

Minimum spanning trees and active contours for identification of the pectoral muscle in screening mammograms

Mariusz Bajger
Flinders University
School of Informatics and Engineering
PO Box 2100, Adelaide SA 5001, Australia
mbajger@infoeng.flinders.edu.au

Fei Ma
Flinders University
School of Informatics and Engineering
PO Box 2100, Adelaide SA 5001, Australia
ma0029@infoeng.flinders.edu.au

Murk J. Bottema
Flinders University
School of Informatics and Engineering
PO Box 2100, Adelaide SA 5001, Australia
murkb@infoeng.flinders.edu.au

Abstract

Image segmentation based on minimum spanning trees (MST) is used to identify the pectoral muscle in screening mammograms. The segmentation found using the MST is used to initialise an active contour for finding an anatomically reasonable estimate of the boundary of the pectoral muscle. The error is reported in terms of the number of incorrectly assigned pixels. Out of 83 images, 25 images have error rates less than 5 percent and 56 images have error rates less than 10 percent. The nature of the errors encountered indicates that the accuracy of computer algorithms for this task is approaching its practical limit.

1. Introduction

Several different methods have been explored for using graph theory ideas to improve image segmentation [1], [5], [6], [8], [3]. Methods based on graph theory are appealing because of their potential to incorporate global image information into the segmentation process. These methods work best if the regions forming the intended segmentation are well defined in terms of intensity or other derived image property. In many images, physical objects appear as regions in the image that are not uniform with respect to a well defined property. In x-ray images, for example, objects of interest may be of non-uniform thickness or density or may be partially obscured by other objects. In such cases, the object of interest may be viewed as comprising a number of regions each of which is relatively uniform with

respect to intensity.

In this study, a graph theoretical method based on minimum spanning trees (MST) [1] is used to recognise the pectoral muscle in screening mammograms. Finding the pectoral muscle is important for automatic image registration and because some automatic processes for detection of anomalies associated with breast cancer are designed to behave differently in the region of the pectoral muscle than in the breast region.

The densest part of the pectoral muscle is often the brightest object in medio-lateral oblique view mammograms. However, in many cases, the muscle varies in thickness and, in its entirety, is not restricted to an exclusive range of intensity values (Figure 1). Segmentation based on MST alone results in anatomically unrealistic representation of the pectoral muscle in many mammograms. By incorporating geometric information, the segmentation is adjusted to arrive at a largely correct delineation of the pectoral muscle. The remaining error lies in the detail of the boundary of the region identified. The detail is subsequently corrected by using the initial segmentation to initialise an active contour.

The MST algorithm is explained in Section 2 and the active contour adapted for this study is presented in Section 3. The data are described in Section 4 followed by results in Section 5 and discussion in Section 6.

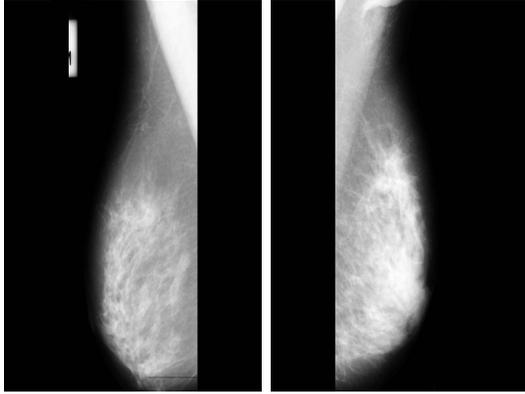


Figure 1. The pectoral muscle is the large, bright triangular region in the upper right corner of the left image (image mdb099) and the upper left corner of the right image (image mdb090). Although the pectoral muscle is clear in both these cases, the intensity of the pectoral muscle in the right image varies greatly and much of the muscle is not as bright as large parts of the breast region.

2. Image segmentation using MST

2.1 The MST algorithm

The MST process starts with a graph, $G = (V, E)$, where the set of vertices, V , is the set of pixels in the image to be segmented and the set of edges is

$$E = \{(v_i, v_j) : \|v_i - v_j\| \leq d\} \quad (1)$$

where $\|\cdot\|$ denotes a suitable norm. The set E determines the level of connectivity to be considered. If $d = 1$ and the norm is $\|w\|_1 = \max_k |w^k|$, where w^k indicates the k -th component of w , then every vertex is connected to its immediate eight neighbours. With this norm and $d = 2$, every vertex is connected to its immediate 24 neighbours.

The method proceeds by constructing a new graph, $H = (V, F)$, with $F \subset E$ in such a way that the (connected) components of H correspond to the desired segmentation of the image. Initially, $F = \emptyset$. This means there are no edges and so each vertex in V (each pixel in the image) is a separate component. In subsequent iterations, the elements from E are considered one at a time for inclusion in F .

The criteria for inclusion depend on weights assigned to edges according to the weight function

$$w((v_i, v_j)) = \begin{cases} |I(v_i) - I(v_j)|, & (v_i, v_j) \in E \\ \infty, & \text{otherwise} \end{cases} \quad (2)$$

where $I(v_i)$ represents the image intensity (or other image property) at the vertex v_i .

The steps of the algorithm are as follows [1]

1. Sort elements of E by edge weight to form a sequence e_1, e_2, \dots, e_m with $w(e_i) \leq w(e_j)$ for $i \leq j$.
2. Set $H^0 = (V, F^0)$ where $F^0 = \emptyset$.
3. Repeat step 4 for $q = 1, \dots, m$, where $m = |E|$.
4. Construct $H^q = (V, F^q)$ given $H^{q-1} = (V, F^{q-1})$ as follows. Let v_i and v_j denote vertices connected by the q -th edge in the ordering, i.e. $e_q = (v_i, v_j)$. If v_i and v_j are in disjoint components of H^{q-1} and $w(e_q)$ is small compared to the internal variation of both components, set $F^q = F^{q-1} \cup \{e_q\}$. Otherwise, set $F^q = F^{q-1}$.
5. Return $H = H^m$.

To understand these steps, formal definitions are needed for components, as well as internal and external variation.

The graph $C = (V', F')$ is called a connected component of $H = (V, F)$ if $V' \subset V$, $F' \subset F$, and for all $v_i, v_j \in V'$, there exists a sequence of vertices v_k , $k = 1, 2, \dots, n$ such that $v_1 = v_i$, $v_n = v_j$, and $(v_k, v_{k-1}) \in F'$ for $k = 2, 3, \dots, n$. In other words, its possible to visit all the vertices in V' by following edges in F' .

The *internal variation* of a component $C \subset H$ is defined as the maximum weight edge in any minimum spanning tree of C . Hence,

$$Int(C) = \max_{e \in MST(C)} w(e), \quad (3)$$

where $MST(C)$ stands for a minimum spanning tree of C . The *minimum internal variation* for components C_1 and C_2 is defined by

$$MInt(C_1, C_2) = \min(Int(C_1) + \tau(C_1), Int(C_2) + \tau(C_2)), \quad (4)$$

where the threshold function τ is given by

$$\tau(C) = \frac{k}{|V'|}. \quad (5)$$

Here k is a positive integer, and $|V'|$ is the number of vertices in V' , where $C = (V', F')$.

Finally, the *external variation* between components $C_1 = (V'_1, F'_1)$ and $C_2 = (V'_2, F'_2)$ is defined as

$$Ext(C_1, C_2) = \min_{v_i \in V'_1, v_j \in V'_2} w((v_i, v_j)). \quad (6)$$

Thus in step 4, edge $e_q = (v_i, v_j)$ is included in F^q if and only if $v_i \in V'_1, v_j \in V'_2, V'_1 \cap V'_2 = \emptyset$, and

$$Ext(C_1, C_2) \leq MInt(C_1, C_2). \quad (7)$$

The constant k in the definition of τ controls the degree of similarity between the components and hence the number of components (and average size) found.

2.2 Preprocessing and cropping

Before applying the MST algorithm, the images were automatically cropped to a rectangle just large enough to contain the general region of the pectoral muscle. Cropping is expedient because the MST algorithm is computationally intensive and because removing competing bright components from consideration simplifies the processing.

In the rest of the paper, processing step will be described for left breast images, meaning that the pectoral muscle appears in upper left corner of the image.

For the data used in the study (Section 4), images did not appear flush with the left side of the film edge. To find the left edge of the image (and hence the left edge of the region containing the pectoral muscle), a simple threshold was set at 100 (image intensity range 0 to 255). To find the right edge of the region containing the pectoral muscle, the top row of the image was inspected for intensity changes consistent with the right edge of the pectoral muscle. The top row was used because the pectoral muscle is most distinct at the top of the mammogram and because the pectoral muscle is oriented so that the top row provides the right most extent of the pectoral muscle in the image. In the top row, an interval of length at least two was considered consistent if the intensity difference between consecutive pixels was less than two (allowing ramps and noise, but not large jumps). The intensity of a consistent interval was taken to be the intensity the left most member of the interval. The left most consistent interval with intensity less than a threshold was taken to signal the right edge of the cropped regions containing the pectoral muscle. For a given threshold, there may or may not exist a consistent interval with satisfying this condition. Accordingly, increasing values of the threshold were tested starting from 120 until a value was found for which such jumps existed. The lowest such threshold was used. The vertical extent of the cropped region was determined analogously. This simple scheme was designed to take advantage of the fact that the pectoral muscle boundary is distinct near the top of MLO view mammograms.

The cropped image was subsampled ($4 \times 4 \rightarrow 1$) and smoothed using a one-dimensional Gaussian filter with $\sigma = 0.5$ and oriented at an angle $\theta = \frac{-\pi}{4}$ to match, roughly, the expected angle of the pectoral muscle boundary.

2.3 Implementation of MST

Implementation of the MST algorithm requires setting E in equation (1) to establish the level of connectivity and choosing the parameter k in equation (5). In the experi-

ments described below, the radius in the definition of E was set to $d = 2$, and the norm was the maximum norm, $\|\cdot\|_{max}$. Thus every vertex was connected to its immediate 24 neighbours. The value of k was set to 300. Values for these parameters were found empirically.

The algorithm described above segmented the image into a number of components. The largest component that overlapped the 10×10 square flush with the top of the image but centered left to right was used to determine the pectoral muscle boundary. In most cases, the resulting component corresponded, roughly, to the whole pectoral muscle area. In a few cases, the selected component covered 80 percent or more of the cropped image. These cases were detected automatically and were addressed by running the segmentation algorithm a second time with $d = 1$ instead of $d = 2$ in the definition of E . This naturally results in a smaller number of components merging during the segmentation process and consequently produces a finer segmentation. In some cases, the resulting component still exceeded 80 percent of the cropped image. These cases (about 5 percent of the images) were also detected automatically and were treated as described below.

The errors leading to overly large pectoral muscle components stem from the fact that the lower part of the pectoral muscle is often poorly defined in the image. If other dense tissue appears near the lower part of the pectoral muscle, then the two components may fuse (Figure 2). In these cases, the preliminary components showed characteristic “fjords” corresponding to the portions of the pectoral muscle boundary where the boundary was strong (Figure 2). The following steps were used to separate the pectoral muscle portion of this component from the breast region. Starting at the top of the image, the first row intersecting the preliminary component was searched left to right for the first “break point”, meaning the first pixel within the component such that the next two pixels fell outside the component. This pixel corresponded either to the right edge of the component or to the left edge of the fjord. In subsequent rows, the break point was defined as in the first row, except that it was restricted to be no farther to the right than the break point of the previous row. In this way, an anatomically reasonable extension of the pectoral muscle is found in places where the boundary is weak (Figure 2).

3 Active contours

The MST method provided estimates of the pectoral muscles that corresponded well to the true pectoral muscle in terms of location and general shape. However, boundaries were ragged and generally appeared somewhat to the left of the visually apparent boundary (Figure 3). To improve the detail of the segmentation, the boundaries found

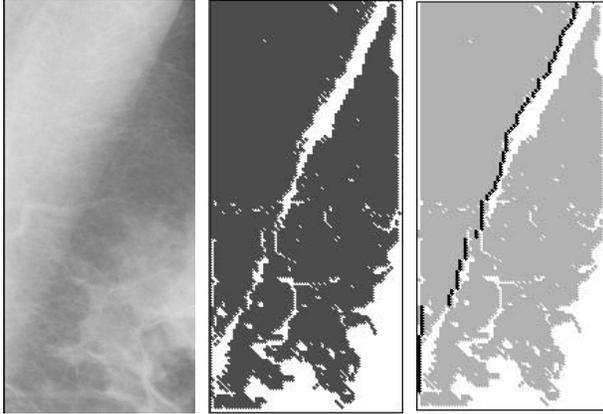


Figure 2. On the left is a close-up view of the pectoral muscle in image mdb118. The center image shows the component associated with the pectoral muscle (black region). This component has fused with the component associated with the natural dense tissue in the breast. The image on the right shows the result of adjusting the pectoral muscle boundary by requiring that the boundary point in each row appears in the same column or to the left of the boundary point in the row immediately above.

by MST were used to initialise a local segmentation method based on active contours.

The algorithm for implementing active contours will be described as it applies to the left breast. Let $V = v_1, v_2, \dots, v_N$ denote the set of pixels forming the current pectoral muscle boundary. Let (x_i, y_i) denote the coordinates of v_i , $i = 1, 2, \dots, N$. The adaptive deformable contour model works by moving the boundary through the spatial domain of the image to minimise a measure of energy based on the following formulas.

$$E_i = \alpha E_{in,i} + \beta E_{ex,i}, \quad (8)$$

where α, β are two weights controlling the internal and external energies $E_{in,i}$ and $E_{ex,i}$. The internal and external energies are given by

$$\begin{aligned} E_{in,i} &= a_1 V'(v_i) + a_2 V''(v_i) \\ E_{ex,i} &= -|I_x(v_i)| / \max(I_x), \end{aligned}$$

where $V'(v_i)$ and $V''(v_i)$ are the first and second derivatives of the contour V at v_i , I is the image, and

$$I_x = \frac{\partial I}{\partial x}. \quad (9)$$

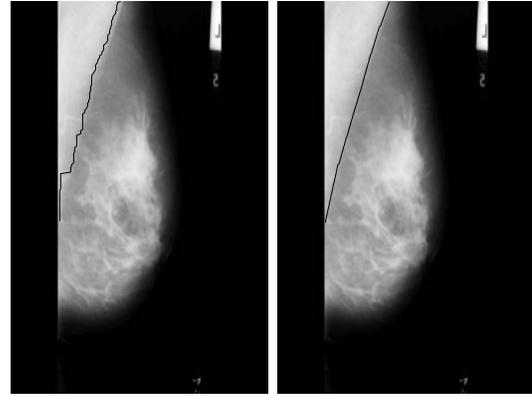


Figure 3. On the left is the boundary drawn by the MST algorithm without further processing. On the right is the boundary after using active contours to adjust the position and smoothness. The image is mdb118, the same as in Figure 2.

The weights a_1 and a_2 are used to control the relative contributions of $V'(v_i)$ and $V''(v_i)$ and were fixed for this study at $a_1 = 1$ and $a_2 = 2$.

The internal energy serves to reduce the curvature of the contour. This is important since the pectoral muscle has a general smooth straight shape. The external energy drives the contour toward strong edges in the image. This is important since the pectoral muscle generally appears much brighter in the image than other tissue.

The snake is an open curve that only moves horizontally in the image. This simplification takes advantage of prior information regarding the general appearance and location of the pectoral muscle in the image. The advantages are a reduction of complexity and the convenience of using row numbers as a fixed index for points at which the snake is evaluated.

At every point v_i , the energies are computed on the asymmetric neighbourhood

$$\Omega_i = [(x_i - 5, y_i), \dots, (x_i + 3, y_i)]. \quad (10)$$

Asymmetric neighbourhoods are used since the initial pectoral muscle boundary usually appears closer to the chest wall than the true boundary.

The weights for internal and external energy, α and β , are adjusted automatically as follows.

$$\begin{aligned} \alpha &= |x_i - x_{i-1}| + |x_{i+1} - x_i| - 2 * d \\ \beta &= \exp((\max_{\Omega_i} |I_x| - \min_{\Omega_i} |I_x|) / \text{mean}_{\Omega_i} |I_x|) \\ d &= (x_1 - x_N) / N. \end{aligned}$$

The parameter α measures the local deviation of the slope from the global average slope. When this value is small, the external forces are allowed to dominate in order to push the snake toward boundary. When α is large, the internal forces dominate to straighten the snake. Similarly, β forces the snake to converge to the boundary quickly when the snake is far away, but favours local straightening of the snake over the rate at which it approaches the boundary when it is close. As a result the snake attains an anatomically realistic local shape as it moves toward the boundary.

The elements of Ω_i will be denoted by $e_j, j = 1, 2, \dots, 9$, and the internal and external energies at these points will be denoted by $E_{in,i}^j$ and $E_{ex,i}^j$ respectively. Thus

$$\begin{aligned} E_{in,i}^j &= a_1 V'(e_j) + a_2 V''(e_j) \\ E_{ex,i}^j &= -|I_x(e_j)| / \max_I(I_x), \end{aligned}$$

where $V'(e_j)$ and $V''(e_j)$ are the derivatives along the curve obtained by replacing v_i by e_j .

To allow comparison between the different energy terms, it is necessary to rescale them to the range $[0, 1]$.

$$\begin{aligned} \hat{E}_{in,i}^j &= \frac{E_{in,i}^j - E_{in,i}^{\min}}{E_{in,i}^{\max} - E_{in,i}^{\min}}, \\ \hat{E}_{ex,i}^j &= \frac{E_{ex,i}^j - E_{ex,i}^{\min}}{E_{ex,i}^{\max} - E_{ex,i}^{\min}}, \end{aligned}$$

where the superscripts min and max denote the minimum and maximum of the respective quantities over the domain Ω_i . Thus the contour is driven to minimise

$$\hat{E}_i = \alpha \hat{E}_{in,i} + \beta \hat{E}_{ex,i}. \quad (11)$$

The energy of the contour is minimised iteratively. Each iteration consists of minimising \hat{E}_i for $i = 1, \dots, N$ consecutively. At a given step, the point v_i will be replaced to the point e_j , if

$$\hat{E}_i^j = \min \hat{E}_i^k, k = 1, 2, \dots, 9. \quad (12)$$

In this study the number of iteration was fixed at 30 although experiments showed that a stable contour was generally reached in fewer than 8 iterations.

4 Database

In order to compare the MST method with other methods in the literature, 84 images were obtained from the Mini-MIAS database of mammographic images [7]. More specifically, the same images were selected as those used in the study on identifying the pectoral muscle conducted by Ferrari, et al. [2]. The spatial resolution of these images is

200 μ m and depth resolution is 8 bit. The images in the database are 1024 \times 1024 pixels in size. For this study, the images were further subsampled to 256 \times 256 pixels.

In the work by Ferrari, et al. [2], boundaries found by the algorithm were compared to boundaries drawn by a radiologist. The coordinates of the lines drawn by the radiologist in that study were kindly provided by R. M. Rangayyan so that our results could be compared to the same standard.

5 Results

One of the files containing the coordinates of the pectoral muscle boundary drawn by the radiologist had an error and could not be used in the study. In two of the images the pectoral muscle was essentially missing in the image (Figure 4). Accordingly, the results presented here are based on 81 images rather than 84.

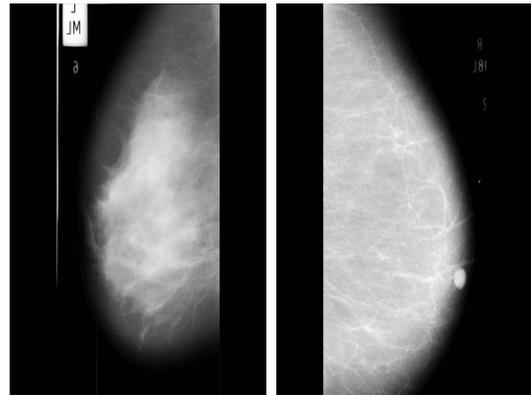


Figure 4. Unlike most MLO view mammograms (Figure 1) the pectoral muscle is essentially absent in these two images (mdb109 on the left and mdb098 on the right).

For each image, the number of pixels predicted to form part of the pectoral muscle by the algorithm but that fell outside the pectoral muscle region determined by the radiologist was recorded. This number was normalised by the total number of pixels in the pectoral muscle region determined by the radiologist to obtain the false positive (FP) score for the image (Table 1). Similarly, the number of pixels predicted by the algorithm to lie outside the pectoral muscle but were found by the radiologist to lie within the pectoral muscle were used to compute false negative (FN) score for the image. This method of reporting the results was chosen to allow direct comparison with results obtained by Ferrari et al. [2] and by the authors in a previous study [4]

Table 1. Comparison of pectoral muscle detection using MST with other methods reported in the literature using a common set of images. The results quoted for the Hough transform method and a method based Gabor wavelets were obtained by Ferrari et al. [2]. The results for adaptive pyramids (AP) are from previous work by the authors [4]. The columns labeled FP and NF list the average FP and NF scores computed over all the images in the study with the exception noted at the beginning of this section. The column labeled < 5 lists the number of images in the study with less than 5 percent total error (FP+FN). Similarly, the columns labeled 5-10 and > 10 list the number of images with total error in the ranges 5 to 10 percent and greater than 10 percent, respectively.

Methods	FP	FN	< 5	5-10	> 10
Hough	1.98	25.19	10	8	66
Gabor	0.58	5.77	45	22	17
AP	3.23	5.73	50	18	15
MST	1.64	12.03	25	31	25

6 Discussion

Table 1 shows that, in terms of assigning percentages of pixels correctly to the pectoral muscle, the MST algorithm does not perform quite as well as either the method based on Gabor wavelets or the method based on adaptive pyramids. Graph theory based methods are computationally intensive and are therefore in their infancy compared to the multi-resolution filters on which the Gabor method is based. The version of the MST algorithm described here is in some sense quite primitive and it is likely that improvements can be made. For example, only image intensities have been used to measure the within and between variation of clusters. The method lends itself to including other criteria such as texture measures, shape, position, etc. These extensions have not been explored.

Although improvements are clearly possible, from a practical point of view, the MST, Gabor wavelet, and AP algorithms are all able to determine the pectoral muscle close to what can reasonably be expected from any such scheme. The pectoral muscle boundary found by the MST algorithm in most cases agrees very well with the boundary drawn by the radiologist (Figure 5). There are two reasons why the error encountered in these cases can be viewed as negligible. In the first place, while the radiologist drawn boundaries

provide the best available gold standard, they are subjective. The visible boundary in the image represents the outer extent of a rounded surface, surrounded by soft tissue. As a result, the boundary is, in principle, not sharp at the pixel level. The notion of a correct boundary at this resolution is artificial. In the second place, the level and distribution of error along the boundary seen in the examples of Figure 5 will have little or no adverse effect on image registration or region based image analysis which are the ultimate objectives that motivate delineation of the pectoral muscle.

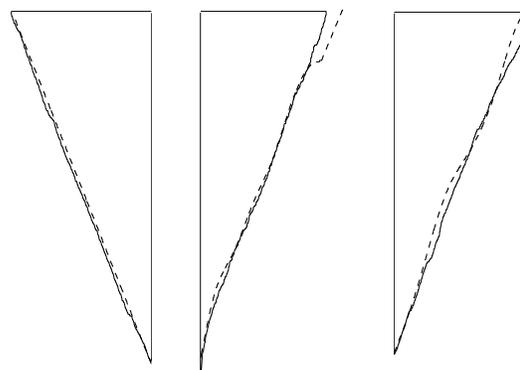


Figure 5. The solid lines show the region of the mammogram defined by the frame of the image and the radiologists drawn boundaries of the pectoral muscle. The dashed lines show the boundaries found by the MST algorithm followed by active contours. On the left is the result for image mdb099 and the result for image mdb090 appears in the centre. These are the same images as in Figure 1. The appearance of the pectoral muscle in image mdb099 (left) is without complications. Most MLO views are of this type. The pectoral muscle in image mdb090 (centre) is complicated by an anomaly near the top of the image (Figure 1). This bright spot resulted in the top portion of the predicted boundary being shifted to the right. On the right is image mdb118, the example of Figure 2. Despite the appearance of bright breast tissue near the pectoral muscle and the complication illustrated in Figure 2, the final result is acceptable.

Another pitfall in determining the pectoral muscle is the presence of false boundaries within the pectoral muscle (Figure 6). Depending on severity of such lines and the relative intensities on either side, the MST algorithm may or may not identify the correct boundary. In cases where

the MST algorithm adopts a false internal structure as the boundary of the pectoral muscle, the final error is substantial. A large amount of the error reported in Table 1 stems from a small number of images.

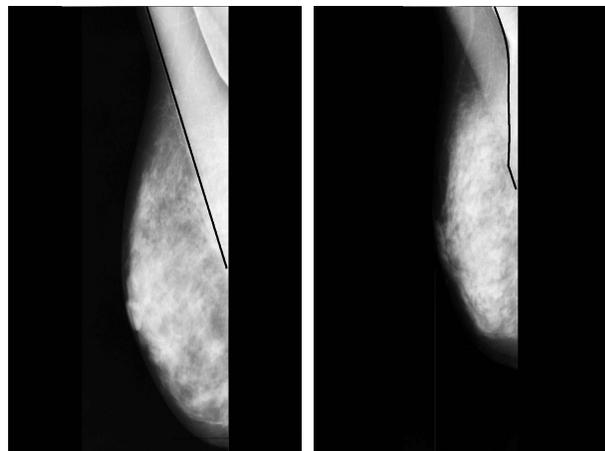


Figure 6. Examples of images showing boundary-like structures within the pectoral muscle region. The grey lines are the boundaries found by MST followed by active contours. The predicted boundary is acceptable in the left image (image mdb123) but the boundary in the right image (image mdb039) follows the false internal structure near the top of the image.

References

- [1] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *Int. Jour. Computer Vision*, 59(2):167–181, Sep 2004.
- [2] R. J. Ferrari, R. M. Rangayyan, J. E. L. Desautels, R. A. Borges, and A. F. Frère. Automatic identification of the pectoral muscle in mammograms. *IEEE Trans. Med. Im.*, 23(2):232–245, 2004.
- [3] J. Jolion and A. Montanvert. The adaptive pyramid: A framework for 2d image analysis. *Computer Vision, Graphics, and Image Processing*, 55(3):339–348, May 1992.
- [4] F. Ma, M. Bajger, and M. J. Bottema. Extracting the pectoral muscle in screening mammograms using a graph pyramid. In B. C. Lovel and A. J. Meader, editors, *Workshop Proceedings, APRS Workshop on Digital Image Computing (WDIC2005)*, pages 27–32. University of Queensland, February 2005.
- [5] M. Pavan and M. Pelillo. A new graph-theoretic approach to clustering and segmentation. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1063–1069, 2003.
- [6] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.*, 22(8):888–905, 2000.
- [7] J. Suckling, J. Parker, D. Dance, S. Astley, I. Hutt, C. Boggis, I. Ricketts, E. Stamatakis, N. Cerneaz, S. Kok, P. Taylor, D. Betal, and J. Savage. The mammographic images analysis society digital mammogram database. *Excerpta Medica, International Congress Series 1069*, pages 375–378, 1994.
- [8] Z. Whu and R. Leachy. An optimal graph theoretic approach to data clustering: Theory and its application to image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.*, 15(11):1101–1113, 1993.