Background estimation in CT lung images with applications to perfusion visualisation and lobe separation

Pieter Bruyninckx[1,2], Dirk Loeckx[1], Frederik Maes[1], Dirk Vandermeulen[1], and Paul Suetens[1]

[1] Medical Image Computing, Department of Electrical Engineering (ESAT), K.U.Leuven, Belgium *
[2] Pieter.Bruyninckx@uz.kuleuven.ac.be

Abstract. This article proposes a method for estimating the local properties of the lung background or parenchyma, i.e. local mean and standard deviation of the intensities, from CT images. The properties estimating method uses an iterative estimation scheme with outlier detection and exclusion applied in a sliding window. The iterative optimisation method allows for the inclusion of prior knowledge about the background characteristics, thereby increasing the robustness.
Two applications for the algorithm are presented. The first application is the visualisation of the background pattern in the lungs. The second application is the separation of the lungs lobes by using a soft vessel segmentation. The estimated background properties allow the detection of the smaller vessels without the need for a fixed threshold. The background estimation method has been validated on a set of 50 phantom images with and without simulated vessels. The applications have been tested on 3 data sets from 3 different patients. The results have been evaluated by a radiologist. He considers the perfusion method helpful for reading the images and finds the lobe segmentation results acceptable.

1 Introduction

Radiologists have to review an increasing amount of CT images of the thorax, due to increased scanner speed and resolution, and an increasing demand. The evaluation of these images is most often performed by looking at 2D slices, using the standard lung window (from -1300 HU to 300 HU). In order to enhance the quality of the diagnosis and reduce the time spent looking at images, CAR/CAD methods are useful. This article presents two methods for improving the reading of lung images: visualisation of the lung perfusion and of the lung lobes. Both of these methods are based on the estimation of the properties of the lung background. The lung voxels can be subdivided in low intensity background voxels (parenchyma) and higher intensity foreground voxels (mainly vessels).

* This work was supported in part by the Research Fund K.U.Leuven under Grant GOA/04/05 and in part by I.W.T.-Vlaanderen under Grant IWT/SBO/60819-QUANTIVIAM.

The visualisation of the lung perfusion is mainly useful for the assessment of mosaic perfusion. Mosaic perfusion is a symptom common in the lungs, (mildly) present in about 80% of all healthy people. It manifests itself in a CT image by inhomogeneous background intensities. There are however multiple factors that can cause mosaic perfusion. [1] Diseases, e.g. pulmonary emphysema or small airways disease, may decrease the blood perfusion in certain areas of the lung, which causes a local decrease in intensity. Gravitation causes a slight increase of blood in the lower parts of the lung, which results in a different intensity pattern according to the patient's orientation (prone vs. supine). Other factors, e.g. the age of the patient and the phase of the respiratory cycle, globally affect the intensities. When the standard lung window is used to view the image it is not easy to perceive the inhomogeneous background pattern.

Herzog et al. have proposed a tool to visualize this background pattern in order to assess pulmonary embolism. [2]. They detect the background by using a global adaptive threshold that preserves the 80% of the lung voxels with the lowest intensities. A smoothed version of these background voxels is shown as a coloured overlay upon the original image.

We propose an improvement computing the properties of the background by using an iterative parameter estimation algorithm with outlier detection and exclusion, applied locally within a sliding window. The parameters under consideration are the mean and standard deviation of the intensities. Due to the iterative procedure this algorithm allows for the incorporation of prior knowledge concerning the lung background intensities.

As a first application of this algorithm the perfusion in the lungs can be visualised by showing the estimated mean as a coloured overlay.

We also propose to use the background knowledge to segment the lung lobes. A segmentation of the lung lobes is useful because it facilitates orienting oneself in the image and because lobe boundaries are often also a border to various pathologies.

Segmenting the lung lobes has traditionally been performed by segmenting the fissures themselves, like in [3]. This is also the method a physician applies when he looks for different lobes in the image. However the fissures are not always visible, or might even be absent [4]. The lobes can also be segmented using the vessel tree, as has been done by [5]. The authors first segmented the vessel tree, and looked for the surface that optimally separated the vessel trees. The exact location of this surface was refined using the fissure information if it was present in the image. Zhou et al. used a similar approach, but they also included the segmentation of the bronchi [6].

We present a lung lobe segmentation using the estimated background properties. First a soft vessel segmentation is performed, by detecting the vessels as outliers to the background. By applying the fuzzy distance transform (FDT) [7] starting from manually indicated points near the hilum the lung lobes are separated using their vessel trees.

The lobe segmentation method presented improves over the previously mentioned methods because it uses the estimated background properties to be robust

in case of mosaic perfusion. Hence it shows the advantage of using a suited background estimator.

The next section presents the methods. A validation of the background estimation method on phantom images and of the applications on real images is presented in Section 3. A discussion of the presented methods and hints for future work in Section 4 conclude this article.

2 Method

The algorithm consists of the following parts. The first step segments the lungs. Next the local background parameters, i.e. the local mean and the local standard deviation of the intensities, of the lung parenchyma are estimated. These parameters are used for two applications, i.e. visualising the lung perfusion and the separation of the lung lobes. All algorithms operate in three dimensions.

2.1 Lung Segmentation

In order to be able to focus on the relevant parts of the image a basic lung segmentation algorithm is applied. At first it selects all air in the image by applying a fixed threshold at -200 HU. The lungs are extracted using a connected component analysis of the whole image. The dominant connected component in the slice at one third from the top in the area $[\frac{1}{4}n_x, \frac{1}{2}n_x] \times [\frac{1}{4}n_y, \frac{3}{4}n_y]$ will denote the air in the lungs, with (n_x, n_y) the slice dimension. The smaller vessels are already included in this volume due to partial volume (PV) effects. To include also the larger vessels, a closing operator using a cuboid of size $10 \times 10 \times 10\text{mm}^3$ is applied. The speed of the lung segmentation is increased by applying it on a subsampled image. The image is subsampled by keeping the first of every two voxels in every dimension, reducing the total image size by a factor 1/8. If the only interest is visualising the background, the algorithm can be speed up by omitting the closing operation.

2.2 Estimating the Background Properties

The background of the lungs consists of the low intensity (around -900 HU) voxels that make up the parenchyma. The foreground consists of everything else with a higher intensity. The vessels make up most of these higher intensity structures. The background intensities are modeled by a normal distribution, characterised by its mean μ and standard deviation σ.

The background properties μ and σ are computed in a sliding window. For each voxel a surrounding region is selected. The diameter of this region should be chosen larger than the largest vessel inside the lung mask. For this application we have chosen the diameter to be 10 mm. Such a region typically contains only two tissues, i.e. the background and the vessels, so the histogram is expected to consist of two Gaussian peaks. For a healthy lung the background peak is centered at about -900 HU, while the vessel peak is centered at about 200 HU,

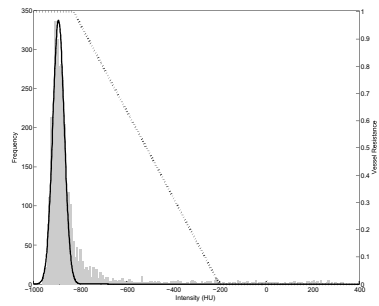

Fig. 1. Histogram of an area of the lung (bars). The full line is the Gaussian fitted to the background peak. The dotted line shows the vessel resistance function defined in (1).

assuming contrast images. However due to partial volume effects, the vessel peak can degenerate into a slope of the background peak. A typical histogram of such a region is shown in Figure 1.

To be able to separate this slope on the right from the background on the left we need to know the properties of the Gaussian at the left, which are in fact the afore mentioned μ and σ. This could be done using a maximum likelihood expectation maximisation (MLEM) that incorporates PV effects like [8].

A more efficient way to get this result is by calculating the mean and variance with outlier detection and exclusion, using an iterative algorithm. First μ and σ are initialised at their default values, i.e. -860 HU and 70 HU. These values are based on the measurement of lung intensities of healthy and pathological lungs. μ and σ are iteratively updated by computing the mean and standard deviation of the voxel intensities in the area that does not deviate more than $n\sigma$ from μ. The estimated standard deviation is biased because it is computed on voxels with a limited intensity range. To prevent a bias on the estimated standard deviation it should be multiplied by 1.13684 for n equal to two, as used in this article. The formula to compute this value is derived in the appendix. The iteration is repeated seven times. Empirical testing has shown that this suffices for convergence on clinical images of the thorax.

This algorithm is speed up by calculating μ and σ on a subgrid, and interpolating the result afterwards. The distance between the voxels included in the subgrid is half the size of the sliding window for each dimension. Since the algorithm acts as a low pass filter this approximation does not cause significant changes to the result. In a pathological image the average absolute differences were 3 HU and 4.3 HU per voxel for μ and σ respectively.

The perfusion is visualised by showing the computed mean μ as an overlay using a rainbow colour map onto the original image. This is shown in Figure 3.

2.3 Lobe Segmentation

The estimated background properties are used to calculate in each voxel a probability-like vessel resistance value f. This value is 1 for background voxels, and 0 for voxels that almost certainly belong to a vessel, while linearly interpolating in between. The following definition has been chosen to incorporate as much image information as possible while minimising the effects of the fissures present in the image.

$$f = \begin{cases} 1 & x \leqslant \mu + k\sigma \\ 1 - \frac{x-\mu-k\sigma}{-200-\mu-k\sigma} & x \in (\mu + k\sigma, -200) \\ 0 & x \geqslant -200 \end{cases} \quad (1)$$

In this equation x is the voxel intensity, while μ and σ are the values computed in section 2.2. The parameter k determines the distance between the centre of the background peak and the start of the slope. This parameter should be chosen to eliminate the background peak from the slope, so a value of $k = 3$ is a safe choice. A lower value might increase the accuracy while a value too low will severely distort the result. For all figures in this paper k is set to 2.3. The end of the slope is fixed at -200 HU, so all voxels with a higher intensity are expected to belong to a vessel, and will have zero vessel resistance. A plot of f is shown in dots on Figure 1.

The lobes are segmented using a FDT of (1) by comparing the distance from multiple starting locations in the image to each lung voxel. The starting locations are placed manually at the start of large vessel trees, near to the hilus.

3 Results

The method for computing the background properties has been validated using a software phantom. The perfusion method as well as the lobe separation have been tested on CT data sets of three different patients, and the results were evaluated by an experienced radiologist.

3.1 Tests on synthetic data

In order to test the accuracy of the background estimation method a set of random phantom images imitating some of the lung's characteristics has been created. The generated phantom images consist of a background with slowly varying intensity (-900 HU to -700 HU, left to right) containing a set of non overlapping vessels of varying size (radius 0.25 voxels to 4.75 voxels, top to bottom). The vessel model is similar to the one presented in [9]. The vessels are modeled as cylinders. The model incorporates partial volume (PV) effects, a Gaussian point spread function (PSF), and different noise characteristics for vessels and background. The PSF is an isotropic Gaussian with standard deviation 0.75 voxels. The foreground (vessel) voxels have a mean intensity of 200 HU and a standard deviation σ_1 of 35 HU, while the background has a standard deviation σ_2 of

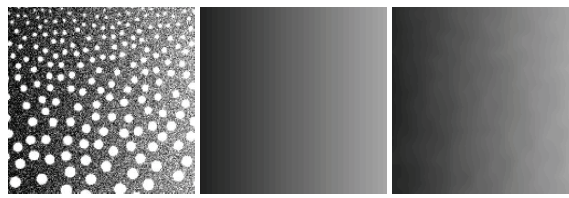

(a) (b) (c)

Fig. 2. Example of the calculation of the background properties on a phantom image. Phantom image (a), background ground truth (b), and estimated background (c). All images have a window level from −970 HU to −600 HU.

50 HU. The standard deviation of the PV voxels is $\sigma_\alpha = \sqrt{\alpha\sigma_1^2 + (1-\alpha)\sigma_2^2}$, where α is the vessel/voxel volume ratio for a certain voxel. A slice of one of the created phantom images can be seen in Figure 2(a).

The background estimation algorithm has been applied to a set of 50 phantom images. The difference between the estimated background intensity and the ground truth background was 7.2927 ± 2.5478 HU The difference between the estimated standard deviation of the background and the ground truth standard deviation was 3.9615 ± 1.5929 HU. The small positive bias of the mean is probably due PV effects. Changing the vessel intensity to 0 HU and 400 HU did not result in significantly different values. For comparison the same values for 50 background-only images were 0.0024 ± 0.7518 HU for the mean, and 0.1799 ± 0.7301 HU for the standard deviation. The standard deviations for all these values over all images was smaller than 0.3 HU.

3.2 Background Perfusion

We have applied the perfusion tool to three datasets of three different patients and presented the results to an experienced radiologist. He said the perfusion tool allows to visualise subtle diffusion variations, like those caused by gravitational effects. Without this tool a tedious amount of twiddling with the window level would be needed for these subtleties to be visible, so therefore he thought the tool would be able to increase his productivity. Since the algorithm is fully automated, it can be executed beforehand, thus enabling the radiologist to show the perfusion by using only one slider.

An example of the result of this algorithm is shown in Figure 3. The image on the left shows the original slice using the standard lung window. The image at the right shows the perfusion using a rainbow color map. The image in the center shows the perfusion map as an overlay onto the original image. In the original image a brighter background is visible when one has a good look at the bottom of the lung on the right hand side. In both other pictures this area is visible at first sight without effort. Moreover the right image shows a soft green

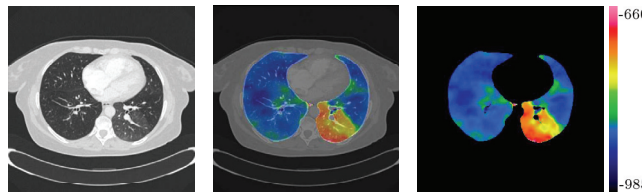

(a) original (b) original with overlay (c) overlay only

Fig. 3. Demonstration of the background properties. The coloured overlay shows the intensity of the parenchyma, thus allowing to visualise subtle diffusion variations.

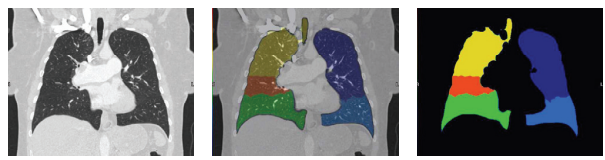

(a) original (b) original with overlay (c) overlay only

Fig. 4. Demonstration of the lobe segmentation. Different colours indicate different lung lobes.

glowing at the bottom of the lung on the left hand side, indicating a slightly higher perfusion due to gravitational effects.

The algorithm has also been executed with less optimal initial estimates for μ. The initial estimate does not influence the values to wich the iterations converge, since a difference by 100 HU yielded a mean absolute difference of less than 1 HU after fourteen iterations.

3.3 Lobe Segmentation

We have applied the lobe segmentation to three datasets of three different patients and presented the results to an experienced radiologist. According to him the lobe segmentation tool is adequate for clinical use. The most important advantage of this algorithm is that it is complementary to lobe segmentation by fissure detection, which is how the radiologist normally recognizes different lobes.

The algorithm does not always return the optimal result after the first manual initialisation. In this case an iterative procedure, in which additional starting points are added, is performed.

An example of the results returned by this algorithm can be seen in Figure 4. Figure (a) shows the original CT slice, and (c) shows a colour map depicting the different lung lobes. A combination of both is shown in (b). The lobe segmentation has been performed for $k = \{2.0, 2.3, 3.0\}$. The maximum volume

percent difference between the several results was smaller than 0.75%. These small differences between the different results show that the exact value of k is not critical.

3.4 Timing

The algorithms have been implemented in Matlab R2006b with C++ calls for the most time intensive parts. The Matlab image processing toolbox is used for performing the morphological operations of the lung segmentation step. On a Dell Precision workstation with dual Xeon 2.2 GHz/512k processors the lung segmentation took 37 minutes, estimating the background properties took 51 seconds, and separating the lobes took 270 seconds, for an image of $512 \times 512 \times 601$ voxels.

If one is interested in the perfusion only, and not in the separation of the lung lobes, the efficiency of the lung segmentation step can be improved by omitting the time consuming morphological operations, which account for 36 minutes of processing time.

4 Discussion and Future Work

We have proposed a method for computing the properties of the lung background using a low pass filter with outlier detection and exclusion. The result of this filter can be visualised immediately, showing the lung perfusion. The accuracy of the results allows the physician to see otherwise nearly invisible variances in the background distribution, like those caused by gravitational effects. Showing the result as an overlay results in a simple tool, which was judged by a clinician as very helpful, e.g. because it can speed up the reading. By using the aforementioned background properties we have also proposed a lobe segmentation algorithm, which is robust in cases of mosaic perfusion.

The algorithm has not yet been tried on non-contrast images, but the intensity gap between foreground and background should still be sufficient to yield acceptable results.

To improve the visualisation of the perfusion an automatic detection of the adequate window level would be useful. A fixed window level is not adequate, since the amount of air in the lungs affects the background intensities. The initialisation of the lobe segmentation method could be automated by using a map like in [3]. Once the lobe segmentation has been performed, the perfusion detection can be recomputed on a per lobe basis, increasing the visibility of the contrast differences at the fissures.

Our focus for the future will rather be on using the background estimation algorithm in applications, e.g. vessel segmentation, than on improving the algorithm as such. The methodology for segmenting the lung lobes can also be used to separate the arteries and the veins, yielding promising preliminary results.

A Bias on Estimated Standard Deviation

Let $f(x)$ be a Gaussian probability density function with mean μ and standard deviation σ.

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \tag{2}$$

Suppose the mean and the standard deviation of this probability density function are estimated on a limited interval, i.e. in the interval $[\mu - n\sigma, \mu + n\sigma]$, with $n > 0$. It is obvious that the estimate mean μ' will be the same as the real mean μ. The estimated standard deviation σ' however will be biased.

The estimated standard deviation σ' is computed as

$$\sigma' = \sqrt{\frac{\int_{\mu-n\sigma}^{\mu+n\sigma}(x-\mu)^2 f(x)dx}{\int_{\mu-n\sigma}^{\mu+n\sigma} f(x)dx}} \quad . \tag{3}$$

From (3), $\frac{\sigma}{\sigma'}$ can be calculated.

$$\frac{\sigma}{\sigma'} = \left(-\frac{\left(\sqrt{2}n - \operatorname{erf}\left(\frac{\sqrt{2}}{2}n\right)\sqrt{\pi}e^{\frac{n^2}{2}}\right)e^{-\frac{n^2}{2}}}{\sqrt{\pi}\operatorname{erf}\left(\frac{\sqrt{2}}{2}n\right)}\right)^{-\frac{1}{2}} \quad . \tag{4}$$

References

1. Webb, W.R., Müller, M.L., Naidich, D.P.: High-Resolution CT of the Lung. 3 edn. Lippincott Williams & Wilkins (2001)
2. Herzog, P., Wildberger, J.E., Niethammer, M., Schaller, S., Schoepf, U.J.: CT perfusion imaging of the lung in pulmonary embolism. Acad Radiol **10**(10) (October 2003) 1132–1146
3. Zhang, L., Hoffman, E.A., Reinhardt, J.M.: Atlas-driven lung lobe segmentation in volumetric X-ray CT images. IEEE Trans Med Imaging **25**(1) (January 2006) 1–16
4. Dandy, W.E.: Incomplete pulmonary interlobar fissure sign. Radiology **128**(1) (July 1978) 21–25
5. Kuhnigk, J.M., Hahn, H., Hindennach, M., Dicken, V., Krass, S., Peitgen, H.O.: Lung lobe segmentation by anatomy-guided 3D watershed transform. Volume 5032., SPIE (2003) 1482–1490
6. Zhou, X., Hayashi, T., Hara, T., Fujita, H., Yokoyama, R., Kiryu, T., Hoshi, H.: Automatic segmentation and recognition of anatomical lung structures from high-resolution chest CT images. Computerized Medical Imaging and Graphics **30**(5) (July 2006) 299–313
7. Saha, P.K., Wehrli, F.W., Gomberg, B.R.: Fuzzy distance transform: Theory, algorithms, and applications. Computer Vision and Image Understanding **86**(3) (June 2002) 171–190
8. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: A unifying framework for partial volume segmentation of brain MR images. IEEE Trans Med Imaging **22**(1) (January 2003) 105–119

9. Wörz, S., Rohr, K.: Segmentation and quantification of human vessels using a 3-D cylindrical intensity model. IEEE Transactions On Image Processing **16**(8) (August 2007) 1994–2004