A Statistical Approach Towards Ribcage Segmentation (Summary of PhD Work in 2012)

Áron HORVÁTH
Advisor: Gábor HORVÁTH

I. Introduction

Ribcage segmentation is an essential step for bone shadow removal on chest radiographs. Ribs are surrounding the thorax, therefore most ribs have two shadows on an image. In this paper I will focus on finding the posterior part of the ribs, but some techniques can be applied to the anterior case as well.

There are several approaches towards ribcage segmentation. One of these is to analyze vertical intensity profiles of the images. This method was introduced by Souza [1], who synchronized these profiles by shifting. Then summed them into a single one-dimensional profile, on which he applied a t-test to enhance the rib borders. After local minima and maxima selection he got approximal rib borders. Here, I would like to present a new procedure which can be regarded as a revision of Souza's work. I will focus on the part which is related to Souza's work, the overview of my full solution can be found in [2]

II. Description of the method

One problem with the aforementioned method is that simple shift can not put these profiles into exact synch, because each rib has a slightly different slope. To overcome this I fit a smooth slopefield to the image gradients and synchronize the vertical sections based on it. The base section of the synchronization will be the yellow dotted curve of fig. 1d, and the intensities are averaged over the red curves.

The averaged intensity profile is depicted by the blue curve of fig. 2a. Instead of t-tests I applied a box filter (fig 2b), differentiation (fig 2c) and normalization, which encompasses a highpass filter and sliding RMS compensation (fig. 2d).

A simple peak-picking routine is not sufficient to locate the rib edges, because usually there are more extrema than rib borders, only images of young and healthy individuals lead to clear-cut cases.

To address this problem a dynamic programming approach is introduced. The mechanics of the search are as follows. The algorithm passes through the signal from the left to the right. Some points on the left side are possible upper borders. If intercostal spaces and rib thicknesses are supposed to change gradually, picking one point as upper border determines the possible position range for the corresponding lower border. Choosing a lower counterpart determines the permitted range for the next rib thickness. The third step, where the upper border of the second rib is selected, constrains the possible spacing range for the following ribs, and it goes like this. These ranges were defined by analyzing a set of manually segmented ribcages.

4 different statistics were created (fig. 3). All of them describe a dependence on the vertical position inside the lungfield: rib thickness, spacing between ribs, consecutive rib thickness ratio, consecutive spacing ratio.

The vertical position is measured from the apex of the lungfield along the midcurve (yellow dotted curve in fig 1d). This curve is not always perpendicular to the ribs, thus the distance measured on it between two rib borders can be a slightly more than the thickness of the rib. This is compensated by measuring distances based on the slopefield, accounting only the part which is perpendicular to it. The one dimensional signal is created utilizing this measure, too.

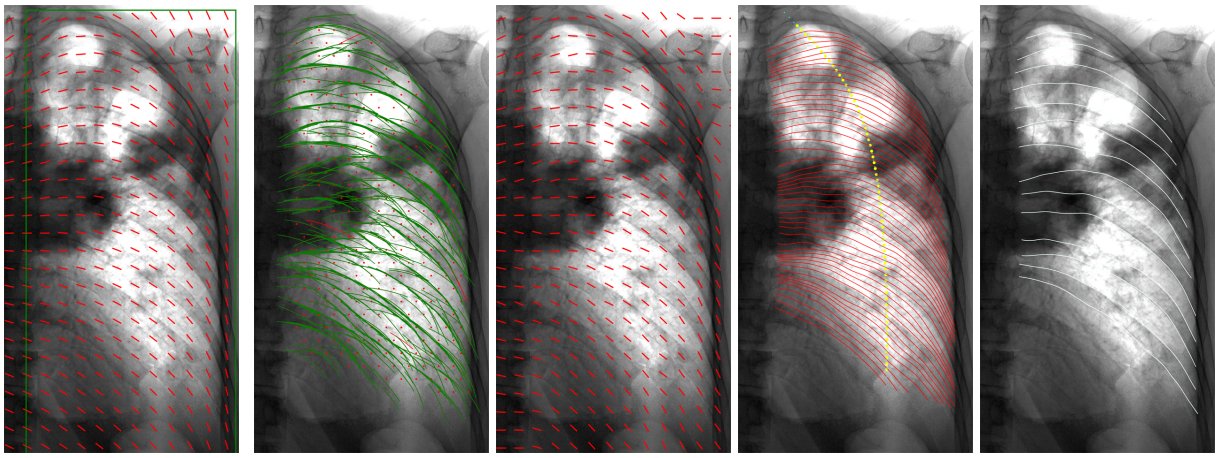

(a) Average slope field (b) Fitted arcs (c) Fitted slope field (d) Possible curves (e) Selected curves

Figure 1: Overview of posteriar ribcage segmentation. Arcs are fitted to the image to finetune an average slopefield, then some of the curves generated from this slopefield are selected.

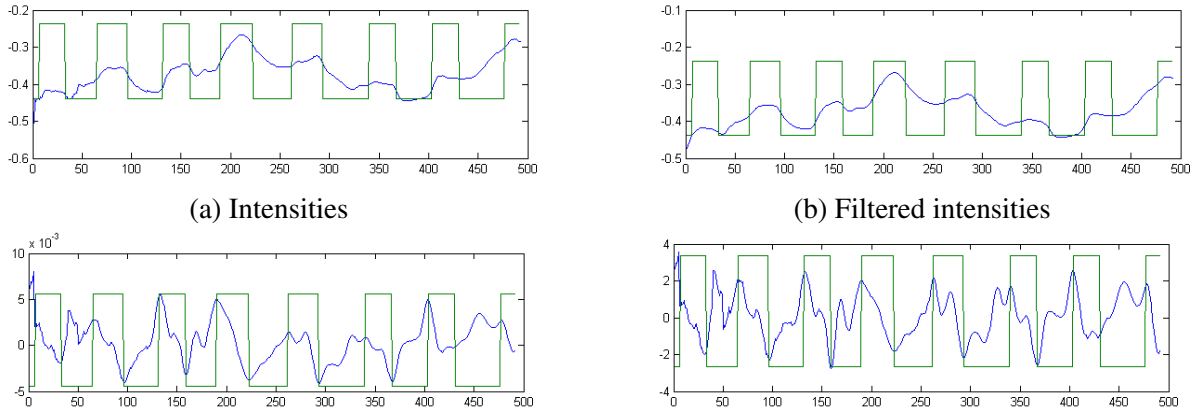

(a) Intensities (b) Filtered intensities

(c) Differences (d) Normalized differences

Figure 2: Processing the intensity signal. The output of the segmentation is indicated by green.

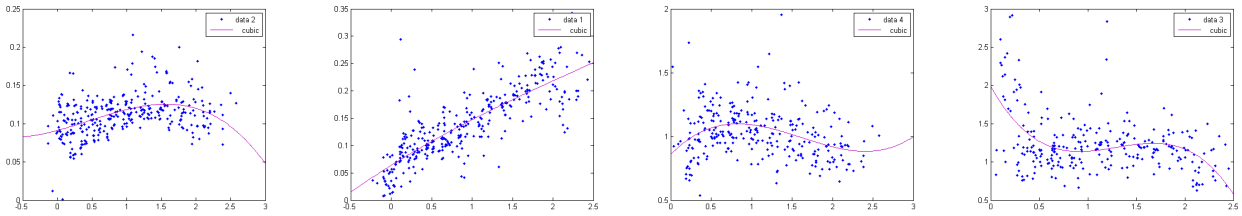

(a) Thickness (b) Spacing (c) Consecutive thickness (d) Consecutive spacing

Figure 3: Statistics collected from manual ribcage segmentations

As can been seen in figure 3 the thickness and the spacing are highly dependent on this vertical position, thus the permitted range for these quantities are gradually changing during the search. This graduality is assured by a cubic polynomial which is denoted by a purple curve in figure 3. Calculating the standard deviation from this polynomial gives the width of the permitted range. Values of $1\sigma < w < 2\sigma$ were proved to work well.

These constraints are incorporated into a dynamic programming problem shown in figure 4. Simplifying the problem by focusing only on one of the borders of a rib we can ignore the thickness. The vertical axis in fig. 4 represents the vertical position, the horizontal corresponds to the spacing. Thus each point on the plane defines a vertical position and a spacing value. The white region is permitted, the purple is forbidden. The possible starting and terminal points are located in the green and yellow area respectively. The aim is to find the least cost path between the green and the yellow region.

The cost for stepping on a vertical position comes from the differentiated intensities. The cost of stepping on a point with a given spacing comes from the statistics.

The cyan dots represent a path on this plane. A possible step can be calculated by moving downwards by the value of the current spacing (marked by an arc on fig. 4), and then choosing a point from a line segment crossing this point and having a $-45°$ slope. The endpoints of this line segment are determined by the consecutive spacing statistics. Each point on this line can have a different cost based on this statistics as well.

Considering both borders of the ribs we get an extra dimension, and this rectangle on the plane becomes a volume in the 3 dimensional space. In this volume the steps are alternating between the two planes. Stepping from an upper border to a lower border will take place on the $verticalposition-thickness$ plane, stepping from a lower border to an upper border will be on the $vericalposition - spacing$ plane.

The cost of a step consists of 3 main components:

$$c = c_i + c_p + c_d \tag{1}$$

The image cost c_i is the current value at the given vertical position on the differentiated intensity signal. The density cost c_d controls the number of ribs to find. The probability cost c_p has 4 components: thickness, spacing, consecutive spacing ratio and consecutive thickness ratio.

Approximating the density functions of these quantities by a position dependent normal distribution, we can get a probability value for a given thickness, spacing or consecutive spacing at each vertical position:

$$p(x, y) = \frac{1}{\sigma}\phi\left(\frac{y - f(x)}{\sigma}\right) \tag{2}$$

Where p denotes the probability in question, x the vertical position, y the dependent quantity, f the cubic polynomial and σ the deviation.

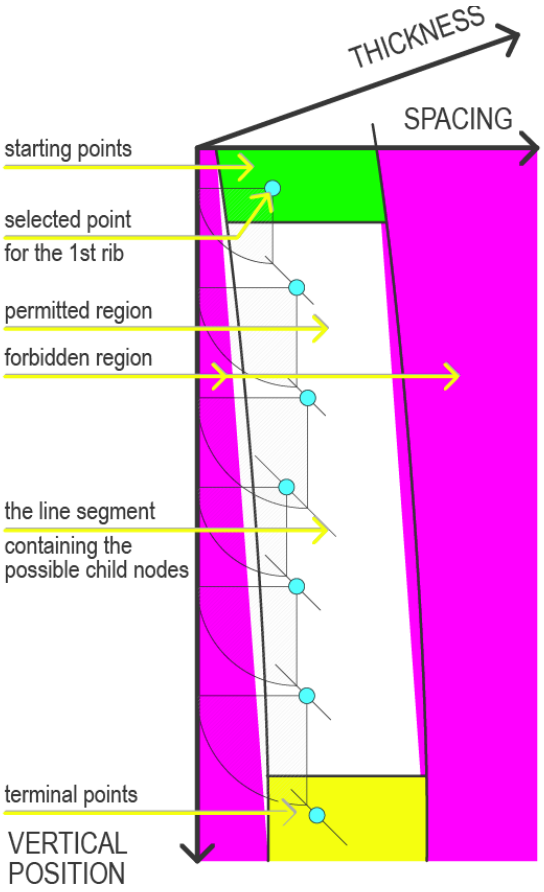

Figure 4: Search space of the dynamic programming problem

Table 1: Results. Found ribs per lungfields.

	true positive	false positive	inaccurate
previous method	7.0	0.28	0.15
proposed method without probability cost (1.2σ)	7.5	0.35	0.05
proposed method with probability cost (1.9σ)	7.2	0.06	0.03

As the dynamic programming accumulates the costs additively, the negative logarithm of the probabilities have to be used as cost. This way the mean squares of the deviations will be cumulated at the end of the path:

$$\sum_{i=1}^{N} c_p^i = -ln \prod_{i=1}^{N} p(x^i, y^i) = -ln \prod_{i=1}^{N} \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{(y^i - f(x^i))^2}{2\sigma^2}} = Nln(\sigma\sqrt{2\pi}) + \sum_{i=1}^{N} \frac{(y^i - f(x^i))^2}{2\sigma^2} \quad (3)$$

N is the number of ribs found, x^i is the position of the i-th rib, y^i can either be the thickness, spacing, consecutive spacing ratio of the i-th rib. The constant term can be grouped to the density cost c_d, as it depends only on the number of found ribs. The 2 in the denominator can also be left out:

$$\frac{\sum_{i=1}^{N}(y^i - f(x^i))^2}{\sigma^2} \quad (4)$$

σ^{-2} becomes the weighting factor between the different probability costs.

This probability cost has equal strength over the full signal, but the signal can be weaker at some places, that is why the aforementioned local RMS compensation was needed.

Unfortunately the cost of a path will depend on the number of nodes as well. If each cost is negative the minimum cost path search tries to break the path into more nodes to reduce the final cost. If each cost is positive the effect is the opposite. A balance have to be made by introducing a density coefficient c_d which will be added to each node's cost.

To keep down the computation time the search volume has to be coarsely quantized. The vertical resolution is set to around 600 depending on the size and shape of the lung. The resolution of the thickness and spacing axes are further reduced by a factor of two and four respectively.

III. Results

The method was evaluated on a set of 39 images independent from the ones used for creating statistics. The results are compared to a previous version where the permitted regions were set manually, and they changed linearly. This version lacked of signal normalization and probability costs as well. Introducing probability costs and statistics based permitted ranges significantly reduced false positives and raised true positives compared to the previous method, however turning off probability cost, and using tighter permitted range yields even more true positives on the expense of more false positives (see table 1). It must be noted that wider permitted range heavily slows down the procedure. The main advantage of this algorithm that its output does not requires further revision steps due to the low false positive level.

References

[1] P. de Souza, "Automatic rib detection in chest radiographs," *Computer vision, graphics, and image processing*, 23(2):129–161, 1983.

[2] S. Juhász, Á. Horváth, L. Nikházy, G. Horváth, and Á. Horváth, "Segmentation of anatomical structures on chest radiographs," *XII Mediterranean Conference on Medical and Biological Engineering and Computing 2010*, pp. 359–362, 2010.