

SEGMENTATION OF BREAST CANCER FINE NEEDLE BIOPSY CYTOLOGICAL IMAGES

MACIEJ HREBIENÍ, PIOTR STEĆ, TOMASZ NIECZKOWSKI, ANDRZEJ OBUCHOWICZ

Institute of Control and Computation Engineering
University of Zielona Góra, ul. Podgórna 50, 65–246 Zielona Góra, Poland
e-mail: {m.hrebien, p.stec, a.obuchowicz}@issi.uz.zgora.pl,
t.nieczkowski@weit.uz.zgora.pl

This paper describes three cytological image segmentation methods. The analysis includes the watershed algorithm, active contouring and a cellular automata GrowCut method. One can also find here a description of image pre-processing, Hough transform based pre-segmentation and an automatic nuclei localization mechanism used in our approach. Preliminary experimental results collected on a benchmark database present the quality of the methods in the analyzed issue. The discussion of common errors and possible future problems summarizes the work and points out regions that need further research.

Keywords: cytological image segmentation, Hough transform, watershed algorithm, active contours, cellular automata GrowCut method.

1. Introduction

In the last decade we have been observing a dynamic growth in the number of research works conducted in the area of breast cancer diagnosis. Many university centers and commercial institutions (Kimmel *et al.*, 2003) are focused on this issue because of the fact that breast cancer is becoming the most common form of cancer disease of today's female population. The attention covers not only curing the external effects of the disease (Boldrini and Costa, 1999; Świerniak *et al.*, 2003) but also its fast detection in an early phase. Thus, the construction of a fully automatic cancer diagnosis system supporting a human expert has become a challenging task.

Nowadays many camera-based automatic breast cancer diagnosis systems have to face the problem of cells and their nuclei separation from the rest of the image content (Lee and Street, 2000; Pena-Reyes and Sipper, 1998; Setiono, 1996; Wolberg *et al.*, 1993). This process is very important because the nucleus of the cell is the place where breast cancer malignancy can be observed. Thus, much attention in the construction of the expert supporting diagnosis system has to be paid to the segmentation stage.

The main difficulty of the segmentation process is due to the incompleteness and uncertainty of the informa-

tion contained in the image. The imperfection of the data acquisition process in the form of noise, chromatic distortion and deformity of cytological material caused by its preparation additionally increases the problem complexity. The nature of image acquisition (3D to 2D transformation) and the method of scene illumination also affect the image luminance and sharpness. In many cases one must also deal with a low-cost CCD sensor whose quality and resolution capabilities are rather low.

Until now many segmentation methods have been proposed (Carlotto, 1987; Chen *et al.*, 1998; Kass *et al.*, 1987; Otsu, 1979; Su and Chou, 2001; Vincent and Soille, 1991) but, unfortunately, each of them introduces numerous additional problems and usually works in practice under given assumptions and/or needs the end-user's interaction/co-operation (Lee and Street, 2000; Street, 2000; Wolberg *et al.*, 1993; Zhou and Pycock, 1997). Since nowadays many cytological projects assume full automation and real-time operation with a high degree of efficiency, a method free of drawbacks of the already known approaches has to be constructed.

In this paper a group of modified versions of cytological image segmentation methods adopted for fine needle biopsy images are presented, that is, the watershed algorithm, active contours and the cellular automata GrowCut technique. One can also find here a descrip-




Fig. 1. Exemplary fragments: (a) cytological image, (b) appropriate segmentation mask.

tion of denoising and contrast enhancement techniques, pre-segmentation and a fully automatic nuclei localization mechanism used in our approach.

2. Problem formulation

The mathematical formulation of the segmentation process is very difficult because it is a poorly conditioned problem, and in many situations the segmentation process is domain specific. Thus, we give here only an informal definition of the problem we have to face.

What we have on input is cytological material obtained using the Fine Needle Biopsy (FNB) technique and imaged with a *Sony CCD Iris* camera mounted on top of an *Axiophot* microscope. The material comes from female patients of the *Onkomed* medical center in Zielona Góra (Marciniak *et al.*, 2005). The image itself is coded using the RGB colorspace and is not subject to any kind of lossy compression (a raw color bitmap format). The size of the image equals 704×576 pixels.

What we expect on output is a binary segmentation mask with one pixel separation rule which will permit more robust morphometric parameters estimation in our future work. Additionally, the proposed segmentation algorithm should be insensitive to colors of contrasting pigments used for the preparation of the cytological material (see the example in Fig. 1).

3. Image filtering and preparation

3.1. Preprocessing. The color components of an image do not carry as important information as the luminosity does, so they can be removed to reduce processing complexity in stages that require only, e.g., gradient estimation. An RGB color image can be converted to greyscale by removing blue and red chrominance components from the image defined in the YCbCr colorspace. The luminosity component can be determined using (Pratt, 2001):

$$Y = 0.299R + 0.587G + 0.114B. \quad (1)$$

Since the majority of images we deal with have low contrast, an enhancement technique is needed to improve




Fig. 2. Typical histogram and a cumulated sum P (dashed line) before contrast enhancement.

their quality. In our approach we use simple histogram processing with a linear transform of the image levels of intensities. The range of cutoff is defined by $(t_{\text{low}}, t_{\text{high}})$ points (Fig. 2):

$$\begin{aligned} t_{\text{low}} &= \arg_i^{(1)} (P(i) \geq \theta_{\text{low}}), \\ t_{\text{high}} &= \arg_i^{(n)} (P(i) \leq \theta_{\text{high}}), \end{aligned} \quad (2)$$

where P is a cumulated sum of all image pixels described by its histogram H (Russ, 1999):

$$P(i) = \frac{\sum_{j=0}^i H(j)}{\sum_{j=0}^{N-1} H(j)}, \quad i = 0, \dots, 255, \quad (3)$$

and $\arg^{(1)}$, $\arg^{(n)}$ denote the first and the last argument, respectively. θ_{low} and θ_{high} define a plateau in the histogram and in our approach they are equal to 0.01 and 0.99, respectively, which means that 1% of pixels are saturated at low and 1% at high intensities of the input image.

The contrast correction is conducted for each color channel separately resulting in an image being better defined for later stages of the presented hybrid segmentation methods (see the example in Fig. 3).

3.2. Presegmentation.

3.2.1. Background. If we look closely at the nuclei we have to segment, we notice that they all have an elliptical shape (Fig. 4). Most of them resemble an ellipse but, unfortunately, the detection of the ellipse which is described by two parameters a and b ($x = a \cos(\alpha)$, $y = b \sin(\alpha)$) and which can be additionally rotated is computationally expensive. The shape of the ellipse can be approximated by a given number of circles (as shown in Fig. 4(b)). The detection of circles is much simpler in the sense of the required computations because we have only one parameter, which is the radius R . These observations and simplifications form a basis for a nucleus presegmentation algorithm—in our approach we try to find such circles with different radii in a given feature space.




Fig. 3. Exemplary fragments of cytological images (a), (c) and their contrast enhanced equivalents (b), (d).




Fig. 4. Exemplary fragment of a cytological image with circular nuclei.

3.2.2. Circle detection. The Hough transform (Ballard, 1981; Duda and Hart, 1972; Żorski, 2000) can be easily adopted for the purpose of circle detection. The transform in the continuous space is defined by

$$HT_{cont}(R, x_0, y_0) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x, y) \delta\left((x - x_0)^2 + (y - y_0)^2 - R^2\right) dx dy, \quad (4)$$

where f is any continuous function defining the space in which the circles will be looked for and δ is Dirac's delta (an infinite impulse at zero), which defines integration only over the circle.

The transform in a discrete space (with similar meaning) can be rewritten as

$$HT_{discr}(R, i_0, j_0) = \sum_{i=i_0-R}^{i_0+R} \sum_{j=j_0-R}^{j_0+R} g(i, j) \delta\left((i - i_0)^2 + (j - j_0)^2 - R^2\right), \quad (5)$$

where g is a two-dimensional feature image and δ is Kronecker's delta (equal to unity at zero). HT_{discr} plays the




Fig. 5. Gradient masks used in our experiments.

role of an accumulator which accumulates the similarity levels of g to the circle placed at (i_0, j_0) and defined by the radius R .

The feature space g can be created in many different ways. In our approach we use a gradient image as the feature indicating the occurrence or absence of the nucleus in a given fragment of the cytological image. The gradient image is a saturated sum of gradients estimated in eight directions on the greyscale image prepared in the preprocessing stage. The base gradients can be calculated using, e.g., Prewitt's, Sobel's mask methods (Gonzalez and Woods, 2002; Tadeusiewicz, 1992) or their heavy or light versions (Fig. 5).

3.2.3. Final actions. Thresholding the values in the accumulator by a given θ value we can obtain a very good presegmentation mechanism with a lower threshold strategy (see, e.g., Fig. 7). Since the threshold value strongly depends on the database and the feature image g (Fig. 6), the method can only be used as a presegmentation stage. A smaller value of the threshold causes fast removal of unimportant information from the background, and what we achieve is a mask which approximately defines the places where the objects we have to segment (nuclei in this case) are located and where the background is. Such a mask can constitute a basis for more sophisticated and detail-oriented algorithms.




Fig. 6. Influence of the threshold θ on objects' cover, and the lack of differences (left) and overcovering (right) for Prewitt (\times), Sobel ($*$), heavy (\bullet) and light ($+$) base gradient masks (experiments performed on a randomly selected 346 element Zielona Góra *Onkomed* cytological benchmark database for radii in the range 4–21).

3.3. Nuclei localization.

3.3.1. Terrain modeling. The results obtained at the presegmentation stage can lead to the estimation of an av-




Fig. 7. Exemplary results of the presegmentation stage for two different θ threshold strategies: (a) high and (b) low.

erage background color. Such information can be used to model the nuclei as a color distance between the background and the objects, which fulfills the requirements of the lack of any color dependency in the imaged material (the color of contrasting pigments may change in the future). In our research we tried several distance metrics: Manhattan's, Chebyshev's, the absolute Hue value from the HSV colorspace, but the Euclidean one gives us visually the best results (Fig. 8(a) and (b)):

$$D_{\text{euclid}} = \sqrt{(I_R - B_R)^2 + (I_G - B_G)^2 + (I_B - B_B)^2}, \quad (6)$$

where B is the average background color estimated for the input image I .




Fig. 8. Exemplary fragments: (a) cytological image, (b) Euclidean distance to the mean background color, (c) smoothed version of (b).

Since the modeling distance can vary in the local neighborhood (see Fig. 8(b)), mostly because of camera sensor simplifications, a smoothing technique is needed to reconstruct the nuclei shape. The smoothing operation in our approach relies on the fact that this sort of 2D signal can be modeled as a sum of sinusoids (Madiseti and Williams, 1997) with defined amplitudes, phase shifts and frequencies. Cutting off all low amplitude frequencies (leaving only a few significant ones with the highest amplitude) will result in a signal deprived of our problematic local noise effect (Figs. 8(c), and 9). The frequency spectrum is determined using the discrete Fourier transform

$$F(m) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi nm/M}, \quad (7)$$




Fig. 9. One-dimensional representation of the images from Figs. 8(b) and (c), respectively.

where x is a one-dimensional signal created by grouping pixels column by column from the square window which is moved through the image. After the frequency cutoff, the signal is reconstructed using the inverse of the discrete Fourier transform. And what we finally achieve is a three-dimensional modeled terrain where hills correspond to nuclei.

3.3.2. Nuclei search. The localization of objects on a modeled map of nuclei can be performed locally using various methods. In our approach we have chosen an evolutionary (1+1) search strategy (Arabas, 2004; Michalewicz, 1996) mostly because it is simple and quite fast, can be easily parallelized due to its nature and settles very well in local extrema, which is very important in our case.

The search in our approach can be conducted in two versions: single-point and multi-point. In the single-point version it is allowed to have only one marker pointing a nucleus while in the multi-point one it is allowed to have more than one marker pointing the same nucleus.

Single-point version with the background. As a final segmentation method, the watershed algorithm used forced us to create two populations of individuals. The first population localizes the background. Specimens are moved with a constant movement step ($R = 1$) preferably to places with a smaller density of population to maximize the background coverage. The second population localizes the nuclei. Specimens are moved with a decreasing movement step (r^t) to group very fast the population near local extrema in the first few epochs and to finally work on details in the ultimate ones. The movement of individuals is preferred to be directed towards places with a higher population density to create the effect of nuclei localization.

The fitness function ϕ calculates the average *height* of the terrain in a given position including the nearest neighborhood defined by the smallest radius detected by the Hough transform in the presegmentation stage. Such a definition of the fitness function avoids a possible split of the population, localized near a nucleus with a multi-




Fig. 10. Exemplary single-point localization: (a) screenshot after 8 epochs, (b) final result (localization points are marked with asterisks).

Algorithm 1. The (1+1) search strategy for nuclei localization.

```

t = 0
initialize Xt

while t < tmax
  Yt = mutate Xt

  if φ(Xt) < φ(Yt)
    Xt = Yt
  end

  t = t + 1
end

```

modal character of its shape, giving only one marker for a nucleus (Fig. 10(b)).

Finally, the nucleus is localized in the place where the density of the population searching for hill tops in the modeled terrain is locally maximal. As has been mentioned, the method is quite fast and just a few epochs are needed to observe a visible progress in nuclei localization and background coverage (Fig. 10(a)).

The presented algorithm for nuclei localization is given as Algorithm 1, where

$$Y_i^t = X_i^t + r^t N_i(0, 1) \quad (8)$$

defines the mutation and

$$r^t = R_{\max} \left(\frac{1}{R_{\max}} \right)^{\frac{t}{t_{\max}}} \quad (9)$$

defines the decreasing movement step. The algorithm for the population which is covering the background is similar, except $R = 1$ during mutation.




Fig. 11. Exemplary multi-point localization: (a) screenshot after several epochs, (b) final result.

Multi-point version. The algorithms that do not have such tight requirements concerning only one single marker per nucleus (that is, they allow multiple markers pointing the same one, not optimal or even false localization point and can take information about the background location from the presegmentation mask (Hrebień and Steć, 2006)) can use a very simplified version of the above (1 + 1) search strategy. In such cases we can use only one population, i.e., the one searching for nuclei, and the fitness function is simply the terrain *height* in an individual position. The number of iterations of the algorithm can also be reduced, because we need only an approximate localization of nuclei (Fig. 11). Thus, the algorithm is the same as the one given in Algorithm 1, and the only difference with the one described above is the fitness function ϕ and a reduced number of epochs.

4. Image segmentation

4.1. Watersheds. The watershed segmentation algorithm is inspired by natural observations, such as a rainy day in the mountains (Gonzalez and Woods, 2002; Pratt, 2001; Russ, 1999). A given image can be defined as a terrain on which nuclei correspond to valleys (upside down the terrain modeled in previous steps). The terrain is flooded by rainwater and arising puddles start to turn into basins. When the water from one basin begins to pour away to another, a separating watershed is created.

The flooding operation has to be stopped when the water level reaches a given threshold θ . The threshold should preferably be placed somewhere in the middle between the background and a nucleus localization point. In our approach the nuclei are flooded to the half of the altitude between the nucleus localization point and the average height of the background in the local neighborhood. Since the images we have to deal with are spot illuminated during the imaging operation (which results in a modeled terrain higher in the center of the image and much lower in the corners), this mechanism protects basins from being overflowed and, in consequence, nuclei from being undersegmented (Hrebień *et al.*, 2007).

Algorithm 2. A simplified version of the watershed algorithm.

```

// For each basin
 $\forall p \in P$  assign a label  $(i + 1)$ 

for  $\theta \in 0 : \Delta : 1$ 
  for  $\forall p \in P$ 
    color  $p$  to the level  $\leq \theta\Psi(p)/2$ 
  end
end

```

A simplified version of the watershed algorithm is given in Algorithm 2. The coloring to the level of θ implements the flooding operation. It also considers a possible situation of watershed building when there is a neighbor with another label. Δ defines the water level increase in each iteration of the algorithm and Ψ defines the difference between the p valley's depth and the background's height in its local neighborhood.

4.2. Active contours. The active contouring technique (Blake and Isard, 1998; Kass *et al.*, 1987) can be considered as a more advanced region growing method (Tadeusiewicz, 1992). The algorithm groups neighboring pixels when a given homogeneity and similarity criteria are met. All joined pixels create a segment whose boundary spreads in all directions until another segment is met or new candidates for joining introduce unacceptable errors. The algorithm is stopped when all pixels get labels, i.e., the object in the image is separated from the background.

The images we deal with may contain more than a single object per image. Additionally, the assumption of the project is that the segmentation process has to be fully automatic (there is no human operator which manually initializes the method). These two factors force us to modify the algorithm to meet the stated requirements. Thus, the algorithm, which in our case is based on the fast marching method (FMM) (Sethian, 1999), must have a multilabel extension (Steć and Domański, 2005) and the seeding process has to be done without the end-user's interaction.

In the proposed approach the multilabel FMM is initialized with a presegmentation mask and the results obtained at the multipoint nuclei localization stage. The background-object boundary from the presegmentation mask is the initial seed for the background segment. On the other hand, the nuclei localization points are initial seeds for the object segments. The most important issue in this method is that the initialization mask and the nuclei localization points do not have to be perfect—all fake initial markers are fully acceptable and they do not have any

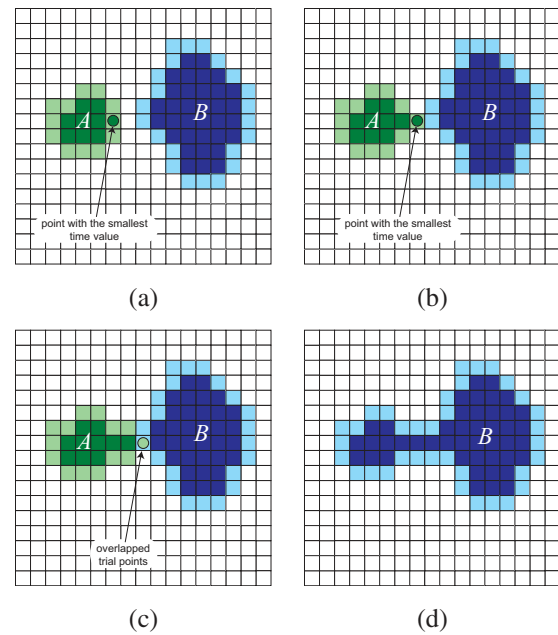


Fig. 12. Illustration of the contour merging operation.

influence on the final segmentation result and its quality.

The contour expansion speed of the multilabel FMM is governed globally by the function (Hrebień and Steć, 2006):

$$F = \frac{1}{|g(x, y) - \bar{g}(i)|^3 + 1}, \quad (10)$$

where $g(x, y)$ is the color under the contour and $\bar{g}(i)$ is the mean color under the i -th segment. Such a definition of the expansion speed slows down the contour near the object (nucleus) boundary. Two very close spreading segments can meet during the algorithm execution. The two meeting segments can be merged (the smaller one into the bigger one) when their mean color difference is below a threshold (Fig. 12). The segments not classified to be merged can bush back the segment with the lower difference between the analyzed pixel and mean color of each segment (Fig. 13). The pushing operation can be performed only once to reduce contour oscillations known from the classical approach and the segment pushed back cannot move further here.

4.3. Cellular automata. The next technique inspired by natural observations is the GrowCut cellular automata segmentation algorithm (Vezhnevets and Konouchine, 2005). It imitates the growth and struggle for domination of rivalry bacteria colonies. Each type of bacteria represents a single type of objects used in segmentation. The GrowCut algorithm was originally developed for multi-label intelligent scissor tasks for photo-editing tools. It requires manual initialization of the seed pixels

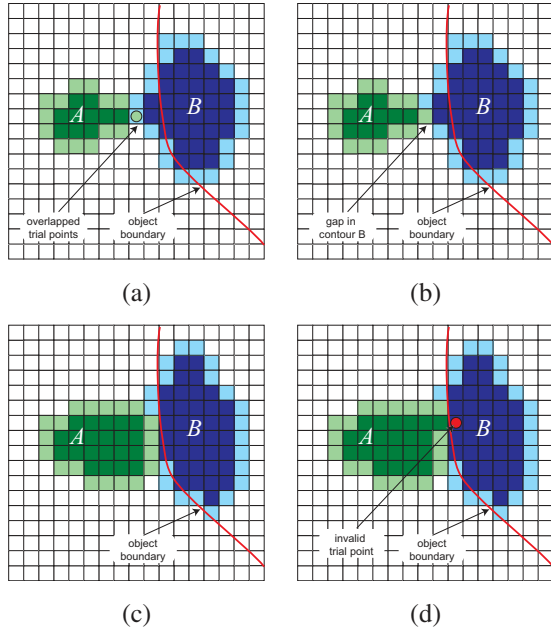


Fig. 13. Illustration of the contour pushing operation.

but, concatenated with a proper presegmentation method, gives a fully automated hybrid segmentation technique.

The GrowCut algorithm defines a cellular space P as a $k \times m$ array, where k and m are the image dimensions. Each of the array cells is an automaton described by the state triplet $(l_p, \theta_p, \vec{C}_p)$, where l_p is the label of the cell, θ_p is the *strength* of the cell and \vec{C}_p is the feature vector of the cell defined by the associated image pixel. An unlabeled image may be then considered as a particular configuration state of the cellular automaton, where the initial states for each $p \in P$ are set to

$$l_p = 0, \quad \theta_p = 0, \quad \vec{C}_p = RGB_p, \quad (11)$$

where RGB_p is the three-dimensional vector of pixel p color in the RGB space. The final goal of the segmentation is to assign each pixel to one of K possible labels. As stated before, we use two labels in the segmentation of cytological images—the *nuclei* and the *background*.

In a single evolution step, each cell (the bacteria) tries to attack all its neighbours. The evolution goal is to occupy the whole image area starting from a group of previously initialized pixels. Cell neighbours are defined by a neighbourhood system. In our approach the Moore neighbourhood system was used:

$$N(p) = \left\{ q \in \mathbb{Z}^n : \|p - q\|_\infty := \max_{i=1, \dots, n} |p_i - q_i| = 1 \right\}. \quad (12)$$

The attack power is defined as a function of the strengths of attacker q and defender p and the distance between their feature vectors, \vec{C}_q and \vec{C}_p . The basic rule of the automaton state change at time $t + 1$ is shown in Algorithm 3.

Algorithm 3. GrowCut algorithm.

```

// For each cell
for  $\forall p \in P$ 
    // copy previous state
     $l_p^{t+1} = l_p^t$ 
     $\theta_p^{t+1} = \theta_p^t$ 
    // neighbors try to attack
    // current cell
    for  $\forall q \in N(p)$ 
        if  $g(\|\vec{C}_p - \vec{C}_q\|_2) \cdot \theta_q^t > \theta_p^t$ 
             $l_p^{t+1} = l_q^t$ 
             $\theta_p^{t+1} = g(\|\vec{C}_p - \vec{C}_q\|_2)\theta_q^t$ 
        end
    end
end
end
    
```

The function g is monotonous, decreasing and bounded to $[0, 1]$. For the purpose of this work, a simple function g was used, as proposed in (Vezhnevets and Konouchine, 2005):

$$g(x) = 1 - \frac{x}{\max \|\vec{C}\|_2}, \quad (13)$$

where $\max \|\vec{C}\|_2$ is calculated as a feature vector length for white pixels ($RGB = [255, 255, 255]$). As the strength of each cell is increasing and bounded, the method is guaranteed to converge. Thus for any seed configuration of the image, after a finite number of evolution steps, all cells are labelled and their states change. Figure 14 shows subsequent steps of the GrowCut segmentation for a manually initialized cytological image.

The GrowCut algorithm requires the initialization of a number of cells with proper labels for each separate, consistent group of pixels (segment seed). To permit the application of the algorithm to the automated diagnostic system, we employ the information from the presegmentation and the nuclei localization stage to initialize the seed pixels. At this point almost any rough segmentation




Fig. 14. Exemplary segmentation of an FNB image with the GrowCut algorithm. White points denote *nucleus* labelled cells, black points represent *background* labelled cells: (a) the seed, (b) Step 2, (c) Step 4, (d) Step 6, (e) Step 19 (end).




Fig. 15. Exemplary segmentation with the GrowCut algorithm initialized with the thresholding result: (a) exemplary image, (b) thresholding result, (c) GrowCut result.

technique (e.g., thresholding) can be also applied as pre-segmentation (Fig. 15): however, our research shows that initialization which leaves unclassified pixels at object boundaries performs better (Hrebień *et al.*, 2006). One of the techniques whose results can be utilized at the GrowCut algorithm initialization stage is the presegmentation mask obtained using the Hough transform. The transform result is a set of circles covering the regions of the image where nuclei are located. Pixels enclosed inside these regions are initially labeled as the *nucleus* ones. The remaining pixels of the image are labeled as the *background*. For this type of initialization, all the image pixels are classified before the first GrowCut evolution step. The goal of the algorithm application is only to adjust the segment edges to real boundaries of objects. Therefore, to enforce a proper direction and a range of label changes within successive evolution steps, associating suitable values of the initial strength for both of the pixel classes is necessary.

The appropriate direction of label changes depends on the threshold value θ used at the presegmentation stage. For lower threshold values, the Hough transform results in a number of background pixels located at boundaries of regions labeled as *nucleus*. These pixels should change their labels to *background* in the process of actual segmentation. Thus, the initial strength value of the *nucleus* labeled pixels has to be less than the strength of the background pixels. For higher values of the threshold θ , a number of the *nucleus* pixels are incorrectly labeled as *background*. In this case the labels of the boundary pixels should be changed to *nucleus*. Therefore, the initial strength of the *nucleus* pixels has to be greater than the *background* pixels.

The GrowCut algorithm can be also initialized with the result obtained at the multi-point nuclei localization stage described above. Due to few initialized pixels of each segment, cell strengths can be set to equal values for both classes. It permits the automation of the segmentation process. However, more uninitialized pixels result in more evolution steps and so a greater computational cost.

5. Experimental results

5.1. Watersheds. Exemplary results of the presented watershed segmentation method and common errors ob-

served in our hand-prepared benchmark database can be divided into four classes:

- *Class 1*: good quality images with only small irregularities and rarely generated subbasins (a basin in another basin) (Figs. 16(a) and (b)),
- *Class 2*: errors caused by fake circles created by spots of fat (Figs. 16(c) and (d)),
- *Class 3*: mixed nucleus types: red and purple in this case, and those of the red which are more purple than yellow (background) are also segmented, which is erroneous (Figs. 16(e) and (f)),
- *Class 4*: poor quality image with a bunch of nuclei glued together, which causes basin overflowing and, in consequence, undersegmentation (Figs. 16(g) and (h)).

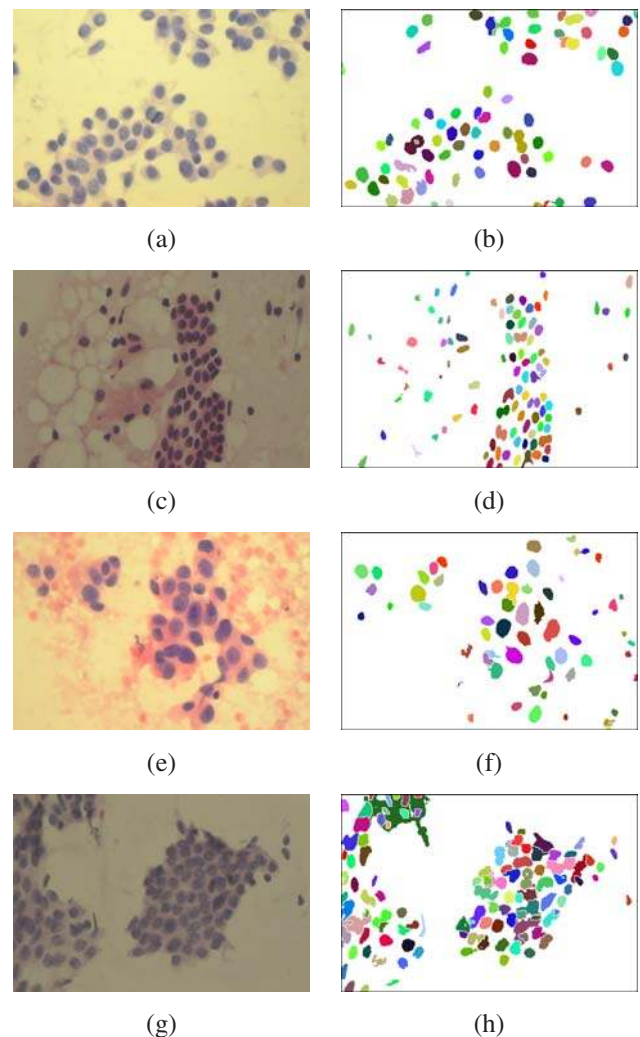


Fig. 16. Exemplary results of the watershed segmentation.




Fig. 17. Exemplary results of active contouring segmentation.

The conducted experiments show that, on average, the watershed algorithm gives a 68.74% agreement with the hand-prepared templates using a simple XOR metric. Most errors are located at boundaries (see, e.g., Fig. 18(a)) of nuclei where the average distance between the edges of segmented and reference objects is about 3.28 pixels. The XOR metric is underestimated as a consequence of a moderate level of water flooding the modeled terrain, but the shape of the nuclei seems to be preserved, which is important in our future work—the estimation of morphometric parameters of segmented nuclei.

5.2. Active contours. The conducted experiments show that the modified multilabel FMM algorithm is very stable and robust to initialization errors. Visually, segmentation quality is promising and yields a good detection of even small objects (Fig. 17). Unfortunately, the algorithm has problems with connected nuclei and detects them as a single object, which is erroneous. The average XOR metric score with the hand-prepared templates is only 22.32%, and the average distance between the edges of segmented and reference objects is about 4.1 pixels.

Despite the above-mentioned problems, the shape of segmented nuclei seems to be represented accurately, and most errors are located at the boundaries of the segmented objects (see, e.g., Fig. 18(b)). This illustrates that a proper selection of the merging threshold and the detection of overlapping nuclei is still a challenge and has to be improved in future works.

5.3. Cellular automata. The conducted experiments show that the influence of the threshold θ and the initial

strength value of the seed pixels on the final segmentation quality can be observed. Figure 19 presents the results for an exemplary image—a number of incorrectly labelled pixels, relative to the initial strength value of the *nucleus* class for a range of the threshold θ . The initial strength of the *background* class had a fixed value equal to 1.0.

It can be observed that for lower values of the threshold θ the optimum initial value of the *nucleus* seed pixels needs to be lower than that of the *background* seeds. It allows changing labels of the boundary pixels from *nucleus* to *background*, so nuclei shapes are adjusted outside-in. The best segmentation results were obtained for the θ threshold $\theta = 0.99$ and the initial *nucleus* strength $S_n = 1.1$.

The percentage of incorrectly labeled pixels (Zhang, 1996) relative to the overall area of another exemplary image is shown in Fig. 20. The best obtained result was 10.4%. The segmentation result is visualised in Fig. 21.

Even better results were obtained for segmentation with the GrowCut cellular automata, initialized with the result obtained at the multi-point nuclei localization stage. Figure 22 presents the initialization mask and the results of segmentation. For the exemplary image the proportion of incorrectly labelled pixels was about 6%. However, the shape of the identified nuclei segments is too ragged, so an additional smoothing post-segmentation stage is needed for this combination of techniques.

The problem with the Hough transform and the GrowCut cellular automata hybrid is that the optimum proportion of the initial *nucleus* pixel strength should be estimated to achieve good segmentation quality. The proportion strongly depends on the analysed image contrast and the pigment used, so a potential automated diagnostic system should be taught beforehand. The second hybrid (the GrowCut cellular automaton), initialized with a multi-point nuclei searching algorithm, can be applied with fixed initial strength values for the *nucleus* and the *background* seed pixels. However, it is much more computationally expensive due to more cellular automata iterations required.




Fig. 18. Exemplary XOR results with a fragment of the hand-prepared segmentation mask for: (a) watershed algorithm. (b) active contouring technique.




Fig. 19. Number of incorrectly labelled pixels for a range of the initial strength S_n of *nucleus* pixels and the threshold θ : 0.95 (+), 0.96 (o), 0.97 (*), 0.98 (●), 0.99 (x). Solid lines—final segmentation results, dashed lines—presegmentation result.

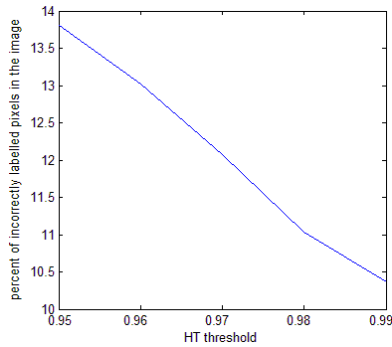


Fig. 20. Percentage of incorrectly labeled pixels relative to the overall area of the image for a range of the θ threshold values.

6. Conclusions

The conducted preliminary experiments show that the Hough transform adopted for circle detection at the presegmentation stage, the (1+1) search strategy used for automatic nuclei localization and the watershed, active contours and cellular automata GrowCut algorithms employed for the final segmentation stage can be effectively used for the segmentation of cytological images. Unfortunately, there are still some problems requiring further research.

The problem regarding fake circles created by spots of fat and unwanted effects it gives as the final output should be considered and eliminated in future work. Images with a mixed nucleus type still constitute a challenge because it seems to be impossible to detect only one type without the end-user's interaction and when there should not be any dependencies and assumptions concerning the color of contrasting pigments used to prepare cytological material. The presented solutions should also be extended to perform better on poor quality images, or a fast classifier should be constructed to reject too poor (or even fake)




Fig. 21. Exemplary result of segmentation for the threshold $\theta = 0.99$ and the initial strength of the *nucleus* labeled pixels $S_n = 1.1$ (dark color underlines fitness to reference objects).




Fig. 22. Exemplary result of segmentation (b) with the GrowCut cellular automata, initialized with multi-point nuclei localization points (a).

inputs. All of the above-mentioned drawbacks of the described algorithms should also be considered and faced in the future.

The described segmentation methods needs about 4–5 minutes on today's machines (Athlon 64 3500+ 2.8 GHz, Pentium 4 2.2 GHz) for an image to give a final response. All the simulations were performed in the Matlab environment so the authors believe that the time consumption can be significantly reduced. Computationally expensive stages (like, e.g., terrain modeling) and the nature of the algorithms used (e.g., the (1+1) search strategy) force and allow us at the same time to implement them using a high level of parallelism to perform more effectively on today's dual core machines and thread oriented operating systems (preliminary lower level implementations show that the whole segmentation process can be closed in several, about 20-30, seconds). One can also consider a dedicated hardware which can perform even more effectively than the appropriate software implementation (according to the authors' estimations, e.g., the above-mentioned terrain modeling stage can be done by DSP hardware nearly 100 times faster on average, including the clock frequency factor, than an efficient software implementation running on today's machines¹).

¹Compared with TMS TM-44 (0.13 μm CMOS Tech., 100 MHz, 8000 mW, 2001), which can perform 1024-point FFT at 11.04 μsec (according to producer information)

Summarizing, the presented solutions are promising and give a good base for our further research in the area of cytological image segmentation. Additionally, all preparation steps including pre-segmentation and the automatic nuclei localization stage can be reused with other segmentation algorithms which need such information.

References

- Arabas J. (2004). *Lectures on Evolutionary Algorithms*, WNT, Warsaw (in Polish).
- Ballard D. (1981). Generalizing the Hough transform to detect arbitrary shapes, *Pattern Recognition* **13**(2): 111-122.
- Blake A., Isard M. (1998). *Active Contours*, Springer, London.
- Boldrini J. and Costa M. (1999). An application of optimal control theory to the design of theoretical schedules of anti-cancer drugs, *International Journal of Applied Mathematics and Computer Science* **9**(2): 387-399.
- Carlotto M. (1987). Histogram analysis using a scale space approach, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **9**(1): 121-129.
- Chen C., Luo J. and Parker K. (1998). Image segmentation via adaptive K-mean clustering and knowledge-based morphological operations with biomedical applications, *IEEE Transactions on Image Processing* **7**(12): 1673-1683.
- Duda R. and Hart P. (1972). Use of the Hough transformation to detect lines and curves in picture, *Communications of Association for Computing Machinery* **15**(1): 11-15.
- Gonzalez R. and Woods R. (2002). *Digital Image Processing*, Prentice Hall, Englewood Cliffs, NJ.
- Hrebień M., Nieczkowski T., Korbicz J. and Obuchowicz A. (2006). The Hough transform and the GrowCut method in segmentation of cytological images, *Proceedings of the International Conference on Signal and Electronic Systems ICSES'06*, Łódź, Poland, pp. 367-370.
- Hrebień M. and Steć P. (2006). Fine needle biopsy material segmentation with Hough transform and active contouring technique, *Journal of Medical Informatics and Technologies* **10**: 25-34, (in print).
- Hrebień M., Korbicz J. and Obuchowicz A. (2007). Hough transform, (1+1) search strategy and watershed algorithm in segmentation of cytological images, *Proceedings of the 5th International Conference on Computer Recognition Systems CORES'07*, Springer, Wrocław, pp. 550-557.
- Kass M., Witkin A. and Terzopoulos D. (1987). Snakes: Active contour models, *Proceedings of the 1st International Conference on Computer Vision*, pp. 259-263.
- Kimmel M., Lachowicz M. and Świerniak A. (Eds.) (2003). Cancer growth and progression, mathematical problems and computer simulations, *International Journal of Applied Mathematics and Computer Science* **13** (3) (Special Issue).
- Lee M. and Street W. (2000). Dynamic learning of shapes for automatic object recognition, *Proceedings of the 17th Workshop Machine Learning of Spatial Knowledge*, Stanford, CA, pp. 44-49.
- Madisetti V. and Williams D. (1997). *The Digital Signal Processing Handbook*, CRC Press, Boca Raton, FL.
- Marciniak A., Obuchowicz A., Monczak R. and Kołodziński M. (2005). Cytomorphometry of fine needle biopsy material from the breast cancer, *Proceedings of the 4th International Conference on Computer Recognition Systems CORES'05*, Springer, Rydzyna, Poland, pp. 603-609.
- Michalewicz Z. (1996). *Genetic Algorithms + Data Structures = Evolution Programs*, Springer, London.
- Otsu N. (1979). A threshold selection method from grey-level histograms, *IEEE Transactions on Systems, Man and Cybernetics* **9**(1): 62-66.
- Pena-Reyes C. and Sipper M. (1998). Envolving fuzzy rules for breast cancer diagnosis, *Proceedings of the International Symposium on Nonlinear Theory and Application, Vol. 2, Polytechniques et Universitaires Romandes Press*, pp. 369-372.
- Pratt W. (2001). *Digital Image Processing*, Wiley, New York.
- Russ J. (1999). *The Image Processing Handbook*, CRC Press, Boca Raton, FL.
- Sethian J. (1999). Fast marching methods, *SIAM Review* **41**(2): 199-235.
- Setiono R. (1996). Extracting rules from pruned neural networks for breast cancer diagnosis, *Artificial Intelligence in Medicine* **8** (1): 37-51.
- Steć P. and Domański M. (2005). Video frame segmentation using competitive contours, *Proceedings of the 13th European Signal Processing Conference EUSIPCO'05*, Antalya, Turkey, pp. 4 (CD-ROM).
- Street W. (2000). Xcyt: A system for remote cytological diagnosis and prognosis of breast cancer, in: (Jain L. (Ed.)), *Soft Computing Techniques in Breast Cancer Prognosis and Diagnosis*, World Scientific Publishing, Singapore, pp. 297-322.
- Su M. and Chou C. (2001). A modified version of the K-means algorithm with a distance based on cluster symmetry, *IEEE Transactions Pattern Analysis and Machine Intelligence* **23**(6): 674-680.
- Świerniak A., Ledzewicz U. and Schättler H. (2003). Optimal control for a class of compartmental models in cancer chemotherapy, *International Journal of Applied Mathematics and Computer Science* **13**(3): 357-368.
- Tadeusiewicz R. (1992). *Vision Systems of Industrial Robots*, WNT, Warsaw, (in Polish).
- Vezhnevets V. and Konouchine V. (2005). "GrowCut"—interactive multi-label N-D image segmentation by cellular automata, *Proceedings of the 15th International Conference on Computer Graphics and Applications GraphiCon'05*, Novosibirsk, Russia, pp. 150-156.
- Vincent L. and Soille P. (1991). Watersheds in digital spaces: An efficient algorithm based on immersion simulations, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(6): 583-598.
- Wolberg W., Street W. and Mangasarian O. (1993). Breast cytology diagnosis via digital image analysis, *Analytical and Quantitative Cytology and Histology* **15**(6): 396-404.

Zhang J. (1996). A survey on evaluation methods for image segmentation, *Pattern Recognition* **29**(8): 1335-1346.

Zhou P. and Pycock D. (1997). Robust statistical models for cell image interpretation, *Image and Vision Computing* **15**(4): 307-316.

Żorski W. (2000). *Image Segmentation Methods Based on the Hough Transform*, Studio GiZ, Warsaw, (in Polish).

Received: 16 June 2007

Revised: 19 November 2007