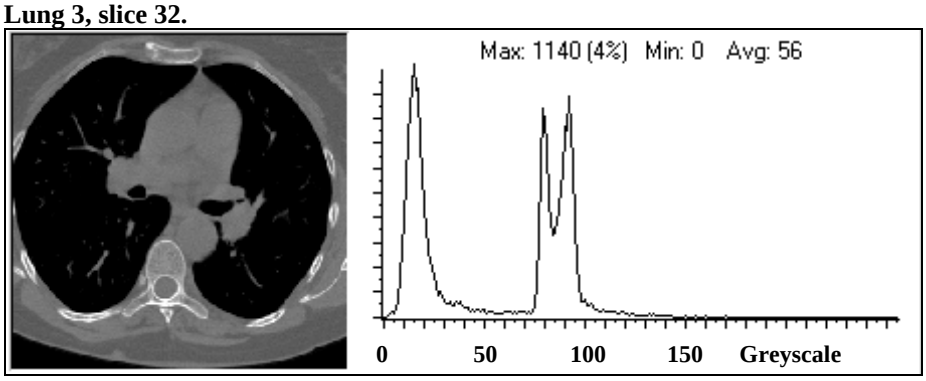

Lung 4, slice 22.

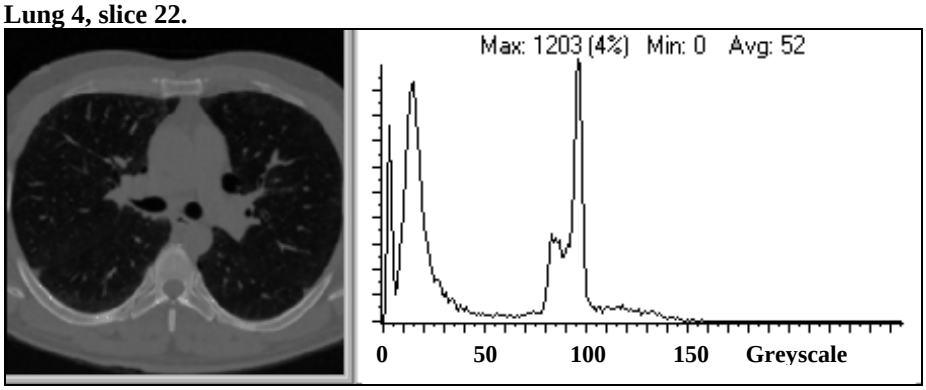

Lung 5, slice 23.

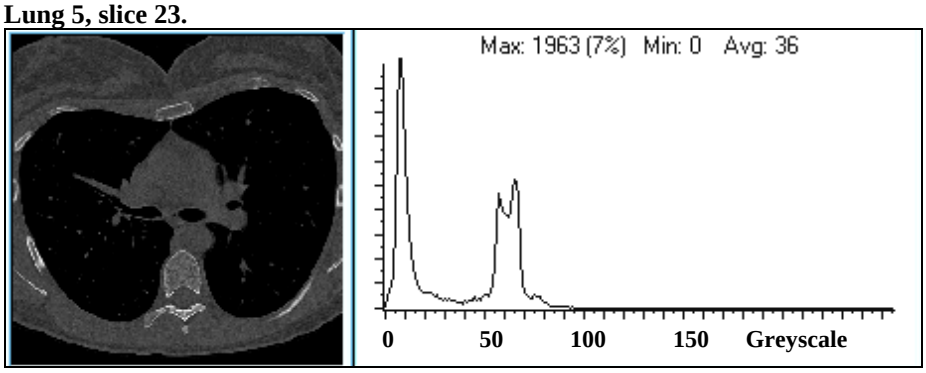

Lung 6, slice 20.

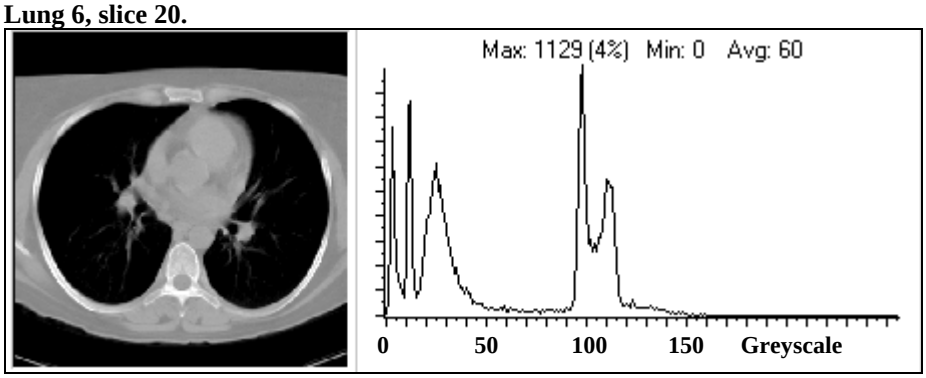

Lung 7, slice 24.

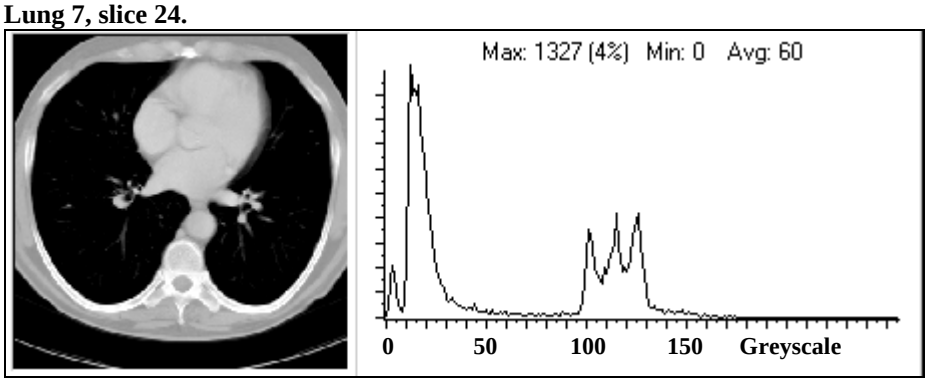

Lung 8, slice 30.

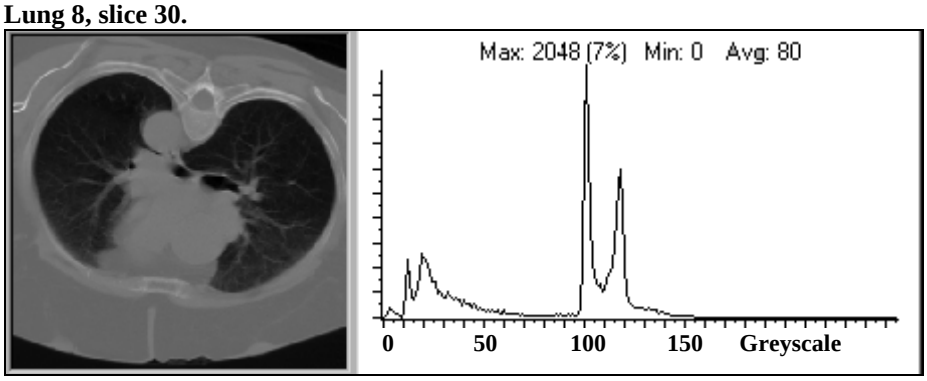

These images show that, though the profiles of greyscale distribution in the selected slices have some broad similarities, there are significant differences in the characteristic

modal values for the common tissue types. Table 1 contains the modal values for normal parenchyma, fatty and muscle tissue etc. for the slices illustrated in figure 1; table 2 provides statistics for the complete data-sets.

Table 1: modal greyscale values for images in figure 1.

Lung - slice	Muscle	Fat	M – F	Spine	Aorta	Parenchyma: ant.	post.
1 - 30	111	98	13	110	110	15	22
2 - 39	81	73	8	81	81	11	–
3 - 32	93	79	14	90	91	13	18
4 - 22	96	83	13	97	95	14	18
5 - 23	66	58	8	66	65	6	10
6 - 20	112	99	13	111	111	–	–
7 - 24	115	102	13	112	120	15	17
8 - 30	116	102	14	117	117 *	21	18

Key: (i) M – F is the difference between the modal values for muscle and fatty tissue.

(ii) Dashes indicate an abnormal region; a possible interpolation figure for 'would-be' normality can be gauged from appropriate data in table 2.

(iii) * The scanning for data-set 8 was carried out in the prone position.

Table 2: statistical details based on modal greyscale values for the eight complete data-sets.

Data-set:		1	2	3	4	5	6	7	¶ 8
Muscle:	max.	113	83	92	97	68	112	116	118
	min.	109	81	90	92	65	111	115	114
	avg.	111	82	91	94	66	111	115	117
	std.	0.98	0.58	0.75	1.96	0.64	0.45	0.43	1.07
	avg.d.	0.63	0.42	0.64	1.80	0.32	0.41	0.38	0.75
Fat:	max.	99	78	79	85	59	99	104	102
	min.	96	71	79	83	58	97	102	102
	avg.	98	73	79	84	58.3	98	103	102
	std.	0.76	1.95	0.0	0.80	0.43	0.61	0.86	0.0
	avg.d.	0.65	1.47	0.0	0.72	0.38	0.50	0.78	0.0
M – F:	avg.	13	10	12	10	7.7	13	12	15
Spine:	max.	111	82	92	97	67	111	115	118
	min.	109	81	90	92	65	111	112	115
	avg.	110	81.5	91	94	66	111	114	117
	std.	0.45	0.51	0.80	1.62	0.80	0.0	0.99	0.84
	avg.d.	0.44	0.49	0.72	1.36	0.69	0.0	0.78	0.66
Aorta:	max.	116	84	92	97	65	112	130	118
	min.	106	80	90	95	65	110	120	116
	avg.	112	82	91	95.5	65	111	124	117
	std.	2.35	1.06	0.75	0.67	0.0	0.71	3.80	0.51
	avg.d.	2.11	0.75	0.64	0.52	0.0	0.50	3.22	0.49
Anterior:	max.	16	12	15	15	11	20	17	# –
(normal)	min.	12	9	11	12	6	16	14	# –
	avg.	14	11	14	14	8.4	18.5	15	# –
	std.	1.32	1.06	1.55	1.00	1.55	1.20	0.90	# –
	avg.d.	1.02	0.93	1.20	0.67	1.42	1.12	0.61	# –
Posterior:	max.	26	18	20	16	12	*27	19	26
(normal)	min.	22	16	16	14	12	*25	15	17
	avg.	23	17.5	18	15	12	*26	17	22
	std.	2.28	0.88	1.33	1.63	0.0	0.71	1.40	3.15
	avg.d.	2.01	0.88	1.04	1.33	0.0	0.63	1.22	2.90

Key: (i) ¶ The scanning for data-set 8 was carried out in the prone position. It follows that the anterior regions are more gravity affected than posterior ones.

(ii) # Too few normal regions in data-set 8 to give meaningful values.

(iii) * Too few normal posterior regions in the right field; readings refer to the left field only.

(iii) 'Avg.d.' is the average deviation, a measure of the variability in a data-set.

(iv) 'Std' is the standard deviation, a measure of how widely values are dispersed from the average value (mean) of the data-set.

The data in table 1 can be used to determine a possible correlation between the modal value of the greyscale distribution for a measurable image feature and the modal values characterizing normal tissue in anterior and posterior regions. An important observation is the closeness in values obtained for muscle tissue, the spinal process and the aorta in a selected slice. These regions will be affected by any use of a contrast medium which, of course, would also affect the radio-density of the parenchyma. The M – F value (see key for table 1) serves as an indicator of the relative effect of any contrast medium, since fatty tissue will not be affected. From this, one might deduce that lung slice 2 – 39 and 5 – 23 result from scanning without the use of contrast, and so the imaged parenchyma, where normal, should have lower modal greyscale values than would otherwise be the case. The data shows this to be so, though slice 5 – 23 has values considerably lower than slice 2 – 39. This difference holds for all the data in table 2 regarding sets two and five.

Overall, these data demonstrate the general consistency in modal greyscale values for similarly located normal tissues throughout a full thoracic scan. However, though this regularity holds for a given data-set, there are considerable differences between one set and another – hence the need for 'normalization'.

Normalization

The basis of normalization is a correlated change of brightness level in a selected region of interest to bring the modal greyscale value for its 'normal' (or would-be normal) parenchyma to a standardized reference value. [11] The selection of this reference value is determined from inspection of data such as that in table 2. It provides a working 'norm' for a particular location within the slice. In this study, the decision involves either anterior or posterior regions for scans taken with the patient in the supine position. The anterior value is probably more reliable as a reference, since patients of widely differing weights will contribute variously to the gravity and pressure effects on dependent tissue during scanning; non-dependent regions are affected much less overall. For example, a glance at the average values for anterior and posterior locations reveals a small effective difference between their densities in sets 4 and 7 compared with those of sets 1 and 2.

One factor informing the choice of a greyscale reference norm is the preference for adjusting brightness values upwards rather than down in order to keep all image data accessible. Any greyscale values reduced to zero lose all information. From the available data, a modal value of 35 (approx. corresponding to –722 Hu) for a normal anterior region looks to be a good practical choice; then brightness adjustments can be found heuristically to bring corresponding regions from each data-set to that value. Visual inspection of the selected region must serve to determine tissue normality. Table 3 shows the results using the slices listed in table 1 – except for set 6, where slice 20 was not normal. In slice 8 – 20, the value for the normal posterior region is used because the patient was scanned in the prone position.

Table 3: % brightness increases for normalization.

Lung - slice	Mode shift	% Brightness increase
1 - 30	15 → 35	8
2 - 39	11 → 35	10
3 - 32	13 → 35	9
4 - 22	14 → 35	9
5 - 23	6 → 35	12
6 - 12	16 → 35	8
7 - 24	15 → 35	8
8 - 30	18 → 35	6

The task now is to find whether a set of extracted feature values correlates satisfactorily with the required brightness shift. Since tables 1 and 2 show similar values for muscle tissue, the spinal process and the aorta in a given set, any of these features could be used. In a correlation exercise, the correlation coefficient for the slice and whole data-set average for muscle values was returned as 0.998. This gives confidence in the use of a single slice to find a reference mode. Further, correlating the muscle values with the corresponding brightness increases necessary for standardizing the normal anterior modes (as set out in table 3) gives a coefficient of – 0.98, whether one uses the single slice or average data-set values; (a minus coefficient because the sets of values are inversely related).

Thus it seems justified to assume a linear relationship between the modal greyscale values for muscle tissue and the empirically derived brightness adjustments required for normalization – certainly for the relatively small range of greyscale values involved. A brightness conversion graph, with its line equation of the form $y = mx + c$, can now be constructed by fitting a straight line (using the method of least squares) to values along a linear trend implied by the known values. The brightness changes are returned along this line for the modified ‘muscle’ values in this linear trend. The parameter values, ‘m’ and ‘c’, are easily found. The graph is shown as figure 2.

Fig. 2: a 'best fit' straight line relating empirically derived brightness 'normalizing' corrections and a linear trend of modal greyscale values of muscle tissue from a series of CT lung scans.

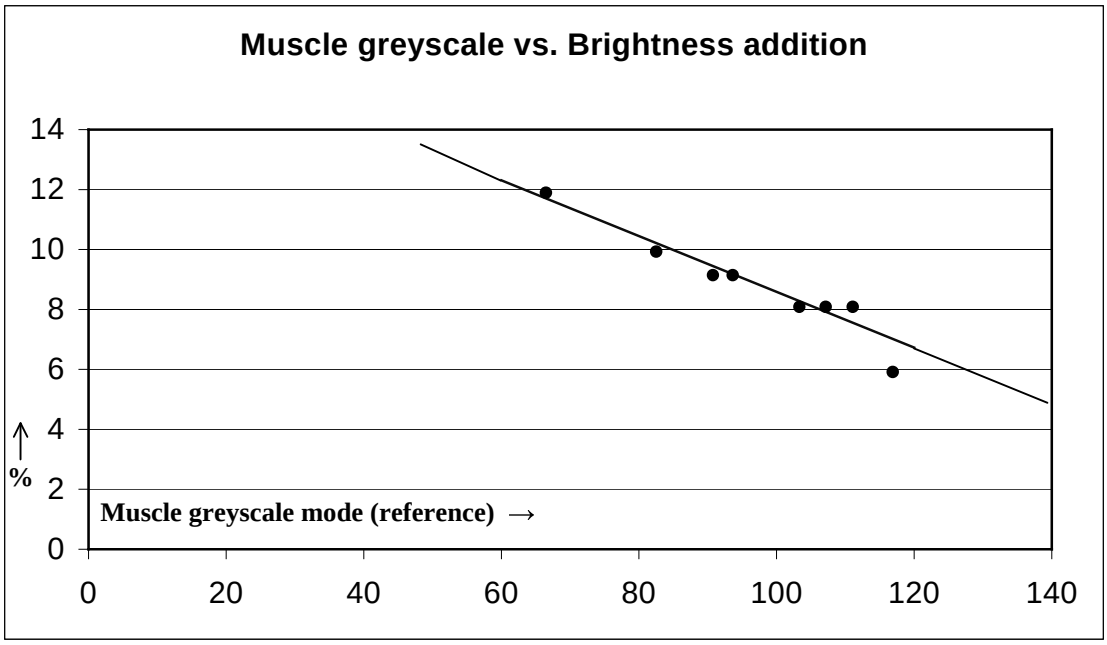

The parameter 'c' has value 17.68 and 'm' – 0.096

The line equation is, therefore, $y = 17.68 - 0.096x$; (where 'y' is the % brightness correction, and 'x' the modal greyscale value for muscle tissue in a sampled pulmonary CT data-set.)

Image smoothing

In image processing, smoothing is a method aimed principally at noise reduction. Employing the simplest algorithm results in some 'fuzzing' of edges in the image, hence its application is sometimes referred to as 'blurring'. The basic operation involves replacing the central pixel in a specified mask with one having the greyscale average of its neighbourhood values. (The neighbourhood as defined by the mask) The smallest mask involves 3x3 averaging and produces the least blurring. In this study, a standard 3x3 mask is used in all 50x50 imaged regions of interest offered for diagnostic quantization. In images with strongly modal distribution(s), one outcome of this averaging is an increase in the number of pixels falling within the modal distribution(s)

together with a clearer definition of the modes. For example, where modal values for fat and muscle tissue are not clearly defined, smoothing may be necessary to separate them in a histogram when determining the reference value for normalization. Importantly, smoothing does not alter the modal greyscale value(s). (Fig. 3)

Fig. 3: a region from lung slice 4 – 30 containing muscle tissue and fat. Image 'a' has no smoothing applied; image 'b' has had two successive smoothing applications using a 3x3 mask.

(a)

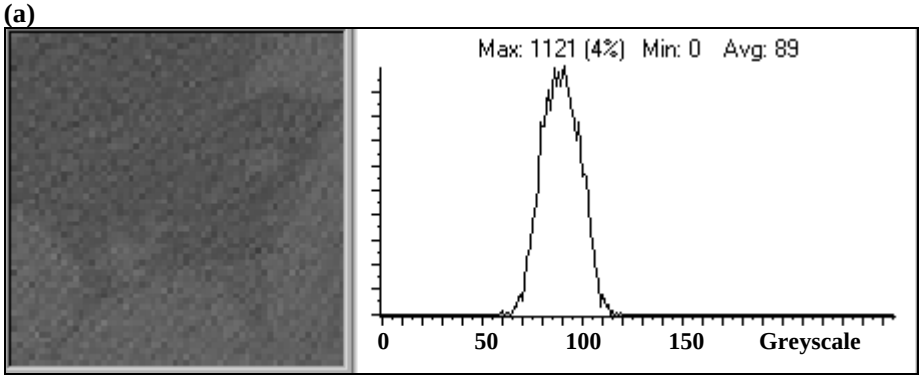

(b)

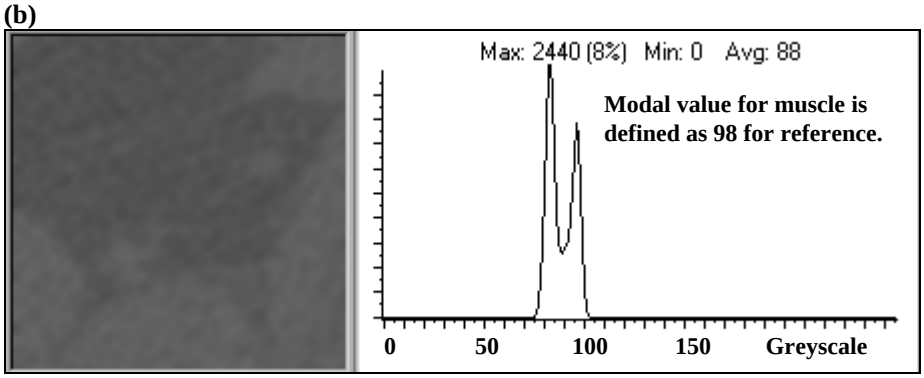

Some impulsive noise will be found in most CT images, which can result in a significant number of black and/or white pixels interfering with the data. Figure 4(a) is a plot of pixels within the greyscale range 0 to 2 in an unprocessed 50x50 anterior region in lung slice 4 – 30. Figure 4(b) shows the plot after the same image has been smoothed.

Fig. 4: greyscale 0 to 2 pixel plots of a 50x50 right anterior region of lung slice 4 – 30. Plot (a) is for the image without smoothing, plot (b) after smoothing with a 3x3 mask.

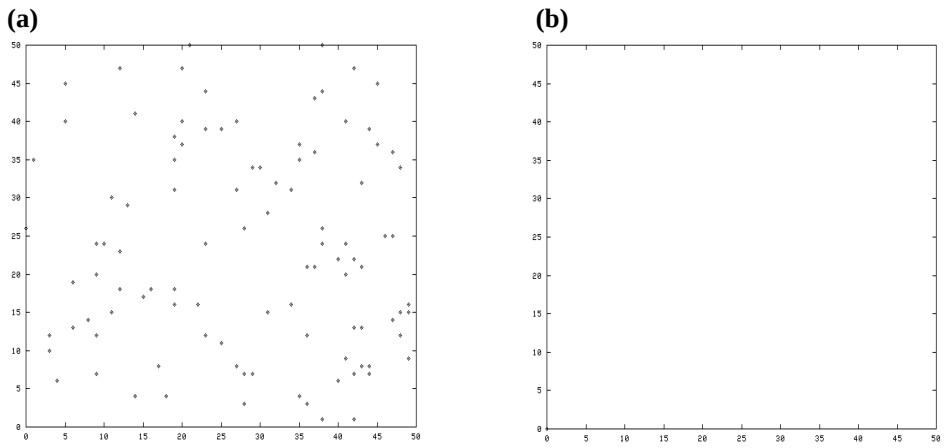

The pixel pattern in figure 4(a) is a random distribution of near-black pixels corresponding to no physiological feature in the scanned lung field; if offered to quantifying software, however, it would be classified as possible low-density tissue. Smoothing removes the pixels into the nearest neighbouring mode. The presence of any emphysematous tissue will appear as a cluster of closely associated low-greyscale pixels, which will not disappear with smoothing. Such regions frequently show bi- or even tri-modal distributions. With smoothing, pixels that are close neighbours of a particular mode will be merged into its characteristic distribution, the mode itself acting as a kind of ‘attractor’. This accounts for the differentiation of fatty and muscle tissue modes, shown in figure 3(b), appearing only after smoothing the distribution shown in figure 3(a).

Figures 5 (a) to (c) show the result of applying the standard image normalizing and smoothing processes to the extracted anterior region referred to above.

Fig. 5: a 50x50 right anterior region from lung slice 4 – 30; image 'a' is as extracted, 'b' with the appropriate brightness increase to normalize and 'c' with smoothing, using a 3x3 mask.

(a)

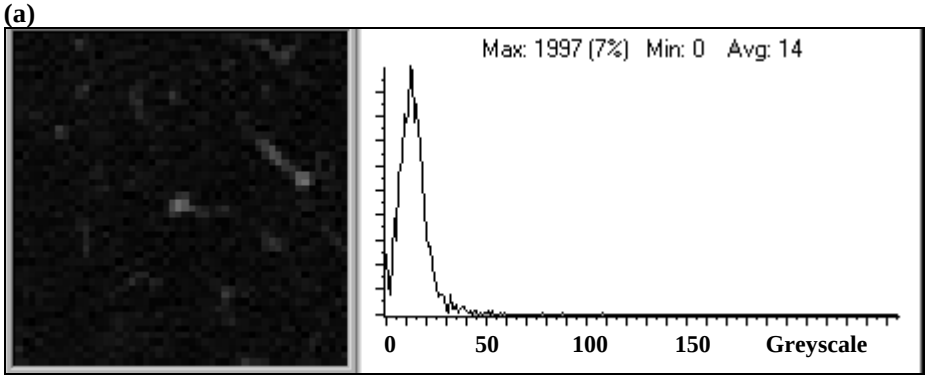

Substitution of the reference value 98 into the line equation, $y = 17.68 - 0.096 \times 98$, suggests a brightness addition of 8.272. (The nearest % brightness addition is 8)

(b)

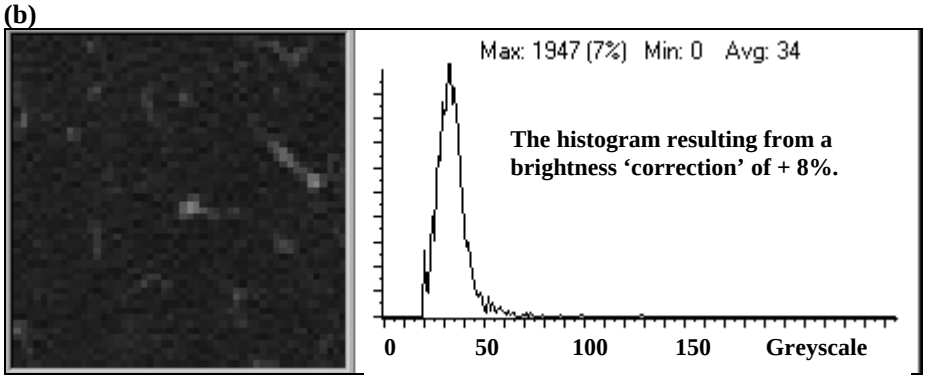

With this brightness addition, the histogram profile is unaltered, but moved about 20 greyscale points to the right. This should normalize the modal value in an image of normal, anterior parenchyma from this set to greyscale 35. Note, the small spike with original mode value 0 (image 'a'), representing noise, has moved to greyscale value 20. This could distort possible tests for emphysema unless smoothed into the normal modal distribution. (See fig. 4, above) The smoothing algorithm, using a 3x3 mask, also produces a more sharply defined mode. (Image 'c')

(c)

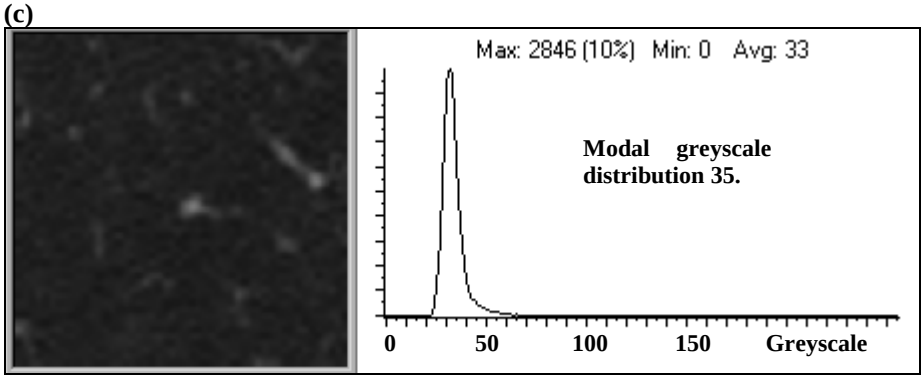

Diagnostic windows

'Window', as used here, refers to a segmentation tool set with precise upper and lower greyscale threshold values. When applied to a normalized region, it creates a quantifiable black and white pattern made of all the pixels from within those limits. The limits are determined heuristically, using the information in images of known physiological type, location and condition. Where a visually 'healthy' anterior region of lung parenchyma has been normalized to a narrow, modal distribution around greyscale 35, it seems reasonable to set threshold limits of, say, 30 and 40 as a first trial for identifying such tissue across all data-sets. Different, empirically determined limits will be needed for posterior regions (when gravity dependent). In the following tests, limits of 35 and 50 are used.

Pixels with greyscale values within the limits might now be called 'normal'; those below the lower limit may represent a typical emphysematous density distribution, but will also comprise any non-registered x-ray attenuations due to the scanning level and range settings; and while those above may result from attenuations due to abnormally dense tissue, there can also be the effect of septa, blood vessels, non-registered attenuations, tissues other than parenchyma etc. Separate 'windows' will be required to

resolve the distributions above and below the 'normal' limits. However, this paper is concerned only with testing a methodology for image normalization as a preparation for reliable, computer-assisted identification of normal anterior and posterior regions. In table 4, the letters 'a', 'b', 'c' and 'd' refer, respectively, to right anterior, right posterior, left posterior and left anterior. (In a supine scan, 'b' and 'c' are 'dependent'.)

RESULTS

Table 4: no. of pixels segmented by diagnostic 'window'.

Lung slice	region	Pixels: in	below	above	Inform. dim.
1 – 16	a	2259	13	129	1.93
	b	1946	0	554	1.87
	c	1995	37	468	1.88
	d	1370	1	1129	1.85
1 – 45	a	2158	16	326	1.90
	b	1825	6	669	1.87
	c	1522	2	976	1.86
	d	2374	0	126	1.93
2 – 33	a	2079	3	418	1.89
	b	821	8	1671	*1.67*
	c	1479	715	306	*1.76*
	d	2311	36	153	1.93
2 – 50	a	2070	11	419	1.88
	b	1196	654	650	*1.75*
	c	1242	1143	115	*1.63*
	d	2277	36	187	1.91
3 – 30	a	2009	28	463	1.90
	b	2358	61	81	1.94
	c	2019	262	219	1.88
	d	2253	8	239	1.91
3 – 48	a	1522	9	344	*1.73*
	b	2026	270	204	1.90
	c	2006	235	259	1.91
	d	1945	44	511	1.86
4 – 10	a	520	27	1953	*1.55*
	b	216	4	2280	*1.09*
	c	699	1	1901	*1.70*
	d	195	0	2305	*1.35*
4 – 15	a	1552	26	922	*1.81*
	b	1443	53	1004	*1.79*
	c	1653	23	824	*1.81*
	d	1127	64	1309	*1.75*
5 – 7	a	2013	19	468	1.90
	b	711	4	1785	*1.67*
	c	2152	73	275	1.89
	d	1925	94	481	1.89
5 – 30	a	2166	24	310	1.91
	b	2324	160	16	1.94
	c	2260	79	161	1.93
	d	2104	23	274	1.91
6 – 12	a	1314	0	1186	*1.82*
	b	1836	0	664	1.85
	c	1873	0	627	1.87
	d	465	1	2034	*1.60*
6 – 30	a	13	0	2487	–
	b	1865	0	635	*1.82*
	c	2050	0	450	1.89
	d	566	0	1934	*1.54*

Key: (i) Pixels 'in' etc. refers to numbers of pixels distributed with reference to the threshold limits taken to signify 'normal'.

(ii) Italicized figures indicate a distribution within those limits characterized by an information dimension below that taken as 'normal'.

7 – 22	a	1912	208	380	1.87
	b	1852	595	53	1.86
	c	2127	251	122	1.90
	d	1923	188	389	1.88
7 – 33	a	1945	268	287	1.87
	b	1168	1263	69	*1.68*
	c	1987	195	318	*1.79*
	d	1959	249	292	1.88
8 – 23	a	864	0	1636	*1.64*
	b	1622	44	834	1.81
	c	2024	0	476	1.88
	d	1715	17	768	1.87

Based on a number of trials, using parenchymal image regions of known normality, an information dimension ≥ 1.85 for pixel distribution within the chosen threshold limits is taken as an indicator of 'health'. Where the imaged 50x50 region is entirely within the lung fields, the 'normal' pixels within the distribution should number at least 2000. A 'good' information dimension assessed for substantially fewer pixels suggests a region partially filled with healthy parenchyma. If the bulk of the 'missing' pixels are below the range, there is the possibility of emphysema; if above, the indication is either the presence of fibrosis or the inclusion within the image of a non-parenchymal region. In general, there will always be some distribution of higher greyscale pixels due to the higher radio-density of larger blood vessels and structural elements, such as septa. There is also the problem of the partial volume effect. [1,11]

The results listed in table 4 suggest the need for closer inspection of image regions 1-16d, 2-33b & c, 2-50b & c, 3-30c, 4-10a, b, c, & d, 4-15 b, c, & d, 5-7b, 6-12a & d, 6-30d, 7-33b & d and 8-34a & d. In each of these, the number of pixels and/or the character of their distribution do not fall within the 'normal' parameters. (See figure 6.)

Of the 60 regions assessed from the eight data-sets, 38 were registered as 'normal', 3 as possibly emphysematous and 19 as 'over dense'. Visual inspection suggests that the first two categories were correctly assessed, as were all but 3 of the 'possibly fibrotic'. Importantly, there are no false negatives.

Fig. 6: images of slices with regions classified as ‘not normal’.

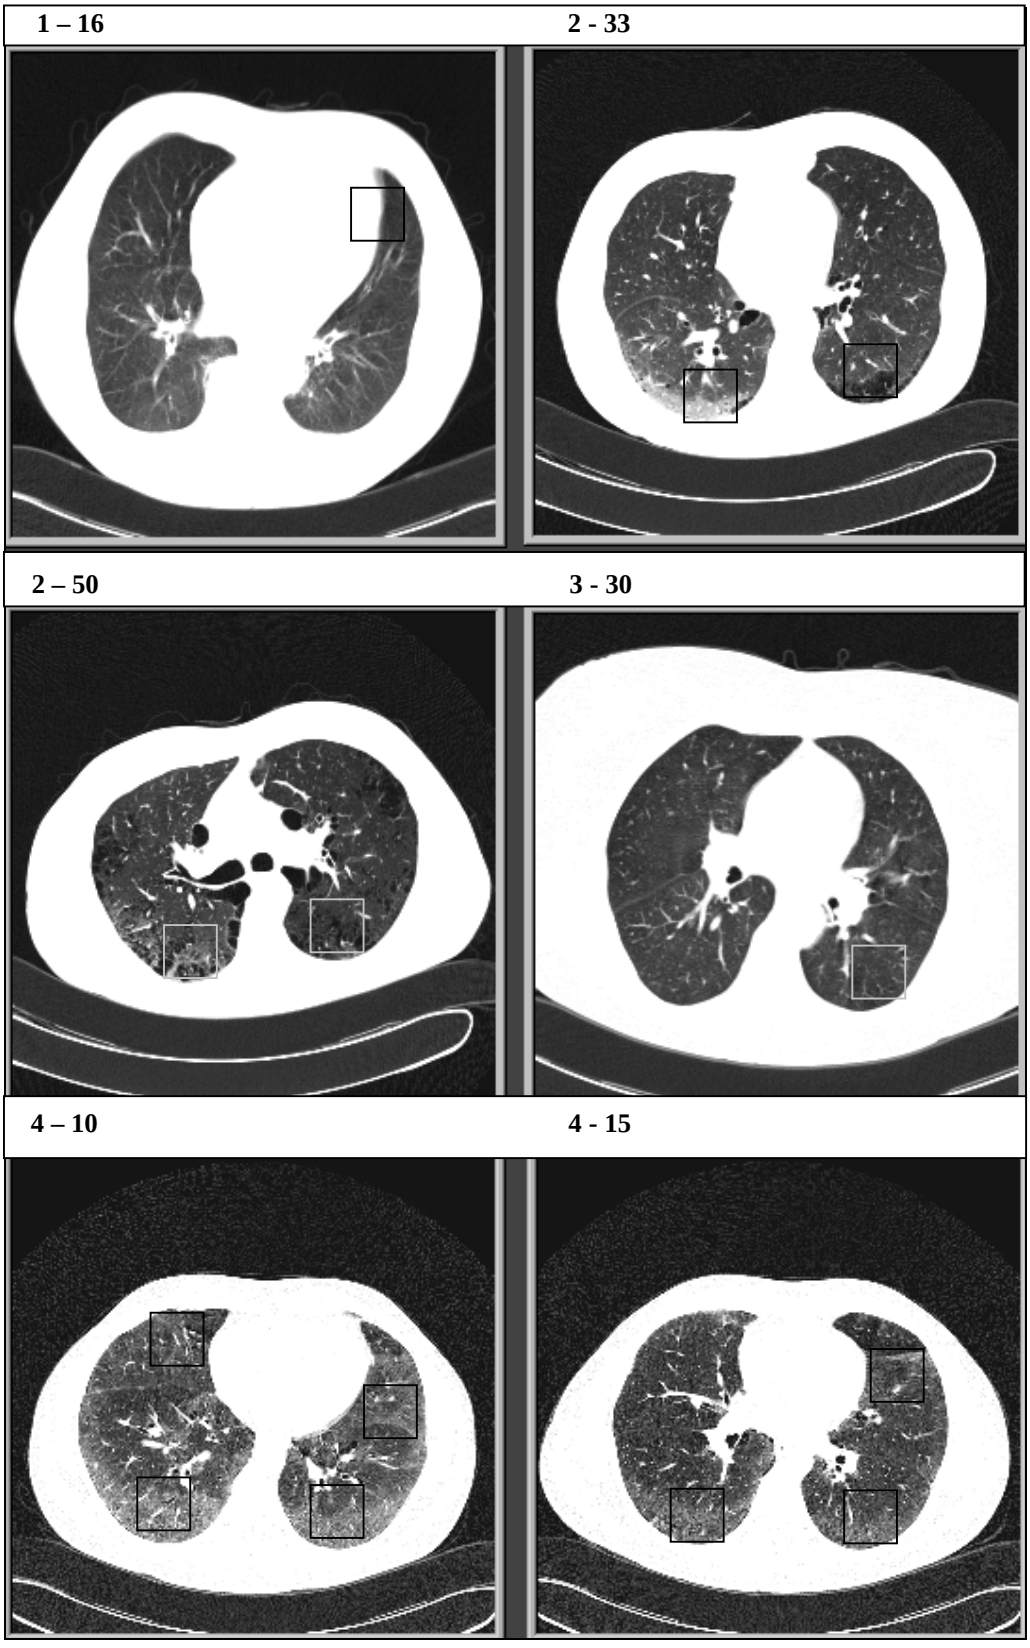

Fig. 6 ctd.

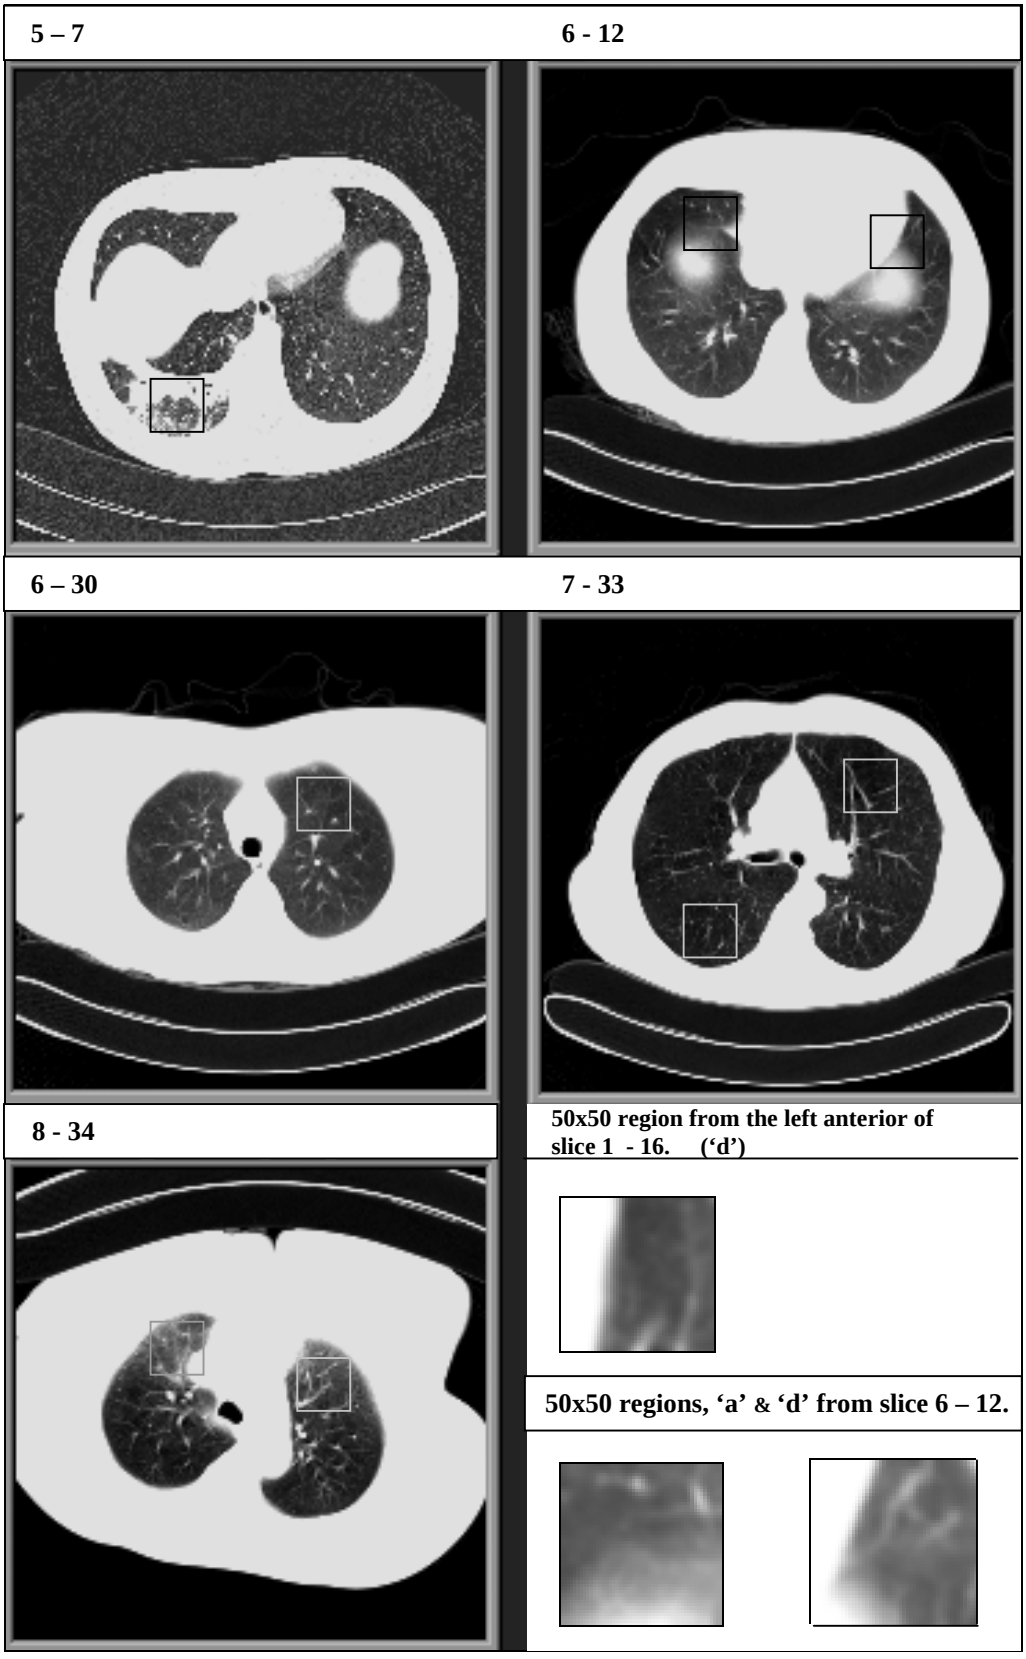
5 – 7
6 - 12
6 – 30
7 - 33
8 - 34
50x50 region from the left anterior of slice 1 - 16. ('d')
50x50 regions, 'a' & 'd' from slice 6 – 12.

DISCUSSION

Image 1 – 16, in figure 6, suggests a location comprising normal parenchyma. Its region ‘d’ referred to in table 4 is assessed as having a ‘normal’ information dimension, i.e. $\geq$1.85, though almost half the pixels have greyscale values above the upper threshold limit. The indication is that the extracted 50x50 anterior image contains a substantial area outside of the lung field; visual inspection confirms this. (Fig. 6)

Image 2 – 33 shows regions of normal, fibrous and emphysematous parenchyma; the two abnormal regions are posterior – fibrosis in the right lung, emphysema in the left. In each case, the information dimension is < 1.83 with a large number of pixels respectively above and below the ‘normal’ threshold limits. The slice from higher in the thorax, 2 – 50, indicates more widespread emphysematous bullae around the periphery of each lung. In region ‘b’ there are about the same number of pixels above and below the ‘normal’ threshold limits suggesting emphysema with scarring, while region ‘c’ has around half the distribution below the limit. Computer and visual analyses concur.

Slice 3 – 30 appears practically lesion-free, though region ‘c’ has a substantial number of pixels below the lower limit. This is also true of slice 48 (not illustrated), which suggests a lower-density feature extending over a considerable portion of the lungs’ posterior. The official clinical report mentions ‘a few tiny bullae peripherally’, a condition to which the program may be responding.

Data-set 4 is from the scan of a patient clinically diagnosed as having ‘appearances that indicate an active alveolitis with early signs of fibrosis’. Images of slices 10 and 15 show clearly the characteristic ‘ground glass’ opacification. Note, incidentally, the amount of noise in the image; no smoothing has been applied to ensure greater clarity. The figures quoted in table 4, however, were obtained using smoothed 50x50 regions.

In every example, the assessed information dimension for normal greyscale distribution suggests abnormality.

A brief inspection of the image of slice 5 – 7 is enough to suggest something seriously wrong at the posterior of the right lung. The low information dimension for the pattern of ‘normal’ pixels, together with the large number of high greyscale ones recorded, point straight to an area of major collapse. Distended bronchi suggest bronchiectasis.

Large numbers of higher than desirable greyscale value pixels are recorded for all regions tested in slices 6 – 12 and 6 – 30. The anterior regions seem most afflicted. The 50x50 images ‘a’ and ‘d’ from slice 12 do have some non-parenchymal portions, but the information dimension for ‘d’ is far from ‘normal’. The features in image 6 – 30 resemble those of slice 12 very closely, though here all values are derived only from an assessment of the parenchyma. The opacity in 6 – 12 ‘d’ could be an artefact due to movement, but that of the anterior regions of slice 30 suggests abnormality. The lungs appear to be hyper-perfused, possibly as a result of hypertension.

The abnormal readings for the images from the posterior of slice 7 – 33, right and left, are more difficult to account for. Visually, the lungs appear sound. It is possible that a small density loss has been detected, though for this study it will be suspected as a false positive diagnosis.

Data-set 8 was obtained from a patient scanned in a prone position. The image of slice 34 has been rotated for more convenient display. However, the threshold limits defining ‘normality’ have been applied appropriately. The anterior regions have unusually high opacity. It is known that there are metastatic deposits present in some

locations, but probably not enough to account for the apparently high densities; again, hyper-perfusion might be suspected.

No attempt has been made in this study to classify abnormality by quantitative means. A modal greyscale value in a normalized CT image of 'healthy' lung parenchyma is used as an index of normality for a similar location in any other data-set; the information dimension for the pixel pattern segmented by thresholding within limits set closely above and below the modal value provides a 'criterion value' – which is an empirically derived measure taken to signify that a particular segmented pixel distribution corresponds to a recognizable physiological condition in the patient. [11] By using only the criterion value for normality, a technician can draw attention to any abnormal regions by removing the apparently normal ones from one's area of concern.

Care is needed when assessing an image restricted to fewer 'normal' pixels than is optimal in a 50x50 format (c. 2000 might be expected); the cause may be an intrusion of a non-parenchymal feature into the image or, with possible pathological significance, the presence of a subset of 'abnormal' pixels. Where an extended patch of higher-than-normal tissue density is suspected from the image data, the diagnosis might be fibrosis. Similarly, a substantial number of associated points below the normal greyscale threshold may indicate a corresponding region of lower-than-normal tissue density – a strong indication of emphysematous change. Visual inspection of the suspect region is generally sufficient to make the diagnosis – especially where the feature in question extends across contiguous slices. However, a more automated diagnostic system should be able to discriminate an 'incomplete' image region from a 'diseased' one, and indicate whether any implied physical abnormality is of the emphysematous or fibrous kind. For this, at least two more threshold 'windows' will be needed – one to segment any pixel pattern corresponding to below-normal tissue density, the other to register any pathologically significant distribution of pixels with higher-than-normal greyscale values. 'Pathologically significant' implies a greyscale range corresponding to parenchyma with either higher than usual opacity or lower than expected density. The

appropriate criterion values (critical information dimensions of the segmented patterns) must be determined empirically and incorporated into a robust algorithm.

Results in this study, obtained using only the one, 'normal', window on normalized images, are encouraging since none of the regions flagged as 'normal', i.e. having an information dimension ≥ 1.85, over a near-full distribution of 'normal' pixels, is found, by visual inspection, to have any abnormality.

APPENDIX
Information dimension (diagnostic criterion)

The CT images used in this study are two-dimensional thoracic slices extracted from complete data-sets. Regions of interest, representing lung parenchyma, are selected in a 50x50 pixel format; the area of this sub-set of points, taking pixel size as unit measure, is, thus, 2500. When the diagnostic thresholds for 'normal' parenchyma are applied, the result is generally a segmented pixel pattern with an irregular distribution. For 'healthy' anterior regions, the greyscale values are normalized closely around mode value 35. The expected number of pixels is close to 2000, the remainder representing blood vessels, septa etc. The distribution will have a fractal dimension < 2 since the pixels do not uniformly fill the square. Because lung parenchyma has near-uniform density, with the available space for gaseous exchange efficiently filled with the branching system of airways and blood vessels, we might reasonably expect segmented image regions of 'normal' tissue to have a fractal dimension close to 1.90. From this, a criterion value of ≥ 1.85 is taken as a heuristic in this study, as an index of health. The implied assumption is that regular lung parenchyma has a near-uniform local density with a strong modality, and that x-ray attenuation values translate more or less linearly into normalized greyscale values.

An assigned fractal dimension is a means of characterizing a distribution of points embedded in a Euclidian space: how spread out or irregular the set is when examined at a particular scale. There are a number of different definitions of dimension, though most practical applications use a form of box-counting. To find a 'box' dimension of a plane set F, a mesh of squares of side 'r' is drawn to cover F and the number of boxes, $N(r)$, counted that contain at least one point of F. A count is made for various small values of 'r', and the dimension taken as the logarithmic rate at which $N(r)$ increases as 'r' $\rightarrow 0$. In practice, the largest value of $N(r)$ is where 'r' = pixel size; the dimension being estimated by the gradient of the graph of log $N(r)$ against –log r.

The information dimension is a version of box-counting in which more account is taken of the way points are distributed within a set. A measure of information gained by observing the occurrence of an event – such as the finding of points from a covered set in a 'box' of given size – is taken as minus the log of the probability of that event. Thus, the probability, 'P', of finding a point from a given set of 's' points in a typical cell of a minimal cover of N cells, linear measure 'r', is given by dividing the number of points found in the cell by the total number of points in the set. For each cell, linear measure 'r', there is a calculable 'surprise' (i.e. information), the average of which for the 'N' covering cells being the product of sampled probability and its surprise value, summed for all the cells. Now, the sum of the probabilities for all cells will equal 1, so average information, $I_{(r)}$, becomes the log of the number of cells, $N_{(r)}$, 'visited' by the notional average point distribution. (This is because average probability for each cell is $1/N(r)$.) We now have an analogue of basic 'box counting', where the count of boxes, scale 'r', all of average probability, just needed to cover the set, is the modified 'box' number that

takes more account of the non-uniformity of distribution. The 'information dimension' is now calculated using the general notion of dimension:

$$D_{(\text{information})} = \lim_{r \to 0} \frac{I_{(r)}}{\log (1/r)}$$

Where cell sizes are very small, the cell count, $N_{(r)}$, will represent saturation and have a value corresponding to the total number of thresholded pixels. At the lower end of the column, the cell sizes used are relatively large and will be too coarse-grained to assess a fractal pattern separately from its background. Between these extremes, the data from box counting are significant and useful for making computations. Rows of stars are used in the readout to indicate the limits.[15] By definition, fractals have the same granularity across scales; but, since images are composed of pixels, it does not take long before the limiting granularity, the individual pixel size, is reached. It follows that only a relatively small range of scales can be used for estimating the fractal dimensions. This is reflected in the small number of relevant measures appearing between the stars in the 'fractal dimension' readout. Because linear sizes of successive grid boxes are scaled by 2, and logs base 2 are used, the differences between successive entries in columns three and four, that is ($\log_2 N_{(r)} - \log_2 N_{(2r)}$ and ($I_{(r)} - I_{(2r)}$), give estimates respectively of capacity (box) and information dimensions – at the specified scale.

From:

(1).... $N_{(r)} = k.(1/r)^d$ and (2).... $N_{(2r)} = k.(1/2r)^d$, we get

(3).…. $\log_2 N_{(r)} = \log_2 k + d.\log_2 (1/r)$ and

(4)…. $\log_2 N_{(2r)} = \log_2 k + d.\log_2 (1/2r)$.

Now, by subtracting (4) from (3), we find:

$\text{Log}_2 N_{(r)} - \log_2 N_{(2r)} = d(\log_2 2r - \log_2 r)$

$= d$ (i.e. dimension).

The final estimate of information dimension is obtained by fitting a 'best' line to the plot of log r against $I_{(r)}$. In effect, this is providing the average slope of the plot, which is an estimate of the fractal dimension for that range of scales.

REFERENCES

1 Goddard P, 1982. Computed Tomography of the Lungs. Bristol, MD Dissertation, ref 'a', pp 77-81 & 134-139; ref 'b', p43

2 Giger M, Bae K, MacMahon H. Computerized Detection of Pulmonary Nodules in Computed Tomography Images. Invest Radiol, 1994; 29, pp 459-465.

3 Archer D, Coblentz C, deKemp R, Nahmias C, Norman G. Automated in vivo quantification of emphysema. Radiology 1993; 188, pp 835-838.

4 Sakai N, Mishima M, Nishimura K, Itoh H, Kuno K. An automated method to assess the distribution of low attenuation areas on chest CT scans in chronic pulmonary emphysema patients. Chest 1994; 106, pp 1319-1325.

5 Zagers R, Vrooman H, Aarts N, et al. Quantitative Analysis of Computed Tomography Scans of the Lungs for the Diagnosis of Pulmonary Emphysema. Investigative Radiology 1995; 30 (9), pp 552-562.

6 Uppaluri R, Mitsa T, Sonka M, Hoffman E, McLennan G. Quantification of Pulmonary Emphysema from Lung Computed Tomography Images. Am J Respir Crit Care Med 1997; 156, pp 248-254.

7 Uppaluri R, McLennan G, Sonka M, Hoffman E, 1998. Computer- based objective quantitative assessment of pulmonary parenchyma via x-ray CT. Http.//everest.radiology.uiowa.edu/spie/spie98/re1/SPIE98.html

8 Penn A, Loew M, 1998. Estimating Fractal Dimension of Medical Images with Fractal Interpolation Function Models. The George Washington University Institute for Medical Imaging and Image Analysis, Washington DC 20052. (www.seas.gwu.edu/~medimages/gwifm12.htm)

9 Heinonen T, Dastidar P, Kauppinen P, Malmivuo J, Eskola H. Semi-automatic tool for segmentation and volumetric analysis of medical images. Med Biol Eng Comput 1998; 36 (3), pp 291-296.

10 Uppaluri Renuka, Mitsa T, Hoffman E, McLennan G, Sonka M. Texture Analysis of Pulmonary Parenchyma in Normal and Emphysematous Lung. Medical Imaging, 1996; Spie 2709, pp 456-467.

11 Horwood A, 2000. Computer Diagnosis of Tomographic Pulmonary Images. Bristol, PhD Dissertation.

12 Kemerink G, Lamers R, Thelissen G, van Engelshoven J. Scanner conformity in CT densitometry of the lungs. Radiology Dec 1995; 197 (3), pp 749-752.

13 Kreithen D, Halversen S, Owirka G. Discriminating Targets from Clutter. The Lincoln Laboratory Journal 1993; 6 (1), pp 25-52.

14 Huda W, Scalzetti E, Levin G. Technique Factors and Image Quality as Functions of Patient Weight at Abdominal CT. Radiology Nov 2000; 217 (2), pp 430-435.

15 Sarraille J, DiFalco P. A Program for Calculating Fractal Dimension; FD3, version 0.3. (ftp://ishi.csustan.edu/pub.....e-mail:john@ishi.caustan.edu)