

Robust Cell Nuclei Segmentation Using Statistical Modelling

Theodoros Mouroutis Stephen J. Roberts Anil A. Bharath

24 February 1998

Abstract

The objective analysis of cytological and histological images has been the subject of research for many years now. One of the most difficult fields in histological image analysis is the automated segmentation of tissue-section images. We propose a multistage segmentation method for the isolation of cell nuclei in such images. In the first stage the Compact Hough Transform (CHT) is used to determine possible locations of the nuclei. We then define a likelihood function which enables us to perform an optimization procedure based on the maximization of this function. Global grey-level histogram information is used throughout the algorithm to reduce the amount of computation and to reject unwanted artefacts. The algorithm was tested on connective tissue images with very encouraging results. Apart from detecting well-separated nuclei in the images, it performed well in separating dividing nuclei into likely substructures. At the same time the algorithm provides us with a measure of uncertainty along the detected boundary, in the form of the value of the likelihood function.

1 Introduction

In the last few years a number of studies have been presented that aim at rendering the interpretation of cytological images more objective [1, 2, 3, 4, 5]. The success of these studies has been boosted by the increase in computational power of modern computer systems. This has allowed the automated performance of tasks which were otherwise performed manually. However, the majority of cytological and histological examinations are still performed manually, by inspection carried out by experts. This way of interpreting cytological information, derived either from individual cells, from a cell population, or from a tissue section is, unavoidably, affected by human subjectivity. Clinicians with various levels of experience may give different opinions on the same samples, while the same expert's opinion may vary depending on factors such as fatigue.

Among the various types of cytological images, the ones that present the highest degree of complexity are those of tissue sections, the main reasons being the diversity of the structures contained in such a section and the intense variations in staining. The majority of the work presented in the field of automated cytology and histology has to do with the analysis of information from individual cells or from biopsies obtained by brushing or needle aspiration. In these cases, the cells are well distinguishable from the background and usually from each other, too.

In this paper, we propose a method for the analysis of images obtained by light microscopy from stained tissue sections. Our goal is to isolate the nuclei of the cells contained in the image for further analysis. The nucleus is a very important structure within the cell, as far as the diagnosis of the severity of a tumour is concerned. In cancer, the cell nuclei undergo significant changes which, if quantified, can allow the diagnosis and, potentially, the prediction of the course of the disease.

A problem we will encounter is that of dividing nuclei, which is similar to the one of overlapping nuclei in thicker tissue sections. Whilst for overlapping nuclei there is a border, in the case of dividing nuclei there is no such border to be detected. We show that, using information from the rest of the nuclear boundary, our algorithm proposes a suitable border in regions where the division has not yet been completed but also gives a likelihood measure which is low (high uncertainty) in these regions.

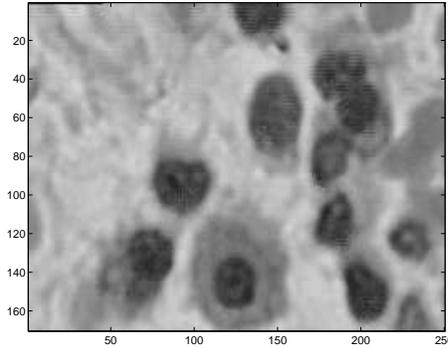


Figure 1: *Image of a section of typical connective tissue. The real size of this region is $56 \times 82 \mu\text{m}$.*

2 Materials and methods

The images used in this work are microscopic images of laryngeal tissue sections. The sections have been prepared and provided by the Department of Histopathology of the University College Hospital in London. Their thickness is $3 \mu\text{m}$. In sections of that thickness we will not encounter cases of overlapping nuclei but cases of dividing nuclei may occur. The slices have been stained using the H&E (haematoxylin and eosin) method and embedded in parafin. This method results in the nuclei appearing darker than the cytoplasm and the cytoplasm darker than the extracellular matter.

The thickness of the slices used is only a fraction of the diameter of a typical nucleus and this unavoidably leads to many nuclei in the sections appearing smaller than they actually are. Ideally, for the extraction of diagnostic conclusions of clinical significance based on nuclear features, more parallel sections of the same tissue sample should be taken in order for the most appropriate plane of view to be determined for each individual nucleus. However, in this study, the main objective is the presentation of the segmentation algorithm, so this issue will not be addressed. We wish, furthermore, to develop a practical automated system which performs well using slices which are prepared according to clinical practice.

The images have been acquired using a Leitz Aristoplan light microscope using a $40\times$ oil immersion lens ($\text{NA}=0.65$). The microscope was connected to a JVC KY-F55B colour CCD camera, whose output was fed into a Neotech image grabber and was stored as a *.tif* image. The camera records using the PAL system and had a S/N ratio of 58 dB. At the magnification used the side of each pixel of the resulting image corresponds to $0.33 \mu\text{m}$. For our image analysis the RGB image was transformed into an 8-bit grey-level image¹, using the conversion formula

$$i = 0.33r + 0.5g + 0.17b$$

where r, g and b are the intensities of the red, green and blue components of the colour image and i the resulting grey-scale intensity.

2.1 Nuclear segmentation

We can see on the typical tissue section image of Figure 1 that even though the cell nuclei appear, in general, darker than the surrounding cytoplasm and extracellular matter, not all the nuclei possess the same chromatin densities, which even vary within the same nucleus. This results in significant grey-level variations inside the nucleus and among different nuclei. The variations are such that segmentation of the nuclei from the rest of the image using simple, or even adaptive thresholding is not sufficient. Furthermore, the existence of other dark regions in the tissue images, corresponding to nuclear debris, other cellular or extracellular structures, as well as artefacts of the staining procedure and image acquisition, make the idea of thresholding even less appealing.

Previous studies showed that the result of nuclear or cell segmentation could be greatly improved if information about the shape of the searched object were incorporated in the segmentation process

¹We used the public domain package XV, version 2.21, developed at the University of Pennsylvania

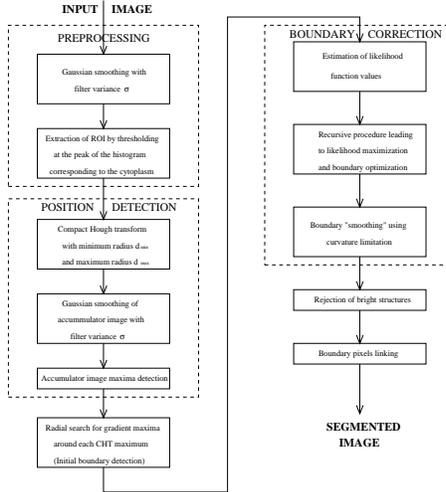


Figure 2: *Block diagram of the proposed segmentation algorithm*

[6, 7, 5]. The method we propose is a hybrid segmentation method in the sense that we use both global grey-scale information and local edge information, while we also take advantage of prior knowledge regarding the morphological characteristics of the cell nuclei, which are expected to be closed shapes with boundary which does not present strong irregularities.

A block diagram of the algorithm we propose in the present study is shown in Figure 2.

2.2 Preprocessing

The preprocessing of each image consists of two stages. In the first stage, we try to reduce the number of image edges irrelevant to the nuclei and suppress the effect of those edges on the detection algorithm. In the second stage we select a region of the image in which we can safely locate the nuclei, using only histogram information. The establishment of this region of interest (ROI) is entirely to reduce the computational time, and has no effect on the computed results.

If we closely inspect images of nuclei, we can immediately remark that within each of the nuclei there are grey-scale variations, some of which can be quite intense. An example is the case of the nucleolus, which typically has a grey-scale distinctly darker than the average grey-scale of the nucleus. To improve the efficiency of our detection algorithm, we wish to minimize or, if possible, eliminate the effects of these local grey-scale variations. Since our detection strategy has an edge-oriented component, the occurrence of edges other than the nuclear boundary will complicate the analysis.

We have chosen a simple approach in which we initially smooth the image with a Gaussian filter [8] (for other approaches, see [9, 10]). Care, however, must be taken in the choice of the filter variance σ which should be large enough to remove undesired edges but not as large as to result in image oversmoothing. The choice of σ is related to the choices of parameters regarding the bounds in the sizes of cell nuclei. We, therefore, discuss the determination of σ in section 2.3.

Based on the histogram of Figure 3, where the peaks correspond to the nuclear interior (μ_n), the cytoplasm (μ_c) and the extracellular matter, if we threshold our image at μ_c and reject the regions with higher grey-scale, we retain all those regions in which nuclei may be located even if their staining is not intense. Optimal thresholding [11] is not appropriate for our application, as it results in a threshold t which leaves part of the nuclei outside this selected region, as shown in Figure 4. It also leaves significant part of the nuclear boundary outside this optimally thresholded region of interest. However, as we will see in section 2.3, the nuclear localization procedure uses edge information which, thus, needs to be preserved.

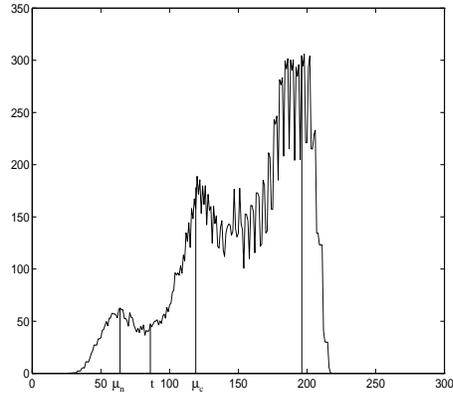


Figure 3: *Smoothed tissue image grey level histogram.*

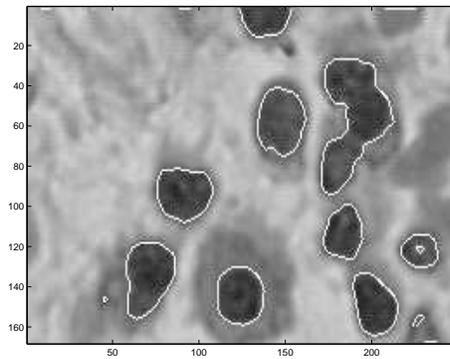


Figure 4: *Result of thresholding the image of Figure 1 at the grey-level value corresponding to the optimal threshold between the grey-level distributions relative to the nucleus and to the cytoplasm. Observe that parts of the nuclei are excluded from the selected regions.*

2.3 Localization and initial segmentation of the nuclei

2.3.1 The Compact Hough Transform (CHT)

One observation that will help us formulate our detection algorithm is that, despite the variations in the shape of the nuclei, we do not come across any extremely elongated ones. Also, we observe that there are no very large variations in the size of the nuclei, e.g. all lie within one order of magnitude of each other. We can, therefore, define a maximum and a minimum possible size that we expect the nuclei to have.

Let us assume that after observing one of our digital images, we determine that for every nucleus there is a maximum (d_{max}) and a minimum (d_{min}) distance at which any pixel on the nuclear boundary could lie from the centroid of the nucleus (Figure 5a). These parameters are obviously related to the magnification at which we are looking at our image, but also, in the case of cancerous tissue, to the degree of the disease. Healthy nuclei tend to have approximately the same size throughout the volume of a particular kind of tissue, while cancerous nuclei may become very enlarged or even very small. By adjusting these parameters at will one is able to search for objects of different sizes. Let us stress at this point that the limits on the distances d_{max} and d_{min} may be extremely loose by setting the two distances much larger and much smaller than the actual maximum and minimum possible distances of the nuclear boundary from the centroid of the nucleus. However, the restriction of the nuclear size to an order of magnitude is necessary, as there are other closed shapes on the same image with size that differs significantly from that of the nuclei and whose detection we are not interested in. The range of the limits d_{max} and d_{min} will obviously also affect the computational load of the algorithm.

The width of the smoothing filter \mathbf{S} may be obtained by noting that, according to Nyquist's sampling theorem, two points \mathbf{x}_1 and \mathbf{x}_2 will merge (share a common extremum) under the smoothing operation if

$$|\mathbf{x}_1 - \mathbf{x}_2| < 2\sigma$$

If the points are the centroids of the two components of a dividing nucleus and we still wish to distinguish between them when they are $d_{min}/2$ apart, then σ must satisfy

$$2\sigma \leq \frac{d_{min}}{2}$$

or, equivalently

$$\sigma \leq \frac{d_{min}}{4}$$

For greatest noise smoothing results we choose the highest possible value of σ , i.e. $d_{min}/4$.

In the majority of studies in which the Hough transform is used for the detection of circles, ellipses or other arbitrary shapes, edge direction information plays a significant role, both because it reduces the computation performed during the implementation of the transform and, mainly, because it renders the algorithm more robust. In order to incorporate edge direction information in the transform, we need to have some information about the shape which we seek to detect. In our case, the nuclear shape is very irregular and we cannot intue precise prior information about the direction of the edges at its border. We can, however, assume that the nuclei are largely convex structures. This means that along the majority of the nuclear border the nuclear interior will lie in only one of the semi-planes defined by the tangent to the boundary. In mathematical terms, we can express this by saying that, if \mathbf{x} is the location of the edge pixel in our image and \mathbf{y} a location in the accumulator image, we increment the element $A(\mathbf{y})$ if the distance condition is satisfied and

$$(\mathbf{x} - \mathbf{y}) \cdot \hat{\mathbf{n}} > 0$$

where $\hat{\mathbf{n}}$ is the gradient vector at \mathbf{x} . Based on the above assumptions, we can initiate an accumulator array A with dimensions identical to those of our image. All the elements of the array are initially zero. For every significant edge pixel of our image, we increment the value of the elements of the accumulator array that lie within the semi-annulus centered at the edge pixel, with inner and outer radii d_{min} and d_{max} respectively, and bounded by the line perpendicular to the direction of the gradient at the edge pixel (Figure 5b).

Our CHT algorithm may thus be summarized as follows

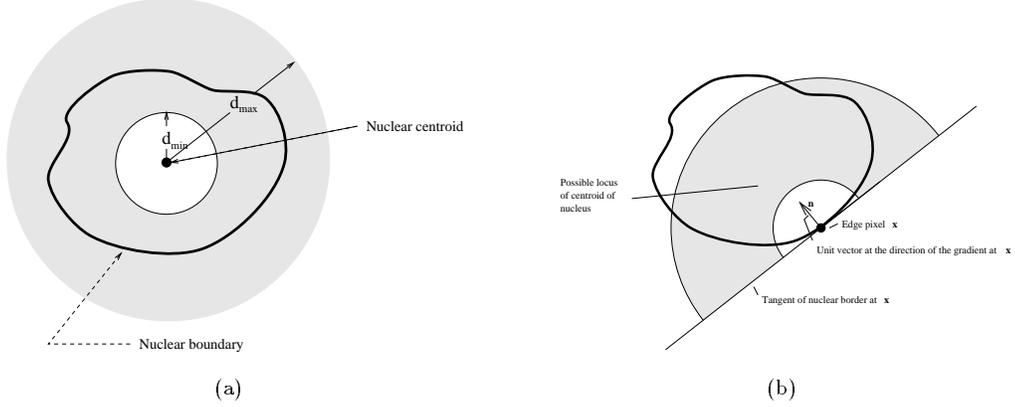


Figure 5: (a) All the boundary points of the nucleus lie at distances between d_{min} and d_{max} from the nuclear centroid. (b) Compact Hough transform: the value of the accumulator array cells in the grey semi-annulus is incremented

for $i=1$ to no_of_rows
 for $j=1$ to $no_of_columns$
 if $G(\mathbf{x})$ is significant
 for all $A(\mathbf{y})$ with
 $d_{min} < |\mathbf{x} - \mathbf{y}| < d_{max}$ and $(\mathbf{x} - \mathbf{y}) \cdot \hat{\mathbf{n}} > 0$
 $A(\mathbf{y}) = A(\mathbf{y}) + 1$

where G is the edge image and A the accumulator image. If we follow Sklansky's notation [12], the CHT may be expressed as the convolution

$$A(\mathbf{x}) = T_{\phi}(\mathbf{x}) * E(\mathbf{x})$$

where E is the binary image generated by setting

$$E(\mathbf{x}) = \begin{cases} 1 & \text{if } G(\mathbf{x}) \text{ is significant} \\ 0 & \text{otherwise} \end{cases}$$

and T_{ϕ} is a binary mask generated from a semi-ring with inner radius d_{min} and outer radius d_{max} centered at the origin, and rotated by an angle ϕ so that the line bisecting the ring is aligned with the gradient direction at \mathbf{x} , i.e.

$$T_{\phi}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \text{ is internal to the semi-ring} \\ 0 & \text{otherwise} \end{cases}$$

Obviously, we need to define what we mean when we use the term 'significant'. Before we do so, however, let us see how our CHT is linked with the original Hough transform (HT) for the detection of circles. The connection is quite simple and is expressed as

$$CHT(\mathbf{x}) = \int_{r=d_{min}}^{r=d_{max}} HT(\mathbf{x}, r) \cdot dr$$

where r is the radial parameter in the circular Hough transform accumulator array [13]. This simple relation would allow the fast adaptation of existing Hough transform calculation methods to the problem of detecting closed curves of *a priori* unknown convex shapes in an image. The three-dimensional parameter space of the original circular Hough transform is reduced to a two-dimensional plane where the value of the accumulator array is a measure of the likelihood of the individual array element being located inside such a closed curve.

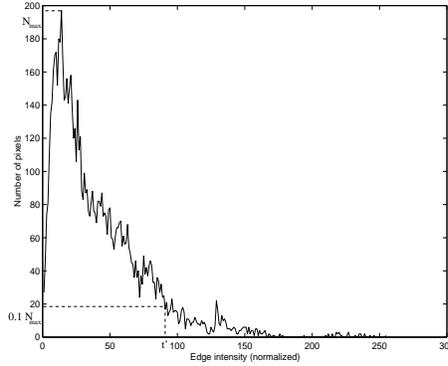


Figure 6: *The distribution of the normalized (in the range 0-255) edge intensities of the image of Figure 1, following smoothing and isolating the region of interest. The edge pixels with values greater than t' are selected as significant.*

To obtain edge information we apply standard Sobel masks to the smoothed image and threshold the gradient image in order to reject the large number of edges due to noise and local grey-level fluctuations. For the application of the CHT we only consider as significant the edges above the threshold shown in Figure 6. In this figure we see the histogram of the pixel edge intensities over the edge image masked by the ROI, normalized in the range 0-255. The threshold t' is set at the edge value for which the corresponding number of pixels has dropped to 10% of its maximum value N_{max} , attained at lower edge intensity values. One can possibly try to determine normal distributions of edges intensities in order to separate significant from non-significant edges, as proposed in [14], but for the use of edge information for the CHT, which is tolerant to missing information, a simple thresholding is sufficient and straightforward. Our choice of incrementing the accumulator image by a fixed quantity and not by a quantity dependent upon the magnitude of the edge as proposed in [13, 10], is imposed by the observation of the non-uniformity of the nuclear staining.

The resultant accumulator image is itself then processed after the CHT has been applied. This processing avoids the subsequent detection of a large number of erroneous candidate points by smoothing the accumulator image with the same Gaussian filter that was used on the original image itself.

2.3.2 Likelihood function

Once the candidate nuclei locations have been determined, we next need to trace the boundary around them. The nuclear boundary detection consists of three stages, i.e.

1. We determine an initial set of boundary pixels, based only on information provided by the image edges.
2. We use a maximum-likelihood algorithm to improve the initial set of pixels.
3. We use a method of posterior probability maximization to determine the optimal set of boundary pixels.

We perform a radial search from each CHT maximum along a number of search lines. In this case, we have chosen 32 directions. A power of 2 was preferable for implementation reasons and this number seemed the most suitable for the nuclear diameters at the given magnification. Clearly, this choice is a trade-off between boundary resolution and computational load. Our initial guess for the nuclear boundary along each search line is the location of maximum gradient. This may not be the actual boundary point, either because there are other edges in the image interfering or because of the smoothing of the image. Each of the 32 pixels is denoted by \mathbf{x}_i , while the entire set of pixels is

$$\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{32}\}$$

The set of CHT maxima for which no local gradient maximum is detected along some search line are rejected. We then associate each of the 32 pixels of the proposed boundary set with a likelihood function $l(\mathbf{x}_i|\mathbf{X})$. This likelihood function is derived as the product of three partial likelihoods associated with the pixels' grey-scale values, distances from the centroid of the structure and location relative to the gradient maxima along the search lines. For all three of these components, we assume Gaussian dependence of the corresponding characteristic and estimators of the parameters of the normal distributions for each of the \mathbf{x}_i are calculated from the set $\mathbf{X}' = \mathbf{X} - \{\mathbf{x}_i\}$. According to our description, the likelihood function $l(\mathbf{x}_i|\mathbf{X})$ has the form

$$l(\mathbf{x}_i|\mathbf{X}) = p_{gr}(\mathbf{x}_i|\mathbf{X}) \cdot p_d(\mathbf{x}_i|\mathbf{X}) \cdot p_e(\mathbf{x}_i) \quad (1)$$

where

$$\begin{aligned} p_{gr}(\mathbf{x}_i|\mathbf{X}) &= \frac{1}{\sqrt{2\pi} \cdot \sigma_{gr}} \cdot \exp\left\{-\frac{(B(\mathbf{x}_i) - \mu_{gr})^2}{2\sigma_{gr}^2}\right\} \\ p_d(\mathbf{x}_i|\mathbf{X}) &= \frac{1}{\sqrt{2\pi} \cdot \sigma_d} \cdot \exp\left\{-\frac{(d(\mathbf{x}_i) - \mu_d)^2}{2\sigma_d^2}\right\} \\ p_e(\mathbf{x}) &= \sum_{i=1}^N \frac{1}{\sqrt{2\pi} \cdot \sigma} \cdot \exp\left\{-\frac{(d(\mathbf{x}) - \mu_i)^2}{2\sigma^2}\right\} \cdot g_i \end{aligned}$$

In the above equations

$$\mu_{gr} = E\{B(\mathbf{X}')\}$$

$$\sigma_{gr}^2 = Var\{B(\mathbf{X}')\}$$

$$\mu_d = E\{d(\mathbf{X}_{\mathbf{6-n}})\} \quad (2)$$

$$\sigma_d^2 = Var\{d(\mathbf{X}_{\mathbf{6-n}})\} \quad (3)$$

$\mathbf{X}_{\mathbf{6-n}}$ being the set consisting of the six neighbours of \mathbf{x}_i in \mathbf{X} , and $B(\cdot)$ and $d(\cdot)$ denoting the grey-scale and distance from the structure centroid functions. Function $p_e(\mathbf{x})$ is constructed as a mixture of Gaussians centred at the locations μ_i of gradient maxima along the search lines. The Gaussians have amplitude g_i proportional to the gradient maximum value and standard deviation equal to that of the original smoothing filter, to counterbalance the uncertainty in the location of edges introduced by the smoothing operation.

2.4 Likelihood maximization

In this stage of the boundary optimisation, we seek to find the set of pixels that maximizes the total likelihood of the system, defined as

$$L = \prod_{i=1}^{32} l(\mathbf{x}_i|\mathbf{X})$$

We achieve this by following an iterative procedure which increases the value of the likelihood function L by modifying the proposed set of pixels, one direction at a time.

Let

$$\mathbf{X}^0 = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{32}\}$$

be the initial set of 32 boundary pixels generated as discussed earlier. For each of these pixels we calculate the individual likelihood functions

$$l_i^0 = l(\mathbf{x}_i|\mathbf{X}^0), \quad i = 1, 2, \dots, 32$$

as well as the total likelihood function

$$L(\mathbf{X}^0) = \prod_{i=1}^{32} l_i^0$$

Let

$$j = ArgMin(l_i^0)$$

be the index of the pixel with the lowest value of individual likelihood function. We search along the j -th radial search line for the pixel \mathbf{x}_m , with m given by

$$m = \text{ArgMax}(L(\mathbf{X}_i^{mod})), \quad k = 1, 2, \dots, N$$

that maximizes the value of the the total likelihood function. \mathbf{X}_i^{mod} is the modified set of pixels created from \mathbf{X}^0 by substituting \mathbf{x}_j with \mathbf{x}_k and N is the number of pixels examined along each search line. Once the pixel \mathbf{x}_m has been determined \mathbf{x}_j is substituted by \mathbf{x}_m in the original \mathbf{X}^0 . If no pixel has been found along the j -th search line that increases the value of L compared to $L(\mathbf{X}^0)$, pixel \mathbf{x}_j remains intact. We then repeat the search along the search line corresponding to the pixel with the next lowest value of the individual likelihood function. Eventually all the search lines have been searched and the resulting set of pixels is denoted with \mathbf{X}^1 .

The process is continued iteratively and every time, after a complete search along the search lines, a new set of pixels, characterised by greater likelihood is generated. The set created after the t -th iteration will be denoted with \mathbf{X}^t . The optimization process is terminated when, after an entire iteration we find that

$$\mathbf{X}^{t+1} = \mathbf{X}^t$$

The convergence of the algorithm we present is guaranteed. Each iteration leads by definition to a non-decreasing likelihood value:

$$L(\mathbf{X}^{t+1}) \geq L(\mathbf{X}^t)$$

Furthermore, $L(\mathbf{X})$ is bounded above by the optimal likelihood value in the (finite) possible set of combinations of the pixels lying on the chosen radial directions and at distances between d_{max} and d_{min} around each CHT maximum. Since there are only a finite number of possible pixel sets, we conclude that the system will converge, at the most, in this finite number of steps. However, it is possible that the state the system reaches after the completion of the maximum likelihood optimization does not actually correspond to the ideal nuclear boundary. The complexity of the texture of a tissue section image is such that errors can often occur for individual boundary points. Also, in the case of nuclei in the process of division or shortly after its completion, the nuclear boundary is either not clear or nonexistent. In many cases this would result in a proposed nuclear boundary which presents irregularities incompatible with the expected shape of a nucleus. This issue of regularity of the nuclear boundary is rectified by the use of a constraint correlating the distances of neighbouring boundary points.

2.5 Prior knowledge inclusion

As previously discussed, in the construction of the likelihood function we have assumed that the distance of boundary pixels from the nuclear centroid follows (locally) a normal distribution. Now, in order to achieve regularity of the nuclear boundary, we express this assumption in the form of a constraint: each boundary pixel's distance *has* to be correlated with its neighbouring boundary pixels in a strict rule of the form

$$\mu_d(\mathbf{x}_i) - \alpha \cdot \sigma_d(\mathbf{x}_i) < d(\mathbf{x}_i) < \mu_d(\mathbf{x}_i) + \alpha \cdot \sigma_d(\mathbf{x}_i)$$

where $\mu_d(\mathbf{x}_i)$ and $\sigma_d(\mathbf{x}_i)$ are given by Equations 2 and 3 and α is a constant, characteristic of the desired regularity. Along each radial search direction we find the pixels that lie within this particular range of distances and among them we choose the one that maximizes the likelihood function.

The selection of α is important, as a very small α will not allow the boundary pixel set to follow the curves of the actual nuclear border and will tend to produce a circle around the nuclear centroid. On the other hand, a very large α will not have the desired regularization effect and the pixel set derived from the likelihood maximization will carry all its irregularities. In practice, we find that $\alpha = 1.5$ gives good results in the regularization procedure. The value of α was set on a training image and was not adjusted later.

This restriction, as well as the obvious smoothing effect of the nuclear border has one more consequence: the regions that are most affected by this process are the ones where the likelihood

has low values, where the border is not well defined. Especially in the case of dividing nuclei, the proposed pixels will be the result of optimization on near random grey-scale fluctuations, where distance will play a minor role. With the constraint, distance information from nearby, well-defined edge pixels is explicitly included while those well-defined pixels are not affected due to their high likelihood function value.

To produce the continuous nuclear boundary from the set of 32 border pixels we use a method similar to the Bresenham line algorithm [15] and the resulting boundary is then coded with an 8-connected chain code [16, 17].

2.6 Grey-level rejection

Now that we have estimated a set of boundaries around each CHT maximum, we proceed to reject structures that do not satisfy a grey-level criterion. By sampling grey-scale values from each candidate nucleus, we can impose a stricter grey-level criterion than the one for the determination of the ROI. Indeed, in reference to the histogram of Figure 3, we have assumed that the population of the samples from the nuclei have a grey-level compatible with that of the first histogram peak at μ_n , while for the cytoplasm, the grey-level samples are from the second peak at μ_c . Assuming normal distributions, the probability of a pixel \mathbf{x} interior to a nucleus being of a grey-level B_i will be

$$p[B(\mathbf{x}) = B_i] = \frac{1}{\sqrt{2\pi} \cdot \sigma_n} \exp\left(-\frac{(B_i - \mu_n)^2}{2\sigma_n^2}\right) \quad (4)$$

and the probability of $g(\mathbf{x})$ being greater than a threshold T is (assuming 8-bit grey coding)

$$p(g(\mathbf{x}) > T) = \frac{1}{\sqrt{2\pi} \cdot \sigma_n} \sum_{i=T+1}^{255} \exp\left(-\frac{(i - \mu_n)^2}{2\sigma_n^2}\right)$$

In the above equations, σ_n is the variance of the nuclear grey-level distribution. A sample of N pixels from a nucleus would have an average following a normal distribution with the same average but with a variance σ_n/\sqrt{N} . The sampling method we followed involved selecting all the pixels along two directions perpendicular to each other and intersecting at the structure centroid. For a typical nuclear diameter of 10 μm this method of sampling gives a total of approximately 50-60 pixels, resulting in a decrease of the variance, for the average grey-level of the samples, by a factor of 1/7. This observation allows us to reject all the structures for which the average grey-level of the sample pixels is higher the threshold t (Figure 3), at the valley between μ_n and μ_c .

3 Results and discussion

We have applied our algorithm to several images, with very encouraging results. In Figure 7a we see the region of interest (ROI) superimposed on the original image. The region of interest comprises, as well as the nuclei, also part of the cytoplasm in the image and omits a significant part of the image which contains edges that would slow down the processing. In Figure 7b we see the final result of our segmentation.

Note the three nuclei at the top right of the image. One of the three nuclei, even though it is not well segmented initially, is brought to a better state by the optimisation algorithm and finally, with the second optimisation conditional to distance priors, is extracted (Figure 8). For the other two nuclei, which were extracted while in the process of division, the algorithm detects the easily identifiable boundary segments and chooses the maximum likelihood solution for the region where the division was still in progress.

For all the images on which the optimisation method we presented was used, two or three iterations of the likelihood optimisation procedure followed by between one and four iterations of the optimisation loop with the constraint presented in section 2.5 were enough to achieve the nuclear segmentation.

As we mentioned earlier, the value of α for this constraint which gave the best result was $\alpha = 1.5$. Figure 9 shows the result produced when α is given the values of 1 and 2. When $\alpha = 1$ the distance restriction is too tight and this results in the boundary not being able to follow the natural curves of the nuclear border. The result is a "rounded" approach of the actual boundary.

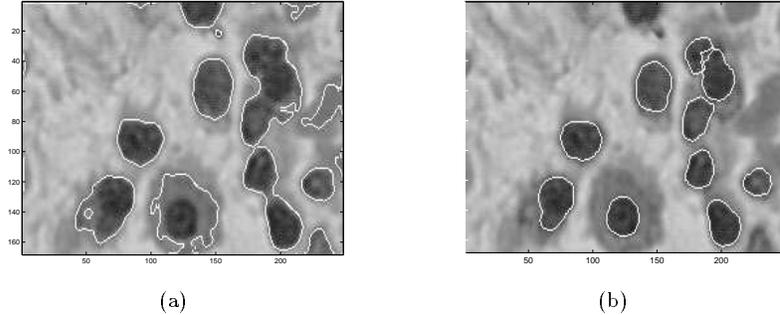


Figure 7: (a) The boundaries of the region of interest (ROI), and (b) Segmented image.

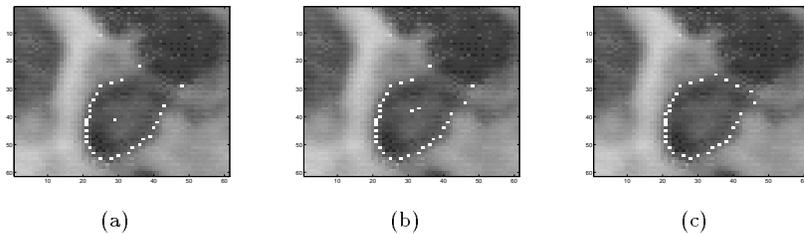


Figure 8: Two stages in the boundary optimisation process applied on the image of Figure 1 (detail): (a) Initial boundary selection, (b) Boundary following maximum likelihood optimisation, (c) Boundary following maximum likelihood optimisation with the distance prior included.

As the value of α increases, the system becomes more tolerant to irregularities of the proposed boundary. For $\alpha = 2$ we observe that the system is no longer able to correct the proposed boundary for the bottom left nucleus of the group of three nuclei previously discussed, while for larger α practically no correction is made to the result of the main optimisation process. Furthermore, we have found that for $\alpha = 1$ the number of iterations needed for the system to converge often exceeded 10. With larger α this is reduced to 1-2, but the effect of the optimization may be very small. We found that with $\alpha = 1.5$ we could successfully process approximately 95% of the images of nuclei in our database for which a reasonable first approximation of the boundary was possible using the gradient information (approximately 200 nuclei).

Figures 10a and 10b show the result of the application of our segmentation strategy on two other test images of different complexity. On the image of Figure 10a the segmentation has been accurately performed. On the more complex image of Figure 10b the segmentation has accurately detected the borders of most of the nuclei in the image. For the nucleus on the upper left, however, the algorithm suggests the existence of two substructures. This may be due to the distribution of the chromatin density inside the nucleus. One nucleus (centred on $[120,40]$) is completely ignored by the algorithm. The nucleus missed by our detection strategy is quite weakly and unevenly stained and it also appears to be partially covered by another nucleus.

In Figure 11a we see the histogram showing the distribution of the log-likelihood values for all the boundary pixels in the image of Figure 1. The majority of the likelihood values are in the range $[\epsilon^{-8}, \epsilon^{-4}]$. By thresholding these values at the position $\epsilon^{-8.66}$, where the histogram presents a local minimum, we can select the boundary pixels with high uncertainty of belonging to a true boundary. In practical applications continuous log-likelihood measures are to be preferred. Figure 11b shows the result of rejecting the boundary seed pixels with values lower than this threshold. We observe that the remaining seed pixels indeed belong to locations where the boundary is clear.

We note, finally, that the introduction of the ROI does not bias the final results of the method.

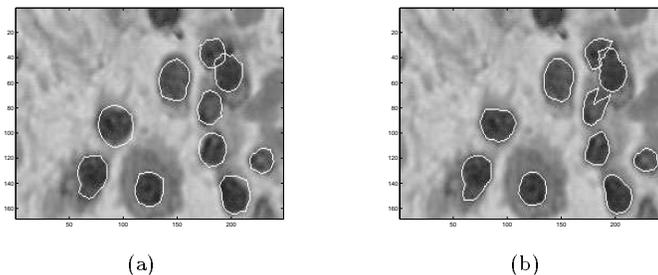


Figure 9: *The result of the optimisation for values of α (a) 1, and (b) 2.*

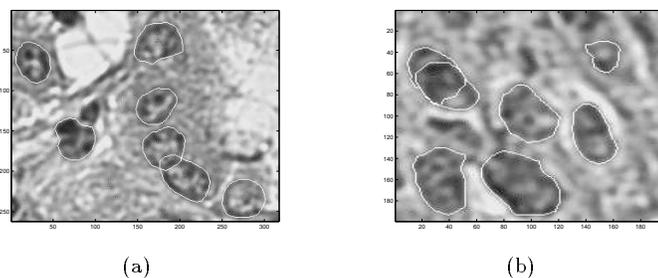


Figure 10: *Segmentation of two test image of (a) medium and (b) high complexity.*

However, it renders the algorithm much faster. In tests we performed, depending on the image, the algorithm with the use of the ROI ran from two to seven times faster than without the use of the ROI. In the particular test image of Figure 1, the number of floating point operations was four times less in the case where the ROI was used. This was due to the fact that the CHT loop (presented in section 2.3) is executed for only a portion of the image. As far as the optimisation is concerned, the ROI does not have any effect, as the edges excluded from the ROI are characterised by grey-level and gradient values that generate very low p_{gr} and p_e values, and hence lead to very low likelihoods. One further benefit from the inclusion of the ROI is the fact that the number of CHT local maxima is significantly reduced and restrained in the ROI, so that the optimization procedure is only performed for the CHT maxima included in the ROI.

Our method consists of two distinct parts: one for the detection of possible nuclear locations via the application of the Compact Hough Transform and one for the boundary optimisation via the likelihood maximisation. These two sub-methods are proposed in this work as alternative to other, already existing methods.

As far as the detection of possible nuclear locations is concerned, one could use simpler methods than the Compact Hough transform, which is responsible for the main load of computations in our proposed algorithm. The local minima of a strongly smoothed image could be used as starting points of the boundary search instead. Again, the strength of the smoothing should, somehow, be determined. Strong image smoothing may cause the merging of neighbouring nuclei and the subsequent loss of image information. The Hough transform approach uses all the edge information which is not lost by the initial noise smoothing. Our approach proposes a location for a nucleus based on all the edge information around this location and, therefore, even with incomplete or unclear boundaries, or even if the stain used is does not give the clearest view of the nucleus possible, it will still give a good estimate for the position of the nuclei. Indeed, this is one of the main advantages in using the Hough transform.

Regarding the optimization procedure, there have been other methods proposed which aim at

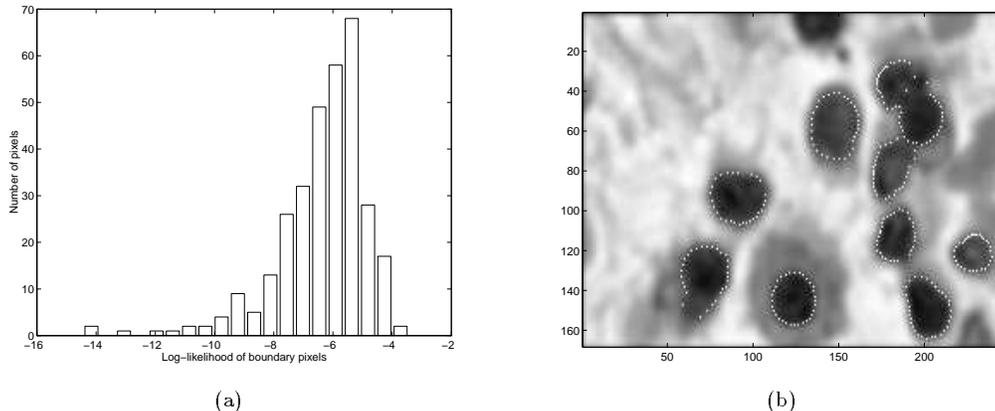


Figure 11: (a) Histogram of the log-likelihood for all the boundary pixels in the image. (b) Image with boundary pixels with likelihoods greater than the threshold of $e^{-8.66}$

minimising an energy functional [10], or a statistical error [18]. The innovation of the boundary optimization approach we propose is that it has been formulated in the framework of a set of likelihood operators. The likelihood function we introduce takes into consideration both gradient and grey-level information and penalizes abrupt transitions along the nuclear boundary. Noise introduced by the staining and scanning procedures is locally smoothed at the preprocessing stage and, therefore, its effect is reduced. Compared to [10] and [18] our method uses simpler and fewer measures for the formulation of the likelihood function, making it computationally more attractive. The problem that local likelihood maxima may be reached, which may not represent the optimal border is still existent and increases with the image complexity. This is a problem found in all existing algorithms and may be partially overcome by subsequent re-optimization excluding pixels of low likelihood.

The concept of seeking a boundary in radial directions from a point of origin was initially been introduced for cell nuclei detection in [6]. However, in this original application the boundary was selected among grey-level isointensity curves and the curvature characteristics were used solely for the purpose of selecting which of the existing proposed set of curves was the less irregular. The only other processing performed was a comparison of the area included within that border with the nuclear area proposed following a straightforward thresholding operation.

The samples we have studied with our algorithm were stained using a commercial H&E stain as this is used in clinical practice. Other stains, such as the Feulgen, have been used in the existing literature for research involving the segmentation of cell nuclei [19, 20]. In [21] haematoxylin stains are proposed as the stains giving the best image segmentation results. In a dedicated experiment, various stains may be used to achieve more intense staining of the nucleus compared to the rest of the cell, possibly leading to better segmentation results.

4 Conclusions

We have presented a multistage nuclear segmentation algorithm which is highly applicable to light microscope images of stained tissue sections. Our method consists of two independent parts:

1. The compact Hough transform application, used to establish the possible positions of nuclei in the image.
2. The likelihood maximization, used to provide a maximum likelihood nuclear boundary.

Our method has produced good results even in the case of nuclei in the process of division, preparing the ground for the extraction of useful morphological measurements from them. At a later stage,

these measurements may be used to classify the nuclei as well as the entire tissue section according to degree of malignancy or other criteria, using conventional statistics or neural networks.

Acknowledgements

The authors would like to thank the “Alexander S. Onassis” Public Benefit Foundation which has partly supported the work of T. Mouroutis. We would also like to thank Mr Ghassan Alusi and Dr Andrew Gallimore for the provision of the tissue sections on which the analysis was performed and Dr Kim Parker of the Physiological Flow Studies Group of Imperial College for his advice. Finally, we wish to thank the referees of this paper for their constructive comments.

References

- [1] R. Bowman, G. Angelini, M. Ota, and B.H. Liwnicz. Grading of astrocytomas using neural networks and linear discriminant functions based on morphometric measurements collected using static image cytometry. *American Journal of Clinical Pathology*, 102:502, 1995.
- [2] H. Galera-Davidson, R. Gonzalez-Campora, A. Mora-Marin, A. Matilla-Vicente, H.E. Dytch, H. Bartels, E. Lerma-Puertas, E. Andrada-Becerra, and M. Bibbo. Cytomorphometric DNA measurements in medullary thyroid carcinoma. *Cancer*, 65:2255–2260, 1990.
- [3] J.S.J. Lee, M.S. Banister, L.C. Kuan, P.H. Bartels, and A.C. Nelson. A processing strategy for automated papanikolaou smear screening. *Analytical Quantitative Cytology and Histology*, 14(5):415–425, 1992.
- [4] R.F. Walker, P. Jackway, B. Lovell, and I.D. Longstaff. Classification of cervical cell nuclei using morphological segmentation and textural feature extraction. In *Proceedings of the 1994 Second Australian and New Zealand Conference on Intelligent Information Systems*, pages 297–301, 1994.
- [5] J.-P. Thiram and B. Marcq. Morphological feature extraction for the classification of digital images of cancerous tissues. *IEEE Transactions on Biomedical Engineering*, 43(10):1011–1020, 1996.
- [6] E. Bengtsson, O. Eriksson, J. Holmquist, B. Nordin, and B. Stenquist. High resolution segmentation of cervical cells. *Journal of Histochemistry and Cytochemistry*, 27:621–628, 1979.
- [7] R.L. Cahn, R.S. Poulsen, and G. Tossaint. Segmentation of cervical cell images. *Journal of Histochemistry and Cytochemistry*, 25:681–688, 1977.
- [8] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis and Machine Vision*. Chapman & Hall, 1993.
- [9] F. Arman and J.A. Pearce. Unsupervised classification of cell images using pyramid node linking. *IEEE Transactions on Biomedical Engineering*, 37(6):647–650, 1990.
- [10] Y.-L. Fok, J.C.K.Chan, and R. Chin. Automated analysis of nerve-cell images using active contour models. *IEEE Transactions on Medical Imaging*, 15(3):353–368, 1996.
- [11] P.K. Sahoo, S. Soltani, K.C. Wong, and Y.C. Chen. A survey of thresholding techniques. *Computer Vision, Graphics and Image Processing*, 41:233–260, 1988.
- [12] J. Sklansky. On the Hough technique for curve detection. *IEEE Transactions on Computers*, 27:923–926, 1978.
- [13] D.H. Ballard. Generalising the Hough transform to detect arbitrary shapes. *Pattern Recognition*, 13:111–122, 1981.
- [14] J. Kittler and J. Illingworth. Minimum error thresholding. *Pattern Recognition*, 19:41–47, 1986.

- [15] J.D. Foley and A. van Dam. *Fundamentals of Interactive Computer Graphics*. Addison-Wesley, 1982.
- [16] H. Freeman. On the encoding of arbitrary geometric configurations. *IRE Transactions on Electronic Computers*, 10:260–268, 1961.
- [17] J.K. Aggarwal, R.O.Duda, and A. Rosenfeld. *Computer Methods in Image Analysis*. IEEE Press, 1977.
- [18] P. Zhou and D. Pycock. Robust statistical models for cell image interpretation. *Image and Vision Computing*, 15:307–316, 1997.
- [19] A. Claeys, C. Van den Branden, and F. Roels. Automated measurements on cell nuclei in tissue sections. validation procedures. *Analytical and Quantitative Cytology and Histology*, 10:276–284, 1988.
- [20] M. Bibbo, D.H. Kim, T. Pfeiffer, H.E. Dytch, H. Galera-Davidson, and P.H. Bartels. Histometric features for the grading of prostatic carcinoma. *Analytical and Quantitative Cytology and Histology*, 13:61–68, 1991.
- [21] M.H. Deverell, J.R. Salisbury, and W.F. Whimster. Comparison of stains for image segmentation and measurement of nuclear parameters by computerised image analysis using IBAS 2000. *Pathology Reserach and Practice*, 185:555–557, 1989.