

**SEGMENTATION OF OVERLAPPING NUCLEI IN LUNG EPITHELIAL
BIOPSY SECTIONS**

By

Roger Kemp

B. Sc. (Physics) Simon Fraser University, 1991

M. Sc. (Physics) University of British Columbia, 1994

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
PHD

in

THE FACULTY OF GRADUATE STUDIES
DEPARTMENT OF PHYSICS AND ASTRONOMY

We accept this thesis as conforming
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

August 2001

© Roger Kemp, 2001

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of Physics and Astronomy
The University of British Columbia
6224 Agricultural Road
Vancouver, B.C., Canada
V6T 1Z1

Date:

2001 November 30

Abstract

This thesis describes the development of an automated image segmentation system for resolving overlapping cell nuclei in crowded scenes such as tissue biopsy section images. The system uses a succession of imaging algorithms that combine to double the number of extracted nuclei from lung epithelial section images compared to selecting free lying nuclei.

1. The system uses the distance transform and watershed transformation to analyse the shape of object clusters and split them into smaller clusters or into individual nuclei. The watershed algorithm reliably separates two overlapping ellipses provided less than 30% of either ellipse's perimeter is occluded by the other.
2. A Hough transform algorithm was created by combining the ellipse center finding routine of Yuen with the least squares ellipse fitting formula of Fitzgibbon. Fitzgibbon's formula was adapted to include a weighting for data points so that strong ellipse edges contribute more in the determination of ellipse fit parameters. The transform was tested on a set of 431 overlapping nuclei in cytological images of lung tumour cell lines grown in culture. The Hough transform was able to produce good ellipse fits for 85% of the nuclei in the set.
3. Active contour refinement is used to refine the borders of objects segmented using the Hough transform. It was applied to the cytological image set and reduced the area misfit measure between the true nuclear mask and the Hough ellipse approximations from $8 \pm 4\%$ to $4 \pm 2\%$. The final segmentation of the nuclei created borders that delineated the overlap regions between nuclear pairs. These overlap regions were then measured for cytological and histological images to determine if the mean optical density (OD) in non-overlap regions could be used to predict the

mean OD in overlap regions. It was found that the overlap regions contain 60-70% of the predicted OD. This result was used as an empirical factor in the development of an OD apportioning scheme for the reconstruction of individual nuclei from overlapping pairs based on a maximum likelihood probability model.

The complete segmentation system was used to automatically segment biopsy section images for the purpose of recovering intact nuclei for morphometric analysis. Experiments on a set of ten tissue section images revealed that an average of 55 free lying nuclear shaped objects can be extracted from typical section frames, 83 can be extracted by applying the watershed algorithm and 102 can be extracted using the complete segmentation system. An experiment was performed on a set of nine biopsy section images that had been manually segmented and given a score using a morphometric index (MI) scoring system, which categorizes image frames based on nuclear irregularity. A decision tree was created to select "valid" nuclei based on their feature values and these nuclei were then classified with the existing MI system. The automated MI scores were compared to the manual ones and agreed for three of the nine cases. The automated system was less likely to identify abnormal appearing nuclei than normal ones. This caused it to disagree with the manual MI system for images containing severe abnormalities. The lack of success in the MI experiment is due the difficulty of combining a decision tree which attempts to throw out irregular objects with a second decision tree that seeks to categorize them. Further work in the classifier design may yield an automated biopsy section analysis system which performs as well as humans.

Table of Contents

Abstract	iii
Table of Contents	iv
List of Tables	vii
List of Figures	viii
Acronyms and abbreviations	xviii
1 Introduction	1
1.1 Background	1
1.2 Format of this thesis	4
1.3 Hypotheses	5
1.4 Description of imaging system and samples	6
2 Thresholding	10
2.1 Definition	10
2.2 Applications to cytology and histology	13
2.2.1 Grey level distribution based thresholding	18
2.2.2 Compactness based thresholding	23
2.2.3 Gradient based thresholding	26
3 Analyzing mask shape	38
3.1 The distance transform	38

3.1.1	Derivation	38
3.1.2	Locating object centres with the DT	44
3.2	The watershed transformation	48
3.3	Resolving overlapping objects in digital images	54
3.3.1	Resolving circles	55
3.3.2	Resolving ellipses	61
4	Finding nuclei from grey scale information	74
4.1	Edge detection	74
4.1.1	The Sobel filter	74
4.1.2	Contrast based filtering	78
4.2	Hough transform — Finding straight lines	79
4.3	Hough transform — Finding ellipses	83
4.4	A proposed method for ellipse detection	89
4.4.1	Sharpening peaks in the accumulator	90
4.4.2	Detecting peaks in the accumulator	94
4.4.3	Least squares recovery of ellipse parameters	97
5	Refining nuclear boundaries and assigning nuclear optical density	107
5.1	Active contours	108
5.2	Testing the optical density hypothesis	119
5.3	A model for assigning ODs for overlapping objects	125
5.4	A sequential approach to breaking up nuclear clusters	131
6	Biopsy section processing experiments	135
7	Conclusions	154

A Derivation of Overlap Measures for Circles	163
B Derivation Optical Density Distribution for Overlapping Objects	166
Bibliography	168

List of Tables

4.1	Comparison of uniform versus weighted voting for detection of 431 overlapping nuclei using the ellipse Hough transform.	94
6.1	Confusion matrix for the decision tree created to select valid nuclei for morphometric analysis.	142
6.2	Results of the automated segmentation of the 10 tissue section image training set using different segmentation options. The number of valid nuclei is obtained by passing all the segmented objects through the nucleus decision tree.	145
6.3	Morphometric index values for nine tissue section images obtained through manual analysis of the images. The uncertainty values are the standard deviations of the four MI values calculated on two separate manual segmentation attempts by two technologists.	147

List of Figures

1.1	Images of overlapping nuclei in lung cell culture. A $20\mu m$ scale bar is shown in the image in the lower left	8
1.2	Two of the 235 images of lung epithelial biopsies studied in this thesis. Each section image contains a $20\mu m$ scale bar.	9
2.1	An image of a pair of nuclei (a), its histogram of grey levels (b), and the mask of the thresholded image at $T=185$ (c).	12
2.2	Global versus local thresholding for a field, (a), with non uniform illumination. The global histogram, (b), contains two background peaks with minima at $T=102$ and $T=124$. The thresholded images for $T=124$ (c), and $T=102$, (d). Local histograms (e), show the variation of the background from top to bottom. Applying local thresholds based on histograms like those in (e) provides a good final mask, (f).	14
2.3	Portion of a lung epithelial tissue section image. There are large variations in the grey level intensity of different nuclei. Fibroblasts and lymphocytes often appear as very dark objects in the image. A $20\mu m$ scale bar is shown in the top-right.	16
2.4	Grey level histogram for the image shown in figure 2.3.	17
2.5	The bimodal Gaussian fit, (a), to the grey level histogram in figure 2.4 leads to an estimate of $T = 198$. Thresholding at this level produces the segmented objects outlined in (b).	19

2.6	The ratio of the number of background pixels to the number of object pixels for 226 tissue section images (a), and a bimodal Gaussian fit demonstrating an error where both modes attempt to fit the background and ignore the object pixels, (b).	21
2.7	Local 30×30 neighborhoods for portions of the lung tissue section image (figure 2.3) and their grey level histograms.	22
2.8	The compactness, C , for several geometric figures.	24
2.9	The failure of compactness based thresholding for nuclear clusters. The polygons that surround the nuclear clusters (a) define a neighborhood for which the histogram of compactness versus threshold is shown (b). Setting the thresholds at the peak value produces objects whose masks are too large (c).	25
2.10	The histogram, (a), of the mean gradient value as a function of grey level for figure 2.4. Smoothing this graph and selecting the peak gives $T = 153$. Thresholding at this level produces the segmented objects outlined in (b).	28
2.11	The comparison of thresholds selected by the gradient weighted and mean gradient thresholding methods.	29
2.12	Portion of a tissue section image where $T_{MG} = 99$ (see peak of MG curve) while $T_{GW} = 159$	31
2.13	A comparison of grey level and gradient profiles (b,e) for an out-of-focus and an in-focus nucleus (a,d) selected from figure 2.12. The triangles in (a) and (d) show the cross sectional limits. The product of the two curves gives the local contributions (c,f) to T_{GW} for each image.	32
2.14	Portion of a tissue section image where $T_{MG} = 107$ while $T_{GW} = 171$	34
2.15	The mask, (a), for the image shown in figure 2.14 thresholded at $T = 226$, and the gradient image, (b).	36

3.1	An object mask (a) with its exact Euclidean distance map (b) and distance transform (c). The DT is shown with only three grey tones, which represent the approximate integer distances from each pixel to the background. For example, the dark rings in (c) represent successive distances of 1, 4, 7, 10... units.	39
3.2	Pixels that touch a central pixel are shown for a grid with 4-connected pixels, (a), and 8-connected pixels, b).	40
3.3	Distances between pixels a central pixel (★) for 4-connected pixels (a), 8-connected pixels (b), and using the Euclidean distance measure (c). . .	41
3.4	Masks containing local distance values for the calculation of the distance transform. (a) and (b) show the form of masks for 3×3 and 5×5 neighborhoods. (c) and (d) are the masks for d_4 and d_8 shown in figure 3.3. . .	42
3.5	Distance transforms for masks of the outlined images (a). The DTs are calculated using the city block and chessboard metrics of figure 3.4 and using the optimal 5×5 local distance scheme of equations 3.5 and 3.7. .	45
3.6	Distance transforms for the two masks which differ by only a single pixel. Panel 1 shows the original mask definition (a) and 5×5 DT (b) for the nuclear cluster from figure 3.5. Adding one extra pixel to the mask definition in 2a causes a local maximum to appear in the DT (2b) for the nucleus on the right.	47
3.7	A curve in one dimension (a) contains several local minima. Each local minimum defines a catchment basin. A drop of water that falls into a particular catchment basin will flow downhill to the local minimum. The points where catchment basins meet are called <i>watershed</i> points (marked with triangles). The watershed algorithm converts the curve into a series of labelled regions, (b), with each region marking a catchment basin. . .	49

3.8	A binary image (a), the distance transform (b) and the watershed segmentation of the DT (c) are shown.	51
3.9	The distance transform and watershed segmentation are shown for three clusters of overlapping nuclei from tissue section images. The arrows show errors in the segmentation. A $20\mu m$ scale bar is shown in the bottom image	53
3.10	The distribution of the maxima of the distance transform for: (a) the cluster of circles and ellipses (figure 3.8), (b) the three nuclear cluster images (figure 3.9).	54
3.11	Measuring the ability of the DT to resolve overlapping circles. The graph (b) shows the minimum separation required to resolve two circles of equal size. In the case where the radius is 20 pixels (c), the DT cannot resolve two circles separated by 13 pixels, but can resolve them at 14 pixels. . . .	56
3.12	The construction for calculating the minimum separation of resolvable circles.	58
3.13	The construction (a) for calculating the minimum separation of resolvable circles where $r_1 > r_2$. Distance transforms (b) for a pair of unresolvable and resolvable overlapping circles.	59
3.14	The best fit ellipses (a) for the watershed of the nuclear cluster in panel 3 of figure 3.9, and the histograms (b) of axes lengths for fit ellipses for all 3 clusters in that figure.	62
3.15	The histogram of major and minor axes lengths for the fit ellipses, (a), and radii of circle fits, (b), to the watershed segmentation of 226 lung epithelial tissue sections.	63
3.16	Possible measures of overlap for Monte Carlo simulations of watershed separation of overlapping ellipses. (a) shows the parameters of the ellipses, while (b-d) show three possible measures based on the parameters.	65

3.17	The three overlap measures for an instance of two resolvable ellipses (a) and two unresolvable ellipses(b).	67
3.18	The graphs of separability of one million randomly generated overlapping ellipse pairs segmented with the watershed versus each of the three measures of overlap. The theoretical results for circles with the same area distribution are overlaid with dotted lines.	69
3.19	A pair of overlapping circles and the watershed segmentation of the mask (a). The MP measure requires knowing OP_i but only the contact length, CL is available. (b), the application of a reformulated version of MP (MP') to the nuclear cluster from figure 3.9.	72
4.1	A cluster of overlapping nuclei (a) is shown along with its Sobel image (b), thresholded Sobel image (c), contrast image (d) and thresholded contrast image(e). The contrast operator enhances interior edges of clusters.	77
4.2	A pair of cultured HL-60 nuclei in (a) are shown along with the Sobel image (b), thresholded Sobel image (c), contrast image (d) and thresholded contrast image(e). Edge recovery fails when the image contains significant nuclear texture.	80
4.3	The Hough transform for finding lines in two dimensions. Each point in the xy plane on the left gives rise to a line in the ab parameter space on the right. The intersection of lines in the parameter space represents points that are collinear in the image space for some value of (a, b) . The intersection point, (a', b') , on the right gives the parameters of the dotted line that passes through (x_1, y_1) and (x_2, y_2) .	81

4.4	The Hough transform for a noisy line image. The two lines in the image give rise to peaks in the transform space that are clearly visible despite the noise (which constitutes 80% of the pixels in the original image). . .	83
4.5	The Yuen method for finding the center of an ellipse. P_1 and P_2 are two points on the edge of an ellipse. Their tangents intersect at T and the midpoint of P_1P_2 is M . Due to the geometry of ellipses, the extension of the line TM passes through the ellipse center, C	86
4.6	The transformation from image space, (a), to vote space (b). The distribution of bright points in (b) represents likely positions of an ellipse center. The voting peaks are shown in (c). Eight of the ten nuclei are located immediately with the Hough transform.	88
4.7	Demonstration of uniform (a,b) and weighted (c,d) accumulator voting strategies for the location of the center of an ellipse.	91
4.8	Two accumulators for two ellipse images, (a), are shown for uniform, (b), and weighted, (c), Hough voting schemes.	93
4.9	An image of overlapping nuclei, (a), the Hough vote space, (b), and the smoothing of the vote space using a Gaussian smoothing operator (c). The vote spaces plotted as surfaces (d,e) show seven and two voting peaks respectively.	95
4.10	The voting edge accumulators, B^i for: (a) the cluster of three nuclei in the top of figure 4.6, and (b) the two ellipse centres recovered from figure 4.9c. The brightness of each pixel corresponds to the number of times it voted (in conjunction with another edge) for the respective ellipse center. . . .	96

4.11	Fitzgibbon's B2AC algorithm for fitting ellipses using edge data. The scattered data used to generate the ellipse parameters are shown with \times 's. The points were selected from the same ellipse and subject to various noise levels. The solid lines show the ellipse fit and the dotted lines show the true ellipse in the absence of noise.	100
4.12	The final result for the ellipse Hough transform applied to: (a) the nuclear clusters in figure 4.6a, and (b) the pair of nuclei in figure 4.9a.	102
4.13	Definition of the measure of difference between the Hough ellipse fit and the true object, (a), and examples showing increasing error levels between the fit ellipses and true objects, (b).	103
4.14	The Hough transform applied to some the cultured cell nuclear images. The errors were calculated by comparing the fit ellipses to the manually segmented masks shown in figure 1.1. The (r) and (g) designations refer to the red and green ellipses.	105
4.15	The cumulative number of nuclei detected versus the percentage area misfit for the 431 nuclei in the cultured cell data set.	106
5.1	Fitting a snake to image data. The contour is initialized either randomly or to an estimate of the image data (a). Optimizing the snake results in one that fits the data subject to the relative sizes of elasticity constraint $\alpha(s)$, (b), and the stiffness constraint $\beta(s)$, (c).	109
5.2	A continuous snake (a), and the planar curve representation (b) as a series of vertices on a predefined grid. The vertices may connected by straight line segments (solid segments) or a piecewise spline (dotted segments). .	112

5.3	Transformation of a nuclear image, (a) using an circular sampling grid (b). The circular transform of the image and the vertical direction contrast based gradient are shown in (c). The nuclear border is a smooth curve in the transform space (c) Transforming back into the image space gives the segmented nucleus (d).	113
5.4	Transformation of a nuclear image of figure 5.3 using an elliptical sampling grid (a). The transformed image and vertical direction contrast based gradient are shown in (b). Using the elliptical transformation the nuclear border is nearly a straight line in the transform space, (c). The final segmentation of the nucleus is shown in (d).	115
5.5	Active contour segmentation of an overlapping nucleus based on the Hough transform ellipse fit to the grey scale image.	117
5.6	Active contour segmentation of two overlapping nuclei based on the wa- tershed segmentation of the thresholded image.	118
5.7	Histograms of the area misfit measure (Err_A) for the cultured nuclei (fig- ure 1.1) prior to active contour border refinement, (a), and after border refinement, (b). The refined borders are shown for five of the images (c).	120
5.8	Experiments in measuring the mean optical density in overlap regions versus non-overlap regions. The non-overlap regions are designated A and B and the overlap region is designated C, (a). The graphs of measured mean OD in overlap regions versus the sum for non-overlap regions are shown for 188 overlap regions in cytological images, (b), and for 204 regions in biopsy sections, (c). The dotted lines represent the model prediction that $\overline{OD}_A + \overline{OD}_B = \overline{OD}_C$	122
5.9	Fifteen images of a nucleus (a) captured at $1\mu m$ increments and the graph of total optical density versus the focal plane position, (b).	123

5.10	Images and manually segmented masks of overlapping nuclei from lung tissue section images. The light grey regions in the masks denote the regions of overlap.	124
5.11	Demonstration of the image reconstruction method. The overlapping nuclei (a) are segmented (b), and the optical density histograms are calculated for the non-overlapping regions (c). The model assumes that pixel OD values are normally distributed and performs a Gaussian approximation to the OD histograms. The OD_C are then apportioned using this model giving the reconstructed nuclei (d).	127
5.12	Grey level histograms for the three regions of the nuclear image in figure 5.11b.	129
5.13	Reconstructed nuclei from figure 5.11 with OD assignments made according to the joint probability density (equation 5.7) rather than using the ML estimates for that equation.	130
5.14	Outline of the sequential procedure for removing nuclei from clumps. A small portion of the border of the clump remainder is refined after each nucleus is removed to allow more nuclei to be retrieved.	133
6.1	Outline of the decision tree steps in the automated MI scoring system. .	139
6.2	Segmenting a grey level image, (a), will lead to different final segmentations depending on the thresholding options used (b-d).	141
6.3	Automated morphometric index score calculations for test images 1, 2 and 3. The solid and dashed lines show the mean and standard deviation for the manual analysis of these images.	149

6.4	Automated morphometric index score calculations for test images 4, 5 and 6. The solid and dashed lines show the mean and standard deviation for the manual analysis of these images.	150
6.5	Automated morphometric index score calculations for test images 7, 8 and 9. The solid and dashed lines show the mean and standard deviation for the manual analysis of these images.	151
6.6	Comparison of the automatic morphometric index calculations to those obtained through manual analysis. The MI score for all 20 different segmentation methods are plotted for each image.	152
A.1	Construction for deriving the measures of overlap for two circles of different radii.	164

Acronyms and abbreviations

AC	Active contours
DT	Distance transform
EDM	Euclidean distance map
GHT	Generalized Hough transform
GW	Gradient weighted (threshold)
HT	Hough transform
MA	Ellipse overlap area measure
MG	Mean gradient (threshold)
MI	Morphometric index
ML	Maximum likelihood
Ms	Ellipse separation measure
MP	Ellipse perimeter occlusion measure
NA	Numerical aperture
OD	Optical density
UEP	Ultimate erosion point
WS	Watershed (algorithm)

Chapter 1

Introduction

1.1 Background

Lung cancer is the leading cause of cancer death in Canada [51] and other developed countries. It accounts for nearly one third of all cancer deaths for men and more than one quarter for women. Despite the advances made in cancer treatment, the five-year survival rate for all people diagnosed with lung cancer is less than 15% [4]. This is largely due to the fact that approximately 85% of people diagnosed with lung carcinoma (in the United States) already have advanced stages of the disease [61], where the therapeutic options are mostly palliative.

The patient prognosis for lung cancer is strongly dependent on how early it is detected. Patients diagnosed with Stage 0 or Stage I disease have a five-year survival of $> 90\%$ and 70-80% respectively. If the diagnosis is made later in the disease's progression, the prognosis is much poorer. Stages II through IV have five year survival rates that range from 40% to less than 5%. These statistics demonstrate a great benefit to discovering the disease early, where more treatment options exist. Early detection of lung cancer can dramatically increase a patient's chance of survival.

Lung cancer is a form of epithelial cancer which accounts for 92% of all cancers detected [71]. The development of lung cancer is a complex and poorly understood phenomenon. Current models suggest that most lung cancers are the end result of the accumulation of genetic damage to the nuclei of cells located in the central airways of the

lung [5, 58]. The genetic damage causes changes in the conformation of the DNA within epithelial cell nuclei that cause them to behave abnormally. They may become immortal as well as proliferate giving rise to a premalignant lesion. The visual appearance of an abnormal arrangement of epithelial cells is termed dysplasia (or neoplasia). Premalignant lesions are typically diagnosed as mild, moderate or severe dysplasia. These categories represent different grades of intraepithelial neoplasia.

Classification of preinvasive neoplasia into such categories is a difficult task. It is common for there to be variations in the grading of a sample by different pathologists. For example, in a study of interobserver variability of grading lung epithelial biopsies by clinical pathologists, the grading categories matched for only 34.8% of the samples, and differed by one or less categories for only 66% of the samples [38]. This suggests that the development of methods to quantitate the analysis of tissue section images may both improve the reproducibility of the diagnosis of lesions and give insights into objective criteria that can aid pathologists in understanding the pathogenesis of cancer.

Precancer diagnoses can be made through observing visual changes in epithelial nuclei. Abnormal cellular nuclei tend to have features that distinguish them from normal nuclei such as being larger, darker, more irregularly shaped, etc. These changes can be measured in cytological preparations of the epithelial cells, where they are extracted, mounted on a microscope slide, stained and observed using a light microscope. Cytometric measurements may be made on the nucleus and, in some cases, the surrounding cellular cytoplasm. The most common cytometric measurements fall into three classes:

- morphometric features relating to nuclear shape,
- densitometric features relating to nuclear optical density (darkness), and
- textural features relating to organization of material in the nucleus.

Precancer diagnoses are also made through analysis of histological preparations of the epithelial tissue. In this case, the tissue is obtained during bronchoscopy as a bronchial biopsy, sliced into sections typically a few micrometres thick, mounted, stained and observed under a light microscope. As well as being able to study the cytometric features of individual nuclei, the organization of cells within the epithelial layer can also be observed.

Efforts to quantitate and automate the diagnosis of precancerous lesions through the analysis of cytologic and histologic images have been made for more than thirty years. [36]. The transition from biological sample to quantitative diagnosis consists of four major steps:

1. Digital images are obtained from the sample.
2. The image scene, consisting of one or more nuclei is *segmented*. Segmentation is the process of dividing an image scene containing pixels (picture elements) into regions or structural units that represent the objects of interest.
3. Features are extracted from the segmented images and multivariate analysis is performed on the features.
4. Classification functions obtained through multivariate analysis are applied to the image dataset to give a diagnosis for the sample.

The second step, segmenting images, is the most difficult stage for quantitative image analysis. For histological samples, there are no completely automatic systems for segmenting nuclei from these images. These samples contain a wide variety of cell types, large amounts of debris and contain such a great deal of nuclear overlap that, invariably, human intervention is required. Free lying nuclei can usually be recognized automatically in these images, but in order to segment more nuclei manual analysis is required.

This thesis focuses on the automation of the segmentation of images containing touching and overlapping objects. The specific problem addressed is how to separate nuclei from each other and the background. Having an accurate automated segmentation system would not only speed up the processing of tissue section images, but it would also allow new questions to be posed in the area of histological analysis. Currently, the manual segmentation step requires 45 minutes per image or longer. This limits the amount of data that can be collected and therefore the nature of the studies that may be performed. The ability to automatically process tissue section images will allow us to expand the scope of histological analysis and give a better idea of the precision of current histological analysis methods.

1.2 Format of this thesis

This thesis is divided into six chapters. Section 1.3 explains the hypotheses that are addressed. Section 1.4 describes the imaging equipment used to obtain the cytological and histological images that were used in the experiments performed for this thesis.

Chapters 2 to 5 describe the development of the automated nuclear image segmentation system and the experiments performed on the individual components of the system. The segmentation system starts with grey level images and generates image masks by thresholding (chapter 2). The image masks are analysed and the shape of the masks is used to initially break large clusters of objects into smaller clusters (chapter 3). The image grey level information is analysed using elliptical Hough transforms to separate individual nuclei from small clusters (chapter 4). The borders of individual nuclei are then optimized using active contour techniques and optical density in overlapping regions was apportioned to obtain the final segmentation of the nuclei (chapter 5).

Chapter 6 describes the application of the segmentation system to the analysis of

lung epithelial biopsy sections. The system was used to isolate individual nuclei from the sections. Current practice uses manual effort to select these nuclei from the images. The results of using the automated system were compared to those obtained manually. Finally, chapter 7 gives a summary to the hypotheses and states the conclusions to the thesis.

1.3 Hypotheses

This thesis examines three hypotheses:

1. Given that lung epithelial nuclei appear elliptical in shape, it is possible to design an imaging system to locate them in cytological and histological images by searching for ellipses in these images.
2. Nuclear optical densities in cytological and histological absorbance microscopy images are additive.
3. It is possible to recover nuclear images from lung biopsy sections in a completely automated fashion and generate equivalent morphometric indices to those determined by manual analysis of the images.

Chapters 2-4 and section 5.1 address the first hypothesis. Section 5.2 details experiments concerning the second hypothesis for measuring nuclear optical density while section 5.3 describes a probabilistic model for reconstructing optical density for overlapping objects. Chapter 6, which describes the morphometric analysis of biopsy sections, addresses the third hypothesis

1.4 Description of imaging system and samples

The cytological and histological images in this thesis were captured using an imaging system designed by the Cancer Imaging Department of the BC Cancer Agency and Xillix Inc [24]. The system uses a light microscope with an objective lens with magnification, $M = 20\times$, and numerical aperture, $NA=0.75$. Samples are illuminated using light that passes through a 70nm narrow bandpass filter centred at $\lambda=600\text{nm}$. Images are captured using a CCD array placed in the primary image plane of the objective. The CCD is made up of an array of 1280×1024 elements $6.8\mu\text{m}$ on edge with a 100% array fill factor. Images are digitised into 8-bit grey scale images so that each image pixel has 256 possible grey level values from black=0 to white=255.

With $20\times$ magnification and CCD array element size of $e = 6.8\mu\text{m}$, the resolution of the device is $0.34\mu\text{m}$ per pixel. This means, for example, that an object that is $34\mu\text{m}$ long will produce an image that is 100 pixels in length. The *depth of field*, d_f , of the imaging system is a measure of the depth in specimen space that appears in to be in focus within the image without readjustment of the microscope [31]. It is given by the equation

$$d_f = \frac{\lambda n}{NA^2} + \frac{ne}{MNA} \quad (1.1)$$

where n is the refractive index of the medium ($n_{\text{air}} = 1.0$). Substituting the system parameters into equation 1.1 gives $d_f = 1.52\mu\text{m}$ for this device.

The cytologic images used in the experiments in this thesis consisted of 212 images of touching and overlapping nuclei from a lung tumour cell line grown in culture. The images were all classified by a cytotechnologist as containing multiple cells. The cells were stained with a Feulgen-Thionin stain which binds to the DNA in the nucleus and leaves the surrounding cytoplasm invisible. Most images contained two nuclei and some contained three. The images contained 431 nuclei in total. Each image pixel was manually

categorized as belonging to one nucleus or to both. Figure 1.1 shows a gallery of 14 of the images. Beside each image is a *mask*, which shows the pixels that comprise the nuclei. The darker grey regions of the masks show the portions of the nuclei that are non-overlapping and the overlap regions are shown with a lighter grey tone.

Cell culture images were used because they tend to be free from debris. This makes the exact nuclear borders easier to detect, which makes it easier to do a precise manual segmentation of the images. The nuclei of cells from a lung tumour line contain abnormal amounts of DNA and therefore tend to be larger and less elongated than normal cell nuclei. Large round objects are easier for algorithms to automatically recognize, so these images provided a starting point for the testing the algorithms used to process the images.

The set of histologic images used in this thesis consisted of 235 images of biopsied lung epithelial tissue. The epithelial biopsies were cut into $7\mu\text{m}$ sections, Feulgen-Thionin stained and then imaged. The images were analysed and most diagnosed as being normal or containing mild atypia while some were diagnosed as containing various stages of precancerous lesions from mild dysplasia to carcinoma *in situ*.

The image set was split into two groups. The first group consisted of 226 images that were used in the development of the segmentation system described in this thesis (Chapters 2–5). The second group consisted of nine images that were manually segmented by histotechnologists in a study of the use of nuclear morphometry in the development of a reproducible method to quantitate the grading of bronchial neoplastic lesions. Figure 1.2 shows examples of two of the nine biopsy sections used in the morphometric analysis study. Each image contains several hundred nuclei, many of which are touching or overlapping with others.




Figure 1.1: Images of overlapping nuclei in lung cell culture. A $20\mu m$ scale bar is shown in the image in the lower left



Figure 1.2: Two of the 235 images of lung epithelial biopsies studied in this thesis. Each section image contains a $20\mu m$ scale bar.

Chapter 2

Thresholding

2.1 Definition

An image is created, whether by photography or electronic imaging, so that the objects of interest stand out from the background. When analysing the image scene it is necessary to separate the regions that we are interested in (foreground) from the background. This is most commonly done by noting the range of brightness values that are part of the foreground versus those that are part of the background, and, setting a limit, called a *threshold*, between them. Thresholding is probably the single most important and widely used procedure in image segmentation. Threshold selection is important because most of the subsequent steps in image analysis procedures depend on having a good initial threshold.

The thresholding procedure converts a grey level image $f(x, y)$ into a binary image $b(x, y)$, where (for dark objects against a light background) b is defined as

$$b(x, y) = \begin{cases} 1 & \text{if } f(x, y) \leq T \\ 0 & \text{if } f(x, y) > T \end{cases}$$

and T is the threshold level. Defined in this manner, the threshold divides the image into regions of object and non-object. An alternative scheme is to define multiple threshold levels, T_i , and categorize pixels into several sets based on mutually exclusive grey level ranges. Although this may have useful application in attempts to segment the darker overlap regions of overlapping objects, this approach is not explored further in this thesis.

The threshold function, T , can be defined in one of three different ways [69]

1. *Global* threshold: $T = T\{f(x, y)\}$

T depends on image data and the same T value is applied to all pixels in the image.

2. *Local* threshold: $T = T\{f(x, y), N(x, y)\}$

T depends on image data and on a function N whose value depends on a local neighborhood around (x, y) .

3. *Dynamic* or *adaptive* threshold: $T = T\{f(x, y), N(x, y), x, y\}$

T depends additionally on some desired property of the image at the co-ordinate (x, y) , such as maintaining the “roundness” of segmented objects or the continuity of object borders in b , etc.

The most common technique for finding global threshold is by analysing the histogram of image grey level values. Histogram analysis relies on the fact that ideal image histograms are bimodal and the threshold, T , should be set somewhere in the valley between the modes. Figure 2.1a shows an image of a pair of nuclei and the histogram of this image (figure 2.1b). The histogram is clearly bimodal, with a sharp peak due to the lighter background pixels that surround the nuclei and a broader peak due to the nuclei. The standard histogram minimum technique [69] sets T at the location where the minimum of the region between the peaks occurs. The histogram is usually smoothed to eliminate false peaks before the minimum is determined. Figure 2.1c shows the image mask created by thresholding the nuclear image at $T=185$ which is the location of the histogram minimum.

Local thresholding is based on treating the image as a series of neighborhoods. A property, $N(x, y)$, is calculated for each neighborhood and a local threshold is determined. The kinds of functions calculated for the neighborhoods are often the same ones calculated

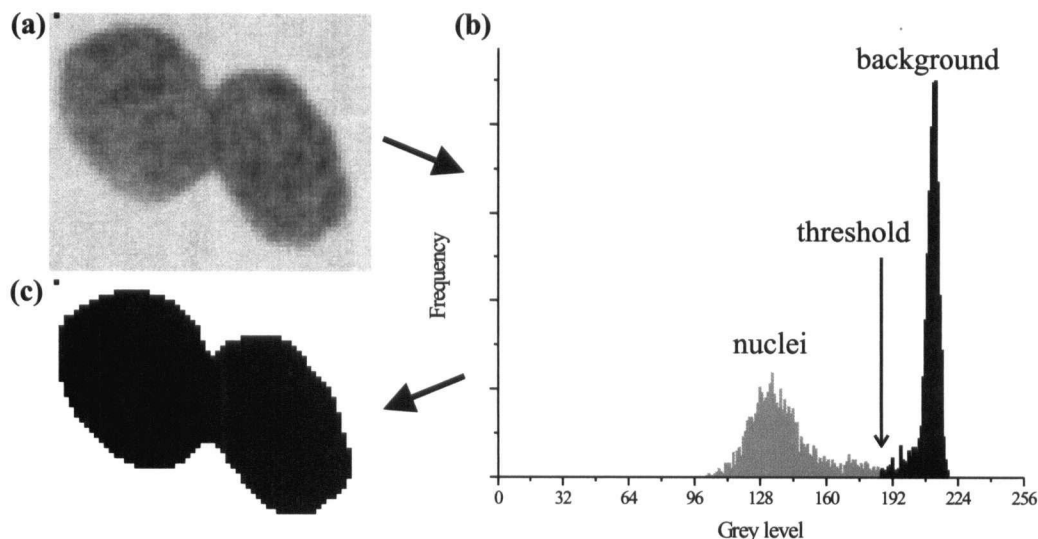


Figure 2.1: An image of a pair of nuclei (a), its histogram of grey levels (b), and the mask of the thresholded image at $T=185$ (c).

for global histograms. By calculating them on smaller regions a threshold may be selected that is more responsive to the nature of the image in that region.

Calculating a local threshold for each neighborhood produces a field of thresholds across the whole image. A problem that can occur is that when two neighboring regions have different thresholds an object at the border of the two neighborhoods is partly thresholded at one T value and partly at a second. This may produce an unsatisfactory object mask that displays banding effects due to the abrupt transitions in threshold value.

This problem can be overcome by applying some sort of smoothing procedure to the field of local thresholds. For example, the thresholding scheme developed by Chow and Kaneko [18] divided the images he used into a grid of 64×64 pixel regions and calculated a local threshold for each region. Then the grid of thresholds was interpolated on a pixel by pixel basis so that each pixel in the image was assigned a threshold. This scheme allowed the threshold to gradually vary across the image. Technically, this is an instance

of a dynamic thresholding [25] scheme since the threshold function, T , is an explicit function of (x, y) .

Local thresholds are typically used when variations in illumination make it impossible to find a suitable global threshold. Figure 2.2a shows an image of a field of nuclei with non-uniform illumination. There is a clear decrease in the illumination from the top of the image towards the bottom. The histogram of the image, (figure 2.2b), contains two distinct background peaks. The histogram minimum technique suggests two possible thresholds, at $T=102$ and $T=124$. Figures 2.2d and 2.2e show the results of thresholding the image at these levels. The upper threshold produces a reasonable mask for the top half of the image but completely fails for the bottom half. Conversely, the lower threshold is suitable for the bottom portion of the image but not for the top portion.

Figure 2.2e shows the histograms for three regions marked with rectangular boxes. Comparing the histograms, we can see how the location of the background peak shifts to lower values as we sample closer to the bottom of the image. We can use these histograms to construct varying thresholds that are suitable to each region. The image is divided into a series of neighborhoods and a local threshold is calculated based on each neighborhood histogram. Figure 2.2f shows the mask generated when the histogram minimum method is applied to neighborhoods consisting of 16 rows of pixels at a time. A total of 45 different local thresholds are determined and each applied to its respective neighborhood. The final mask is a good approximation to the nuclei that are present in the image.

2.2 Applications to cytology and histology

Many thresholding schemes have been studied in medical image analysis [23, 56] and for finding nuclei in cytological cell preparations in particular. The most common approach applied to cytology is to select thresholds based on the shape of the grey level

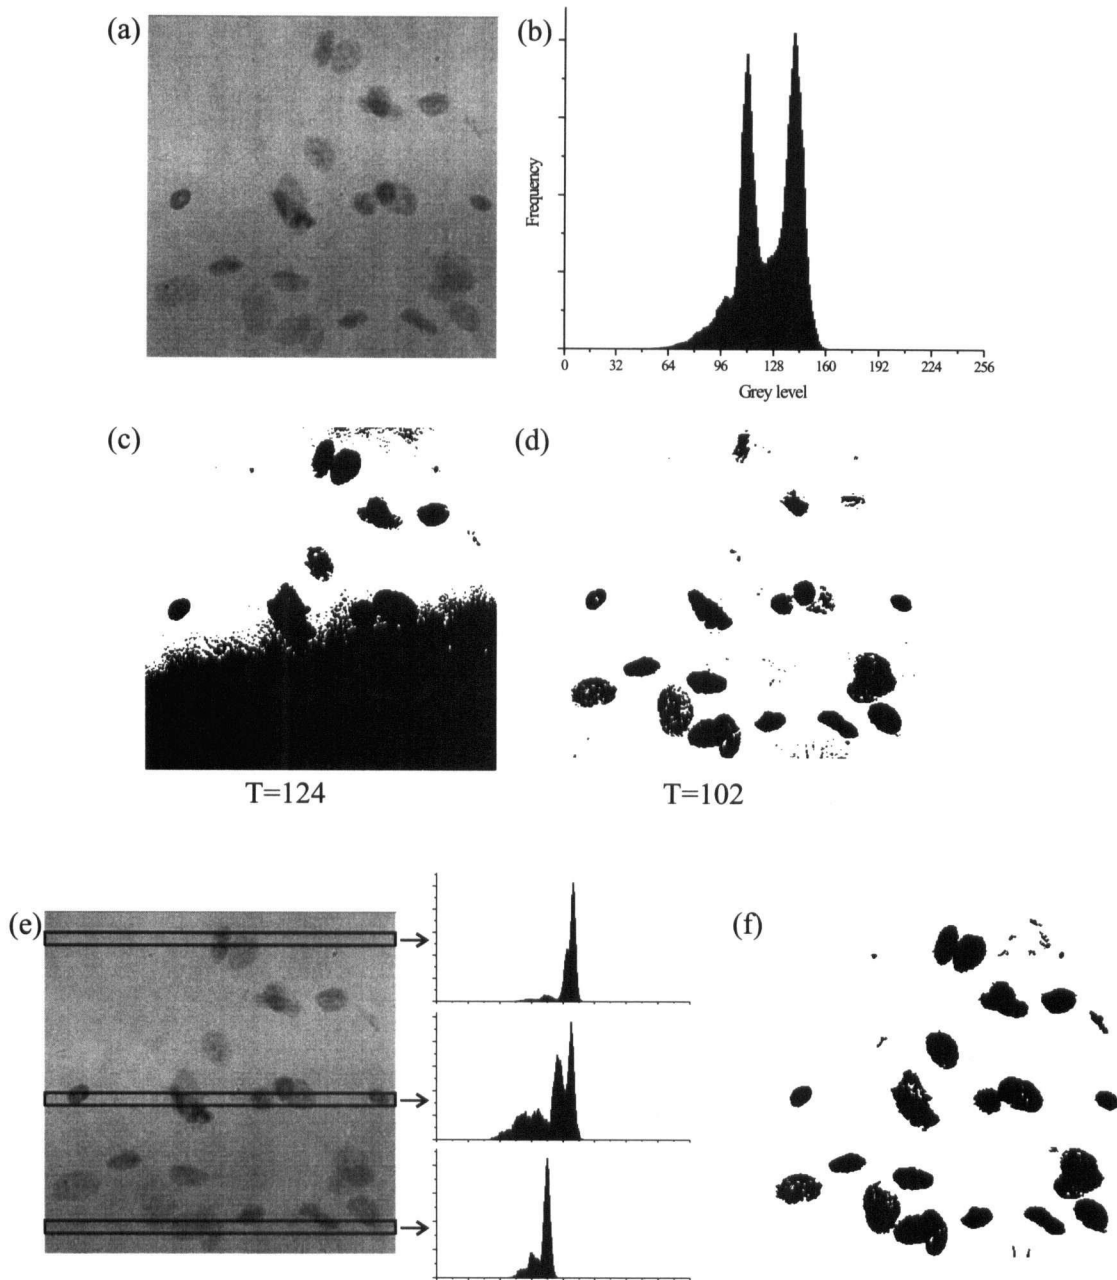


Figure 2.2: Global versus local thresholding for a field, (a), with non uniform illumination. The global histogram, (b), contains two background peaks with minima at $T=102$ and $T=124$. The thresholded images for $T=124$ (c), and $T=102$, (d). Local histograms (e), show the variation of the background from top to bottom. Applying local thresholds based on histograms like those in (e) provides a good final mask, (f).

histogram [11, 14, 28, 41, 42, 59]. Some of the thresholding methods described in these papers and others described below were studied in order to find an effective one for the biopsy section samples used in this work.

Since thresholding is an important component of a segmentation system the initial work in this thesis focused on developing a thresholding scheme suitable for the types of images under consideration. The cytological and histological images studied in this thesis were all Feulgen-Thionin stained. This stain is a quantitative DNA stain that only stains the cell nucleus. For example, for the cytological image in figure 2.1, only the nuclei are seen and the cellular cytoplasm that surrounds each nucleus is suppressed. This should make the task of thresholding these images easier than images of specimens stained using other techniques, such as Papanicolaou staining (which stains the nucleus with one stain and the cytoplasm with another).

Figure 2.3 shows a portion of a Thionin stained lung epithelial tissue section image. The region shown is 443×368 pixels while a full frame is 1280×1024 pixels. The image contains roughly 120 nuclei or nuclear fragments. The nuclei have a wide range of grey level distributions due to being sliced in the biopsy preparation and also because some nuclei are not in the image focal plane. The grey level histograms for two nuclei are shown at the top. For example, the lighter object that is highlighted has a mean grey value of 132 while the darker has a mean grey value of 74. Their grey level distributions hardly overlap. A good threshold for the darker object of $T = 124$ excludes more than half of the pixels of the lighter nucleus.

At the base of the image is the subepithelium, which consists of fibrous material, collagen, etc. This material does not stain with the Feulgen-Thionin reaction, but it contains density variations that cause light to refract as it passes through. The visual effect is to create a textured region that can influence the threshold selected by global thresholding routines. Also present in the subepithelium are a significant number of




Figure 2.3: Portion of a lung epithelial tissue section image. There are large variations in the grey level intensity of different nuclei. Fibroblasts and lymphocytes often appear as very dark objects in the image. A $20\mu m$ scale bar is shown in the top-right.




Figure 2.4: Grey level histogram for the image shown in figure 2.3.

lymphocytes and fibroblasts. These objects, particularly the lymphocytes, tend to be darker than the epithelial nuclei. A few of the nuclei in the image are out of focus. This is common when imaging thicker biopsy sections. They pose a problem for segmentation systems because they display a gradual intensity change from object to background rather than a sharp border.

Figure 2.4 shows the grey level histogram for the whole frame. The background is very uniform, having a standard deviation of only four grey levels. The distribution of nuclear pixels is very wide, extending from roughly 35 grey levels to 175. Between the nuclei and background is a region of grey levels that come from two sources: subepithelial texture and pixels that are very close to nuclear borders. All of the pixels darker than 35 grey levels correspond to the lymphocytes and to regions of nuclear overlap.

The flatness of the nuclear grey level distribution and the presence of a transition zone between nuclei and background means that this kind of image is not amenable to thresholding techniques that look for minima in the global histogram. There is no clear

distinction between the nuclear and background regions, so the histogram minimum technique is not likely to select a good threshold. In fact, the selection of any global threshold requires a trade-off between including or excluding lighter objects versus generating enlarged or optimal borders for dark objects.

2.2.1 Grey level distribution based thresholding

The separation of the histogram in figure 2.4 into different grey level regions suggests that it may be possible to define a probability distribution approach to selecting a threshold. In this approach, we assign a probability model to the distribution of object pixels and background pixels and perform a numerical fit to model. The threshold T is then selected to minimize the probability of misassigning pixels. This is the basis of the optimal thresholding method described by Gonzalez [25]. This method assumes that the pixel values are drawn from two Gaussian populations—one representing the object pixels, with mean μ_1 and noise level σ_1 , and the other representing the background, with parameters $\{\mu_2, \sigma_2\}$. Gonzalez calculates the optimal threshold based on these quantities and the proportions of nuclear and background pixels. This approach requires fitting a bimodal Gaussian distribution to the grey level histogram.

The bimodal Gaussian fit for the histogram of figure 2.4 is shown in figure 2.5a. The Gaussian model does a reasonable job of fitting the background grey level distribution, which although not exactly symmetric, is within 15% of the experimental value at each point. However, since the distribution of nuclear pixels is not Gaussian, and is certainly not symmetric, the Gaussian fit does a poor job of approximating their distribution. This wide grey level spread leads to an estimate of $\sigma_2 = 30$ pixels. The optimal Gaussian threshold for this image, using the method of Gonzalez, is $T = 198$.

This is an extremely high threshold for such an image. The high threshold from the Gaussian model results because the background distribution has a very small standard

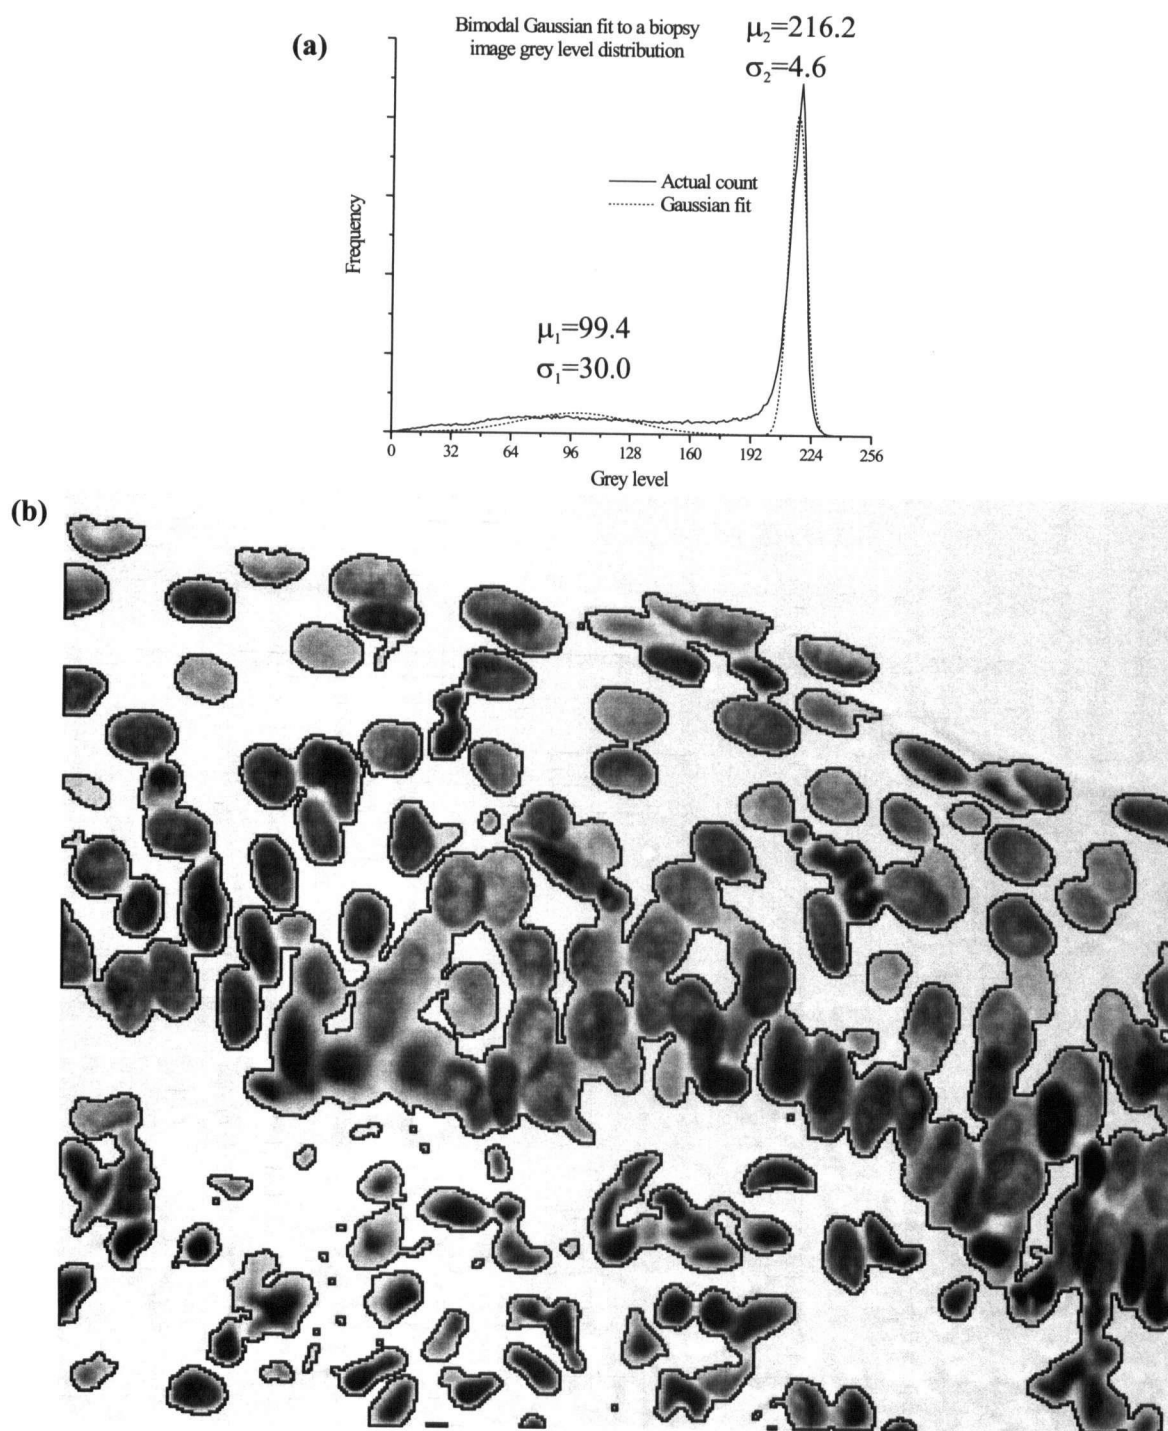


Figure 2.5: The bimodal Gaussian fit, (a), to the grey level histogram in figure 2.4 leads to an estimate of $T = 198$. Thresholding at this level produces the segmented objects outlined in (b).

deviation of $\sigma_2 = 4.6$ pixels, while the nuclear pixels have a standard deviation 6.5 times larger. The method sets T at the grey level which minimizes the misfit according bimodal distribution assumption. The large σ_1 value causes the nuclear Gaussian approximation to be non-zero for very large grey level values. This means that to minimize the probability of misfit, the threshold level is forced upwards, towards the background grey level peak.

Figure 2.5b shows the outline of the segmented objects for $T = 195$. Most of the segmented objects are surrounded by a ring of background pixels that have been included in the mask definition. Segmenting the objects at this threshold distorts the shape of the mask and prevents the splitting of clusters based on nuclear shape, which will be discussed in section 3.2. For this reason, global thresholding using this method was rejected.

A second problem with global thresholding based on a bimodal density model occurs when the number of background pixels greatly exceeds the number of nuclear pixels. When this occurs, the fitting algorithms can reduce the misfit criterion by attempting to fit the background with two distributions and ignoring the nuclear pixels altogether. The ratio of number of background pixels to object pixels for the image shown in figure 2.3 is 2.4:1, which is a relatively small ratio for tissue section samples being studied. Figure 2.6a shows a histogram of the ratio of number of background pixels to nuclear pixels for the set of 226 lung epithelial tissue section images. The ratios range from a minimum of 2:1 for a dense image to a maximum of 49:1 for an image containing only a few dozen nuclei. The mean ratio for the all the tissue sections was 8.5:1.

Figure 2.6b shows the bimodal Gaussian fit to one of the images where the ratio of the number of background to nuclear pixels was 12:1. The fitting procedure minimizes the misfit between the model and the actual data by ignoring the object pixels and fitting the region centred at grey level 213. This leads to an unreasonable estimate of $T = 216$,

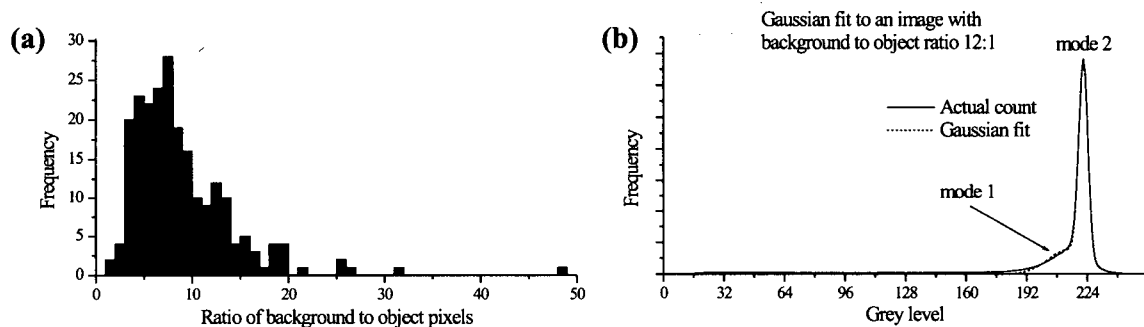


Figure 2.6: The ratio of the number of background pixels to the number of object pixels for 226 tissue section images (a), and a bimodal Gaussian fit demonstrating an error where both modes attempt to fit the background and ignore the object pixels, (b).

a value that is high enough to be well into the background. Since nearly a quarter of the images have background to object ratios higher than this image, it is clear that the bimodal fitting scheme cannot be applied on a global basis.

It can be argued that since the darkness of desirable nuclei varies significantly, applying the bimodal Gaussian model on a local basis may give better results. This is the approach used in the local thresholding method of Chow and Kaneko mentioned earlier. Their method performed bimodal Gaussian fits over a 7×7 mesh of 256 pixel neighborhoods in order to find optimal local thresholds.

Local thresholding schemes require setting the dimension of a window over which the local histograms will be calculated. The window size must be large enough so that valid background and nuclear estimates can be calculated from the histogram. Since the nuclei are 20-30 pixels in diameter, this suggests that a window size of at least 30×30 is required.

Figure 2.7 shows two portions from the tissue section image of figure 2.3. A 30×30 neighborhood is outlined in each portion, and its grey level histogram is shown. Each of

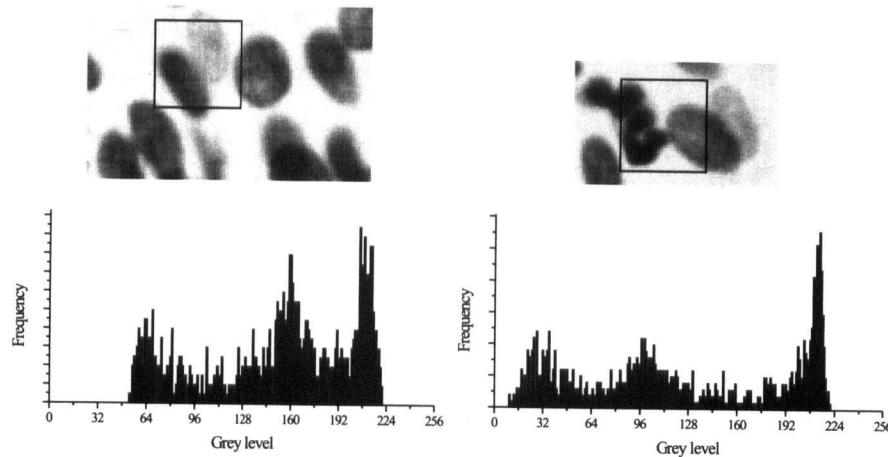


Figure 2.7: Local 30×30 neighborhoods for portions of the lung tissue section image (figure 2.3) and their grey level histograms.

the histograms contains three modes: one for the background, and one each for the light and dark objects within the outlined neighborhood. Fitting a single Gaussian to the nuclear pixels as a whole will not produce a good fit to the data. Thus, we are unlikely to obtain a good threshold for this region.

The problem is not solved by making the neighborhood smaller. This results in fewer pixels available to generate the histograms which leads to less reliable Gaussian fits. As well, when the neighborhood falls on the region between a light object and a dark object, it will still have a histogram like those of figure 2.7 and not lead to a suitable threshold. Yet, it is the regions at the borders of touching objects where it is important to have good local thresholds.

The difficulties demonstrated by images such as those in figure 2.7 demonstrate that local thresholding methods based on only the grey level histogram are not suitable for segmenting tissue section images that are thick enough to contain significant variations in average nuclear intensity. The Gaussian noise model could be replaced with another noise model, but this does not address the root of the problem. The fault lies in the

assumption that a predefined neighborhood gives information relevant to segmenting a single object. It is necessary to look to other indicators of the boundaries of objects, such as using the image gradient, in order to develop an effective thresholding scheme.

2.2.2 Compactness based thresholding

Some of the algorithms that are used successfully in segmenting cytological scenes do not work well for segmenting tissue sections. The nuclei in cytology images and thin histological sections are generally non-overlapping. For all but the thinnest biopsy sections, the section images will tend to contain a large number of overlapping nuclei.

Having distinct nuclei allows the use of shaped based criteria, such as compactness. Compactness is a measure of circularity given by $C = \frac{4\pi A}{P^2}$ where P is the object perimeter and A is the area. Figure 2.8 shows the compactness for several different geometric figures. Circles have the maximum possible compactness of 1.0. All other shapes will have a compactness less than 1.0. The more intricate the border of an object, the larger its perimeter compared to its area. Several authors have described local or global thresholding methods that use the *a priori* knowledge that nuclei are mostly circular and therefore very compact objects. They calculate the compactness of the resulting mask for different threshold levels and set the threshold to the value that maximizes the compactness of the segmented objects. The method has been applied to finding nuclei in cytology [17, 55] and liver biopsies [2].

Unfortunately, this method is not suitable for tissue section images that contain a significant amount of nuclear overlap. When nuclei touch, the compactness of the mask of combined object is lower than for the nuclei alone. Setting the threshold that maximizes the compactness causes the method to choose a threshold that makes the blob, as a whole, rounder. This is accomplished by setting T to a higher value, one that includes some of the background pixels in the object mask to make it rounder.




Figure 2.8: The compactness, C , for several geometric figures.

This behavior can be seen in figure 2.9, which shows the application of the compactness based thresholding to two nuclear clusters. The nuclear clusters are shown in figure 2.9a surrounded by a polygon that shows the neighborhood used in the compactness calculations. For the pixels in the two neighborhoods, the compactness was calculated for the mask that would result for each possible threshold setting. This was plotted in the graphs in figure 2.9b. The histograms have peak values at $T = 202$ and $T = 215$ respectively, which are the optimal thresholds for maximizing the object compactness. These are unreasonably high threshold values.

Figure 2.9c shows the outline of the segmented objects at these thresholds. The reason for the high thresholds becomes clear. The selected thresholds incorporate some of the background pixels to make the object masks rounder. A good threshold would produce an outline that hugs the nuclear edges and creates a relatively long object perimeter. This would correspond to a lower compactness value. The compactness maximizing technique is not suitable for segmenting images where there is a significant amount of nuclear overlap. For the tissue sections studied in this thesis, there may be in excess of a hundred nuclei in a cluster, so another technique was sought.

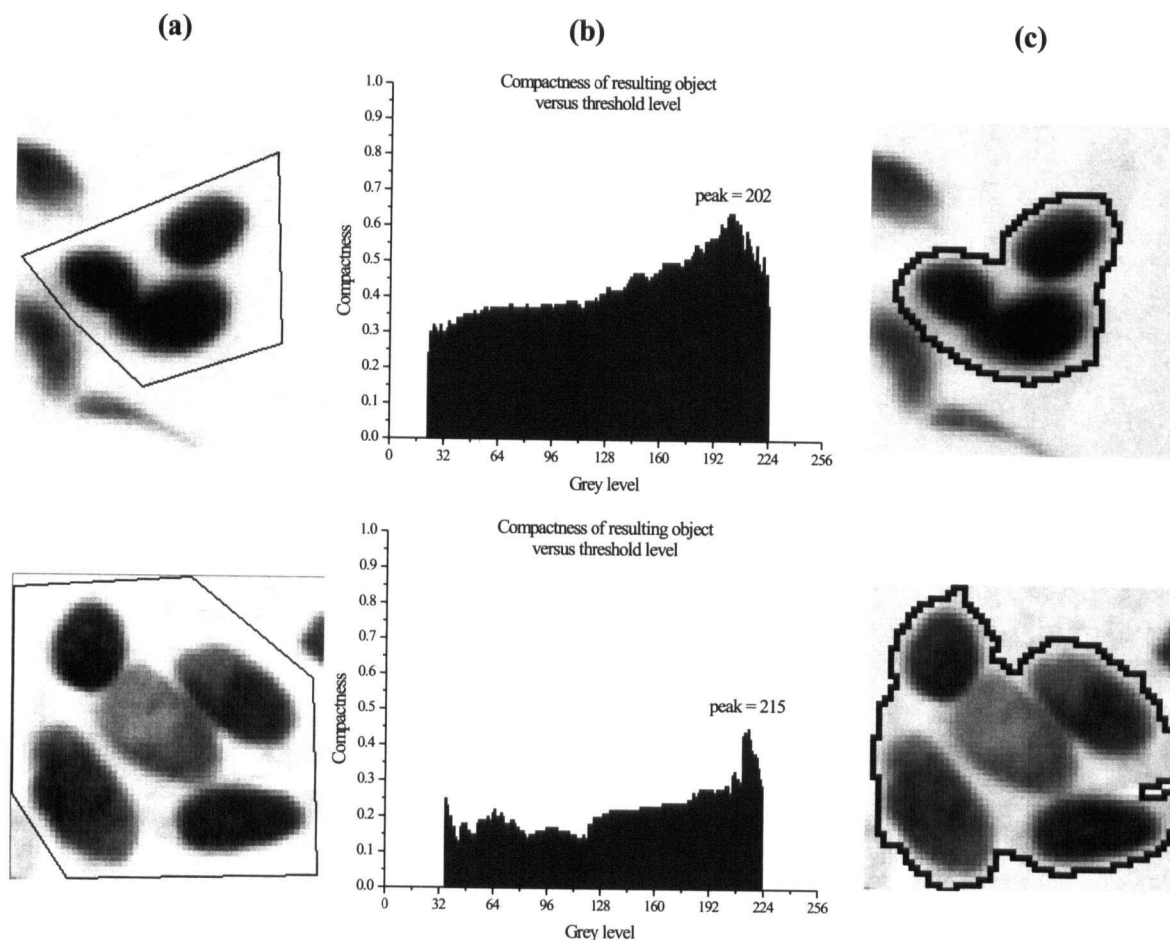


Figure 2.9: The failure of compactness based thresholding for nuclear clusters. The polygons that surround the nuclear clusters (a) define a neighborhood for which the histogram of compactness versus threshold is shown (b). Setting the thresholds at the peak value produces objects whose masks are too large (c).

2.2.3 Gradient based thresholding

An effective thresholding criterion that has been applied to medical image analysis [35, 43, 44, 69] is the gradient weighted (GW) thresholding method. For a greyscale image with values $f(x, y)$, the gradient image $g(x, y)$ contains large values wherever an edge is present in f . The gradient image is usually calculated using the Sobel filter, and will be described more fully in section 4.1.1. An example of the Sobel filter of an image can be seen in figures 4.1a and 4.1b. Letting w and h represent the width and height of the image, the GW thresholding method calculates the threshold T as

$$T = \frac{\sum_{x,y} f(x, y)g(x, y)}{\sum_{x,y} g(x, y)} \quad x = 1 \dots w, \quad y = 1 \dots h \quad (2.1)$$

T may be calculated on a global basis or on a neighborhood by neighborhood basis.

Equation 2.1 can be interpreted as a thresholding scheme that assigns large weighting to grey level values that occur near edges. At one extreme, if all the pixels had the same gradient value, $g(x, y) = j$, then T would be set to the mean of all the grey levels in the image, $T = \frac{\sum_{x,y} f(x, y)}{w \cdot h}$. At the other extreme, if only pixels with a certain grey level $f(x, y) = i$ have non-zero $g(x, y)$, then equation 2.1 assigns all the weight to the one grey value—much like an election where there is only one candidate to vote for.

Watanabe [68] described a thresholding method for cytology where the sum of the gradient values for each grey level are plotted in a histogram. The method sets the threshold at the grey level for which the total magnitude is the largest. This predisposes the method to choose a threshold at a value for which there are many pixels. Weszka [69] modified the algorithm to divide the total gradient by the number of pixels at that grey level to give what will be referred to as the mean gradient (MG) histogram method. The threshold T for this method is given by

$$T = \arg \max_i \left[\sum_{(x,y) \ni f(x,y)=i} g(x, y) \right] / n_i \quad x = 1 \dots w, \quad y = 1 \dots h \quad (2.2)$$

where n_i is the number of pixels having grey level i .

Figure 2.10a shows the graph of the mean gradient values as a function of grey level for the tissue section portion shown in figure 2.4. The trend in the graph is to rise to a peak somewhere between 144 and 170 grey levels. There are also small local peaks at around 35 and 241 grey levels. The lower peak corresponds to the regions of overlap and lymphocytes present in the image. The upper peak is due to a few noise pixels and is not significant for this image. This can be seen by referring to the histogram in figure 2.4, where the pixel count for grey levels above 235 is too small to be seen. Applying a Gaussian smoothing operator to the MG graph with $\sigma = 2$ gives a global maximum at 153, which is where the threshold is set. Figure 2.10b shows the result of thresholding the sample image at this level.

Qualitatively, the segmentation looks reasonable. Most of the objects are selected by this T value and only a few of the lighter ones are broken or excluded. Many of the nuclei are separated by this threshold and those that remain joined have a border which may be separable using shape analysis. In comparison, the GW threshold for this image is $T = 144$, which is 9 grey levels darker. These two threshold values are close enough that there are no significant visual differences between segmenting at these values.

The comparison of threshold levels by different techniques for a complicated scene is a difficult task. Fu [23] and Pavlids [53] comment that an image segmentation is an issue of psychophysical perception and not one that has an analytical solution. One option would be to have the scenes thresholded by an experienced technician and then compare the values to those selected by the two algorithms. However some scenes, with good contrast, are insensitive to the exact threshold level. Choosing a level between say, 120 and 170 grey levels might give roughly the same results. For other images, a difference of 15 grey levels in thresholds can make a significant difference in the number of objects correctly segmented. This makes it difficult to quantify what grey level difference is

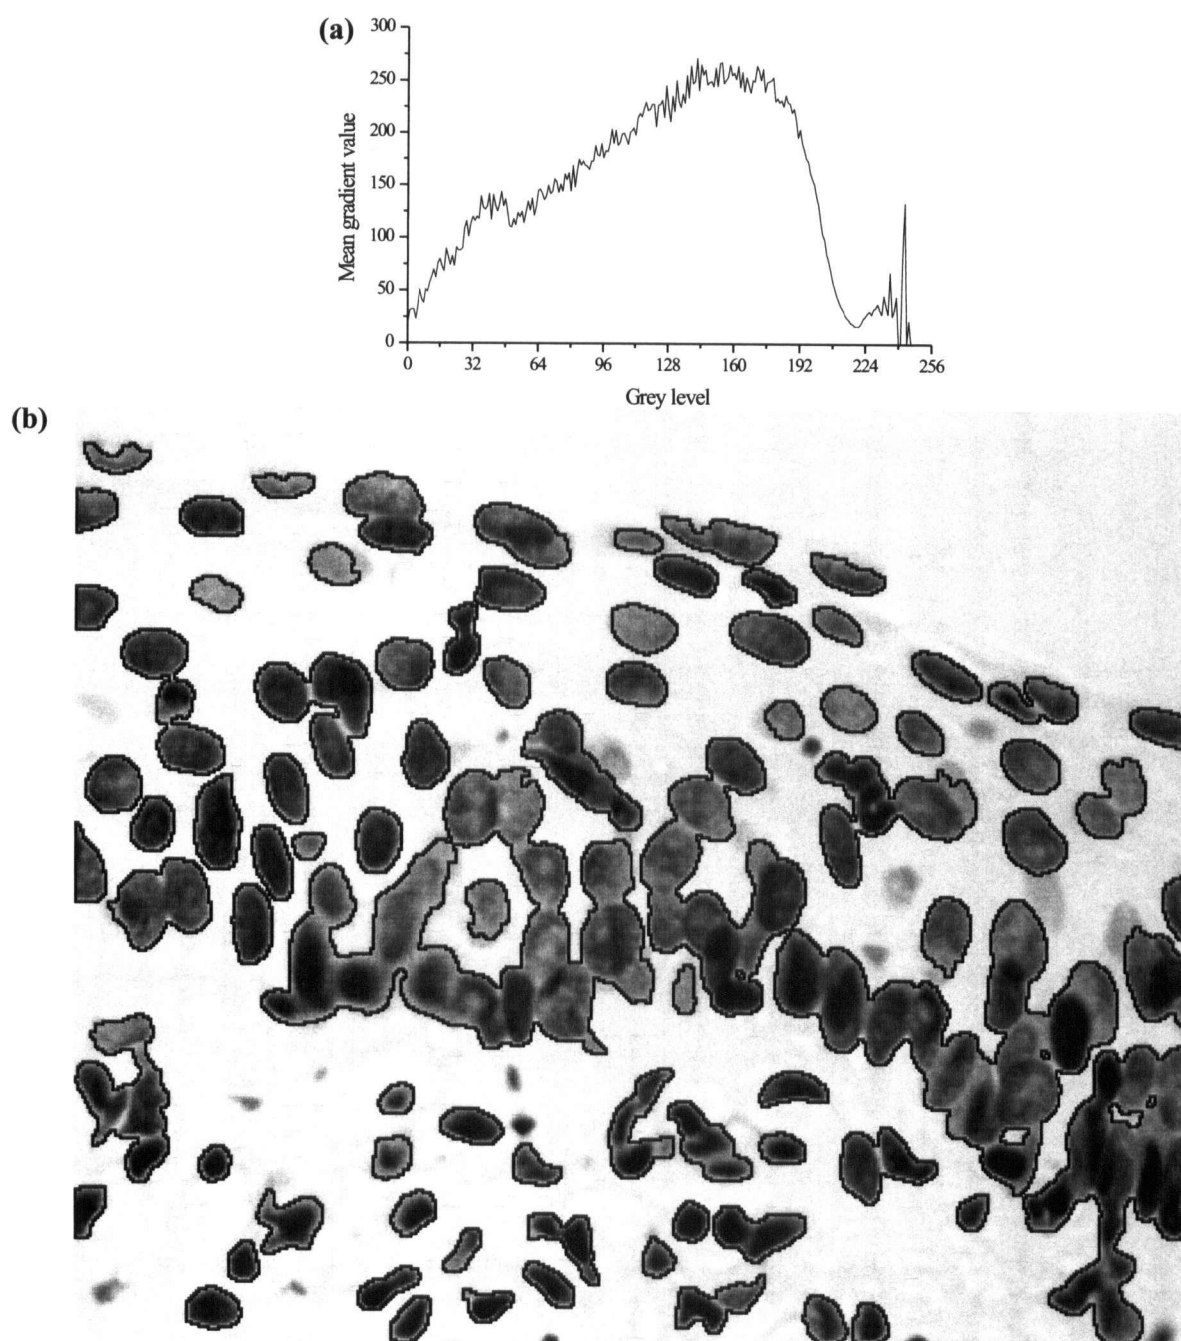


Figure 2.10: The histogram, (a), of the mean gradient value as a function of grey level for figure 2.4. Smoothing this graph and selecting the peak gives $T = 153$. Thresholding at this level produces the segmented objects outlined in (b).

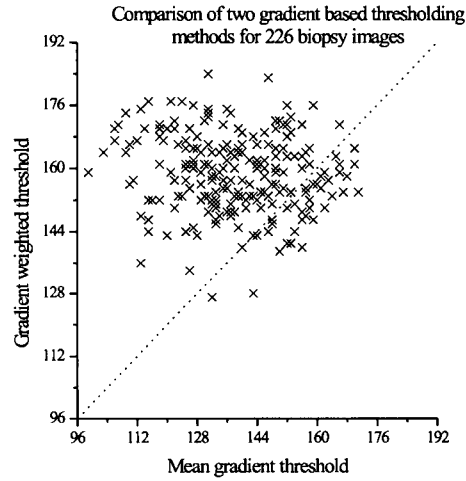


Figure 2.11: The comparison of thresholds selected by the gradient weighted and mean gradient thresholding methods.

significant between two thresholds. A second difficulty is that the threshold selected by one technician may differ from that selected by another. Each will respond to different visual cues and inevitably select a threshold that is biased by their experience in tissue section imaging.

The approach taken was to compare the GW and MG global thresholds for the 226 tissue section data set and to look for instances of large differences between the thresholds selected. The statistics for the mean gradient threshold (T_{MG}) were $\bar{T}_{MG} = 138$ with standard deviation $\sigma_{MG} = 15$ and for the gradient weighted threshold (T_{GW}), $\bar{T}_{GW} = 159$ with $\sigma_{GW} = 10$. Figure 2.11 shows the scatter of the two thresholds for all 226 tissue sections. The diagonal line represents the occurrence of both methods selecting the same threshold.

The most prominent feature of this graph is that $T_{GW} > T_{MG}$ for all but 39 cases—17% of the tissue sections. The difference between the two thresholds is a mean of 20 grey levels, but with a very large standard deviation of 15 grey levels. In eight cases,

T_{GW} exceeds T_{MG} by 60 or more grey levels. The two methods seem to be responding to different visual features in these images.

Figure 2.12 shows a portion of a tissue section where the threshold difference is 60 grey levels. The image was captured by placing the region of interest in focus while ignoring the focus of the other regions. For this image, $T_{MG} = 99$, which is verified by noting the peak of the mean gradient histogram shown in the figure inset. At this threshold, the dark nuclei at the top are included and much of the rest of the material is excluded. The threshold is low enough to cause some of the nuclei with lighter pixels to be cut. On the other hand, $T_{GW} = 159$, a threshold that is high enough to cause the inclusion of a significant number of transition zone pixels in the object mask. Because the nuclei are so close together, the epithelial layer of more than 200 nuclei are segmented into three large masses.

Setting the threshold manually suggests that a suitable threshold is around $T = 110$, which is closer to T_{MG} . The high value of T_{GW} can be explained by looking at the histogram of grey levels (solid line) in figure 2.12. The background peak is asymmetric with a significant shoulder in the region between 192 and 210 grey levels. This shoulder is due to the large number of out-of-focus pixels present in the image. Instead of having a clear transition between object and background, the out-of-focus nuclei have a smooth rise in grey level values. These transition pixels have non-zero gradient values and thereby contribute to the numerator of T_{GW} in equation 2.1.

This point is demonstrated in figure 2.13. The two images (figures 2.13a and 2.13d) contain an out-of-focus and an in-focus nucleus, respectively. Each nucleus is marked with a pair of triangles to demarcate a cross-sectional line of pixels. The grey levels of these pixels and the gradient values at these locations are shown in figures 2.13b and 2.13e. The gradients were calculated using the Sobel filter (section 4.1.1) and reflect the gradient of the 2d image, so the dashed curve is not merely the derivative of the solid




Figure 2.12: Portion of a tissue section image where $T_{MG} = 99$ (see peak of MG curve) while $T_{GW} = 159$.

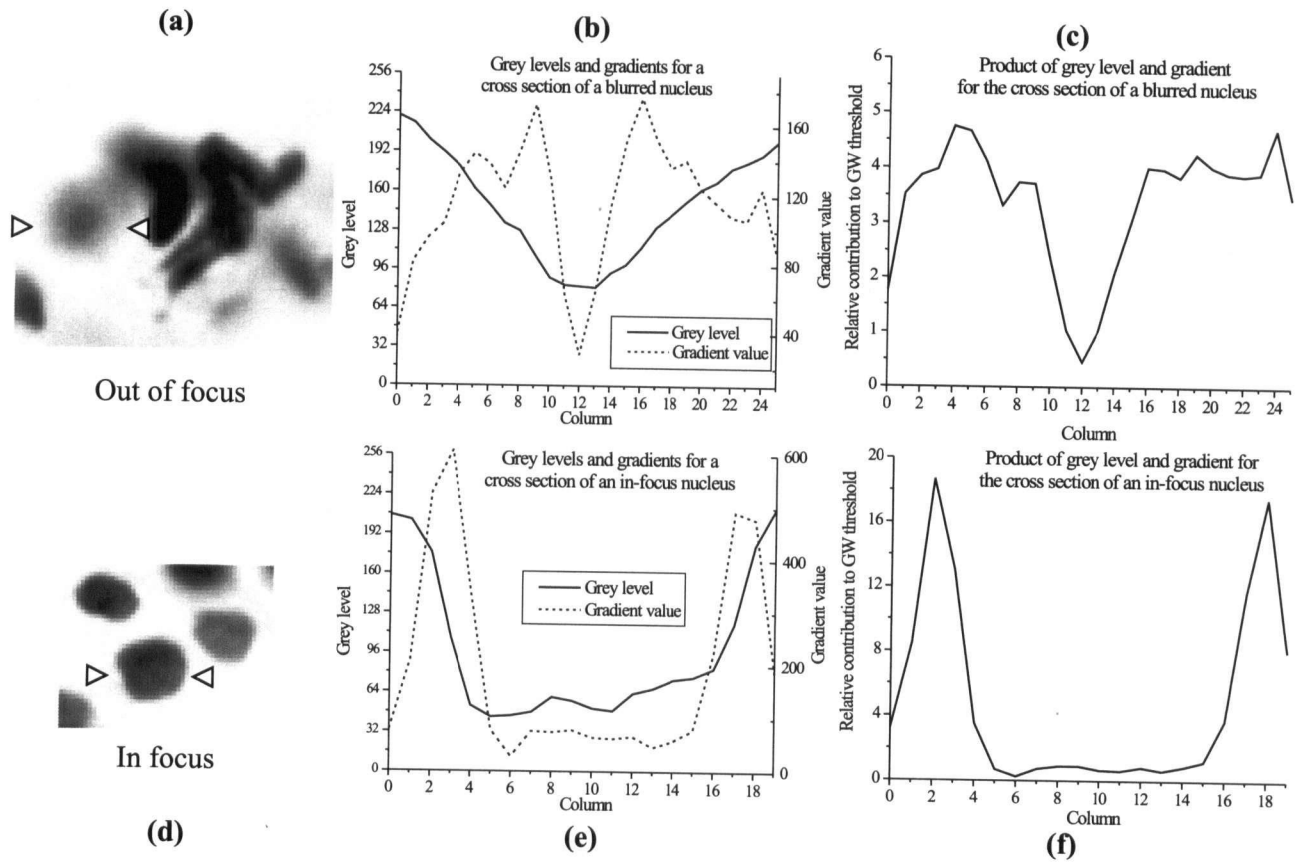


Figure 2.13: A comparison of grey level and gradient profiles (b,e) for an out-of-focus and an in-focus nucleus (a,d) selected from figure 2.12. The triangles in (a) and (d) show the cross sectional limits. The product of the two curves gives the local contributions (c,f) to T_{GW} for each image.

curve. The solid curve in figure 2.13b shows how the grey level changes only gradually across the nucleus. The gradual change causes the gradient to have a significant value at most of these pixels.

Having the gradient spread across several pixels, as opposed having a sharp transition (as is seen in figure 2.13e), raises the GW threshold above what it would be for an in-focus image. The numerator of T_{GW} in equation 2.1 can be equivalently expressed as the product of the number of pixels at grey level i with the mean gradient value at that

level. This is the product of the solid and dashed curves in figures 2.13b and 2.13e, and is shown in figures 2.13c and 2.13f. These curves represent the relative contribution to the numerator of T_{GW} of each pixel in the cross sections. Figure 2.13f shows that for the in-focus nucleus, contributions come almost exclusively from just a few pixels at the nuclear edge while in (c) contributions come from most of the pixels across the cross-section.

Both the MG and GW thresholding methods are affected by the spread of high gradient values among out-of-focus pixels. However the effect of the gradient value spread on the MG graph is to lower and broaden the peak. Provided there is a significant amount of in-focus information in the tissue section image, the effect should not significantly change the location of the peak. As the results above show, the GW threshold value is systematically raised proportional to the amount of out-of-focus information present in the image.

Examination of the images for which T_{GW} exceeded T_{MG} by 60 or more pixels revealed another difference between the methods. Figure 2.14 shows a portion of a tissue section image for which $T_{MG} = 107$ while $T_{GW} = 171$. Once again, the value of T_{GW} seems too high, labelling much of the transition areas as part of the nuclei. Manually setting the threshold indicates that a value around 120 reasonably includes most of the objects without rejecting too much. Most of the image is in good focus, so the threshold level difference is not due to lack of image focus. There is a great deal of bright texture present in the stroma and this appears to be affecting the threshold.

To demonstrate this point, the image in figure 2.14 was thresholded at $T = 226$, which is well above the background peak in the grey level histogram. The mask of the thresholded image is shown in figure 2.15a. Pixels with grey levels greater than 226 in the original image appear as white dots in this image. The large region in the top-left portion of the image contains only background pixels and no objects. The threshold has


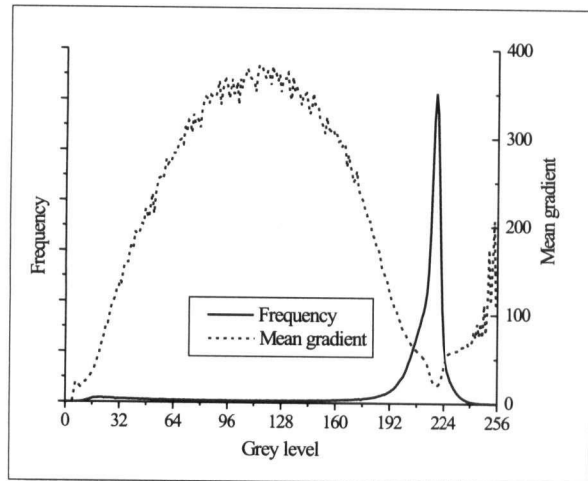


Figure 2.14: Portion of a tissue section image where $T_{MG} = 107$ while $T_{GW} = 171$.

been set high enough so that almost every pixel in this region is beneath the threshold. Notice that there are only a few white dots in the top-left of figure 2.15a, out of the many thousands of pixels in this region. Since the light path is unobstructed as it passes through this region, this should be the brightest region in the image (since the illumination is relatively uniform). However, the entire stromal structure can be discerned in the figure, and every one of these pixels is significantly brighter than the true background.

This effect is caused by refraction of light between the collagen fibres and other material of different refractive index. The refraction has the effect of concentrating the amount of light collected from some regions leading to excessively bright texture. The bright texture has a significant gradient value associated with it, as can be seen in the gradient image in figure 2.15b. According to the histogram in figure 2.14, the pixels brighter than 224 grey levels have mean gradients of roughly 60, which is around 15% of the maximum mean gradient. These pixels inflate the GW threshold, dragging it towards the background.

If a polygonal neighborhood is carefully drawn around just the nuclei at the top so as to exclude the all the stroma, the new threshold is $T_{GW} = 134$, which is closer to the true threshold. This shows that it is necessary to be careful in the application of T_{GW} to the exact region of interest rather than automatically applying it to the whole tissue section frame. The value of T_{MG} does not change significantly whether the polygonal neighborhood or whole image frame is used.

Which of these two thresholding methods is better will undoubtedly depend on the nature of the images being examined. The tendency of the mean gradient method to select lower thresholds appears to make it better suited for separating the dense arrangement of nuclei in tissue sections. However, using a lower threshold gives very tight borders to the objects being segmented and sometimes ignores some of the nuclear optical density at the perimeter of the nuclei. The borders that are generated when technologists manually




Figure 2.15: The mask, (a), for the image shown in figure 2.14 thresholded at $T = 226$, and the gradient image, (b).

segment nuclei tend to include this optical density, so the mean gradient method may not generate nuclear segmentations that match those that technologists do. In section 6, both methods are tested as part of the complete segmentation system applied to tissue section image analysis.

Chapter 3

Analyzing mask shape

Thresholding the greyscale image gives us a binary mask that contains both free lying nuclei and clustered nuclei. We need to determine which of the objects in our mask fall into each of these categories and how to split up those objects that are determined to be clusters. Analysis of the shape of the binary mask is a computationally efficient and powerful technique for performing this task. One of the primary image processing tools for describing the shape of the object is the distance transform (DT), which is described in this chapter. The analysis of shape with the DT and algorithms that use the DT is the primary tool for selecting nuclei from clusters.

3.1 The distance transform

3.1.1 Derivation

The Euclidean distance map (EDM) is a function that converts a binary image—containing object pixels and background pixels—into a real valued image where each object pixel is replaced with a value that represents its distance to the background. Figure 3.1a shows a mask (in black) for an object consisting of two overlapping circles. The EDM for this mask is shown in figure 3.1(b). The pixels in this image are scaled so that the maximum value of the EDM is represented by black and the rest are rescaled accordingly.

A simple way to calculate the EDM is to calculate the Euclidean distance between each object pixel and each background pixel in the image. This would be inefficient and




Figure 3.1: An object mask (a) with its exact Euclidean distance map (b) and distance transform (c). The DT is shown with only three grey tones, which represent the approximate integer distances from each pixel to the background. For example, the dark rings in (c) represent successive distances of 1, 4, 7, 10... units.

prohibitively time consuming for most reasonably sized images. Also, since binary images are constructed from greyscale images using thresholds, they are subject to both thresholding errors and to the quantization errors of the grey scale image. It is of questionable value to calculate an exact distance for each pixel in the mask when other sources of error in the imaging system can lead to larger losses in precision than imprecise distance calculations.

Because the binary images are not exact, it is common to use an approximation to the EDM called the distance transform (DT). The DT is calculated using an algorithmic procedure that replaces each pixel in a digital image with an integer that is proportional to the distance to the nearest background pixel. An example of a DT is shown in figure 3.1c. In essence, the distance transform of a binary image resembles a contour map, with the value of contours increasing for points situated further from the boundary.

Standard techniques for calculating the DT look at only a small neighborhood at a time around each pixel rather than performing a global search for the nearest background pixel. These algorithms approximate global distances by the sum of local distances in order to speed up the calculation of the DT. A straightforward iterative technique for constructing the map is described as follows [57]:




Figure 3.2: Pixels that touch a central pixel are shown for a grid with 4-connected pixels, (a), and 8-connected pixels, b).

1. Assign each background pixel a value of 0. Assign each object pixel an infinite value (in practical terms, a value larger than half of the image height or width).
2. Set a variable N equal to 0.
3. For each pixel with value $> N$ that touches a pixel with value N , set its value to $N + 1$.
4. Increment N and repeat the procedure until all pixels have been assigned a value.

The number of iterations required is proportional to the size of the object. The question that arises is: what is meant by “touches another pixel” in step 3? For a square lattice, it is natural to define pixels to be touching if they are 4-connected or 8-connected. Figure 3.2a shows the pixels that touch a central pixel for a 4-connected object and figure 3.2b shows the same for an 8-connected object.

In the case of the 4-connected definition, we set a pixel the value of $N + 1$ if the pixel above, below, left or right has value N . The distance between two pixels is the sum of the number of horizontal and vertical steps required to move from one pixel to another. More precisely, if (x_1, y_1) and (x_2, y_2) are the locations of two pixels in a 4-connected neighborhood, then the distance between them is given by the *city block* metric, d_4 , which is

$$d_4[(x_1, y_1), (x_2, y_2)] = |x_2 - x_1| + |y_2 - y_1| \quad (3.1)$$

6	5	4	3	4	5	6	3	3	3	3	3	3	$\sqrt{18}$	$\sqrt{13}$	$\sqrt{10}$	$\sqrt{9}$	$\sqrt{10}$	$\sqrt{13}$	$\sqrt{18}$
5	4	3	2	3	4	5	3	2	2	2	2	2	$\sqrt{13}$	$\sqrt{8}$	$\sqrt{5}$	$\sqrt{4}$	$\sqrt{5}$	$\sqrt{8}$	$\sqrt{13}$
4	3	2	1	2	3	4	3	2	1	1	1	2	$\sqrt{10}$	$\sqrt{5}$	$\sqrt{2}$	$\sqrt{1}$	$\sqrt{2}$	$\sqrt{5}$	$\sqrt{10}$
3	2	1	★	1	2	3	3	2	1	★	1	2	3.00	2.00	1.00	★	1.00	2.00	3.00
4	3	2	1	2	3	4	3	2	1	1	1	2	3.16	2.24	1.41	1.00	1.41	2.24	3.16
5	4	3	2	3	4	5	3	2	2	2	2	2	3.61	2.83	2.24	2.00	2.24	2.83	3.61
6	5	4	3	4	5	6	3	3	3	3	3	3	4.24	3.61	3.16	3.00	3.16	3.61	4.24

(a) d_4 metric
(b) d_8 metric
(c) EDM

Figure 3.3: Distances between pixels a central pixel (★) for 4-connected pixels (a), 8-connected pixels (b), and using the Euclidean distance measure (c).

The distances between a central pixel (★) and its neighboring pixels are shown in the left panel of figure 3.3 for a 4-connected neighborhood.

In the case of an 8-connected neighborhood a simple distance measure is the *chess board* metric, d_8 :

$$d_8[(x_1, y_1), (x_2, y_2)] = \max(|x_2 - x_1|, |y_2 - y_1|) \quad (3.2)$$

Using this metric, diagonal pixels also have a local distance of one. The distances between a central pixel and its neighboring pixels are shown in the central panel of figure 3.3 for this local distance scheme.

Both schemes distort distances along the diagonal. This can be seen by comparing the first two grids (figures 3.3a and 3.3b) with the third (figure 3.3c), which shows the exact Euclidean distances. In the 4-connect case, diagonal distances are too large. Diagonally adjacent pixels are separated by 2 units (one horizontal plus one vertical step) rather than the true distance of $\sqrt{2} \approx 1.414$. On the other hand, the 8-connect grid returns distances along the diagonals that are too small; it assigns diagonally adjacent pixels a separation of only one unit.

It is possible to improve the accuracy of the DT values by using approximations to the Euclidean distances shown in the right of figure 3.3. As with the earlier algorithm,

(a)

c_2	c_1	c_2
c_1	0	c_1
c_2	c_1	c_2

(b)

c_5	c_4	c_3	c_4	c_5
c_4	c_2	c_1	c_2	c_4
c_3	c_1	0	c_1	c_3
c_4	c_2	c_1	c_2	c_4
c_5	c_4	c_3	c_4	c_5

(c)

∞	1	∞
1	0	1
∞	1	∞

(d)

1	1	1
1	0	1
1	1	1

Figure 3.4: Masks containing local distance values for the calculation of the distance transform. (a) and (b) show the form of masks for 3×3 and 5×5 neighborhoods. (c) and (d) are the masks for d_4 and d_8 shown in figure 3.3.

object and background pixels are initially assigned values of ∞ and 0, respectively. A mask is defined such as the ones shown in figure 3.4. The size of mask used depends on the accuracy required. Its coefficients, c_n , are the local distances to be propagated through the distance map. Let $v_{i,j}^m$ be the value of the pixel at (i, j) at iteration m . Let $c(k, l)$ be the value of the mask at position (k, l) , with the mask center being $(0, 0)$. The value at each pixel is updated according to the formula

$$v_{i,j}^{m+1} = \min_{(k,l) \in \text{mask}} [v_{i+k,j+l}^m + c(k, l)] \quad (3.3)$$

The procedure is repeated for each pixel until no pixels change value. As before, the number of passes through the image is proportional to the size of the object.

There are many ways to select the mask coefficients c_n . Borgefors [16] reviews previous methods and calculates optimal values of c_n for different sized neighborhoods. His optimality criterion is to minimize the maximum possible difference between the pixel values of the DT and the exact Euclidean Distance Map (EDM). This difference is expressed in terms of a fraction of the maximum distance M from an object pixel to the background.

For example, for the 3×3 mask (figure 3.4a), the optimal mask coefficients are [16]

$$c_1^{\text{opt}} \approx 0.95509 \quad \text{and} \quad c_2^{\text{opt}} \approx 1.36930 \quad (3.4)$$

which gives the maximal difference between DT and EDM values of $\text{maxdiff} \approx 0.04491 * M$. This means that the maximum difference between the approximate and exact distance values is less than 4.5% of the object radius M . For a circular object with a radius of 20 pixels, this would correspond to a maximum possible error of approximately one unit in the distance assigned to the center pixel of the object.

When computing the optimal 5×5 mask, it is not necessary to compute values for c_3 and c_5 . If the value of $c_3 > 2c_1$, then c_3 would not be used since two steps of c_1 would be shorter. If $c_3 < 2c_1$, then c_1 would only be used to mark pixels next to the background and c_3 would be used exclusively afterwards. The same argument applies to c_5 . Eliminating these coefficients from the mask in figure 3.4b results in a mask of the form

$$\mathbf{DT}_{5 \times 5} = \begin{array}{|c|c|c|c|c|} \hline - & c_4 & - & c_4 & - \\ \hline c_4 & c_2 & c_1 & c_2 & c_4 \\ \hline - & c_1 & 0 & c_1 & - \\ \hline c_4 & c_2 & c_1 & c_2 & c_4 \\ \hline - & c_4 & - & c_4 & - \\ \hline \end{array} \quad (3.5)$$

Borgefors derives the optimal 5×5 mask coefficients to be

$$c_1^{\text{opt}} = 1, \quad 1.39463 < c_2^{\text{opt}} < 1.43155 \quad \text{and} \quad c_4^{\text{opt}} \approx 2.19291 \quad (3.6)$$

which gives the maximal difference between DT and EDM values of $\text{maxdiff} \approx 0.01958 * M$.

This difference of less than 2% corresponds to a maximum expected error of less than half a unit for an object with a 20 pixel radius. There are eight instances of c_4 and a total of eight instances of c_1 and c_2 (combined) in $\mathbf{DT}_{5 \times 5}$. This means that the calculation of the DT using the 5×5 neighborhood requires twice as much computational effort as using a 3×3 neighborhood. The reduction in the expected error in the distance transform,

however, justifies the extra computing expense.

Since pixels in digital images are integer valued, it is common to round the local distance approximations in equations 3.4 and 3.6. For the 3×3 case, it is best to use integer pairs for c_1 and c_2 that preserve the ratio of $\frac{c_2}{c_1} \approx 1.36$. Using $c_1 = 3$ and $c_2 = 4$ gives a ratio of 1.33 and an expected maximum error of less than 8%. Replacing the optimal values in equation 3.6 with

$$c_1 = 5, \quad c_2 = 7 \quad \text{and} \quad c_4 = 11 \quad (3.7)$$

results in a maximum error of 2.02% for the 5×5 DT. Borgefors reports that the ratios for this triple of integers is so close to the optimal values in equation 3.6 that it would be necessary to use integers more than 10 times as large in order to obtain any further reduction in the maximum expected error.

3.1.2 Locating object centres with the DT

The distance transform is a valuable tool for finding objects within clusters in binary images. It is a global algorithm that processes pixels in a raster fashion, and is therefore efficient to calculate compared to other shape analysis algorithms. Figure 3.5a shows the distance transforms for four images of overlapping nuclei. The dark outlines show the extent of the masks that are used to calculate the transforms. The transforms were calculated using the city block, chess board and 5×5 local distance methods (figures 3.5b, 3.5c and 3.5d) discussed in the previous section.

In the case of the 5×5 method, the integer coefficients of $c_n = \{5, 7, 11\}$ were used since they give the desired ratio for the local distances $\{1.0, 1.4 \dots, 2.2 \dots\}$. This choice essentially quantizes distances into steps of 0.2 pixels in the DT. For illustrative purposes, the distances have all been rounded to 1.0 pixels. This allows the drawing of the unbroken contours shown in figure 3.5d.




Figure 3.5: Distance transforms for masks of the outlined images (a). The DTs are calculated using the city block and chessboard metrics of figure 3.4 and using the optimal 5×5 local distance scheme of equations 3.5 and 3.7.

Figure 3.5 demonstrates how the DT can be used to locate the nuclei within the cluster images. Each transform image contains local maxima that closely coincide with the centres of the nuclei. These local maxima are called Ultimate Erosion Points (UEP) since the DT using the d_4 and d_8 metrics can be also be calculated (less efficiently) by successive morphological erosion operations. By calculating the DT and locating the UEPs we can obtain an estimate of the number and approximate position of objects that are present in a cluster.

There are sometimes subtle and sometimes significant differences in the DTs obtained using the different methods. For example, in panel 1, both the d_4 and 5×5 methods correctly identify local maxima at the approximate centres of the two nuclei. However, the d_8 method only returns a single UEP. In panel 2b, it is seen that the city block method returns a UEP in the wrong position for the left-hand nucleus. For the three nuclei in panel 3, only the 5×5 method finds a UEP for the nucleus on the left. Panel 3b actually shows three local maxima near each other for the nucleus on the right, which is not a desirable result. Finally, none of the methods are able to find the UEPs for all three nuclei in panel 4. There is an incorrect maximum near another UEP in the center of panel 4b, but no methods find a maximum for the nucleus on the right.

For $\mathbf{DT}_{5 \times 5}$, the missing maximum from panel 4 is not an artifact of truncating the exact DT values. This was simply a case where the degree of overlap of the two nuclei on the right of the image is large enough so that the DT algorithms do not produce a reliable result. Figure 3.6 shows the DT calculated for the same image shown in figure 3.5. Panel 1 shows the original outlined image and 5×5 DT. Panel 2 shows a new outline, where the border has been changed by adding a single pixel to the mask definition. The resulting DT now contains a local maximum at approximately the correct position of the nucleus on the right.

Figure 3.6 demonstrates the extreme sensitivity of the DT to the exact shape of the




Figure 3.6: Distance transforms for the two masks which differ by only a single pixel. Panel 1 shows the original mask definition (a) and 5×5 DT (b) for the nuclear cluster from figure 3.5. Adding one extra pixel to the mask definition in 2a causes a local maximum to appear in the DT (2b) for the nucleus on the right.

object mask. The alteration of a single pixel in a mask containing 1327 pixels affects whether a local maximum is found for one of the nuclei in the image. If the grey level value for the pixel were slightly lower, below the threshold level, the pixel would be part of the mask, and the DT would find an appropriate UEP. This demonstrates that one cannot rely solely on the DT to find the nuclei in a clustered nuclear image. The number and location of UEPs found depends on the distance metric, threshold level used to generate the mask, and pixel noise present in each greyscale image.

Despite these difficulties, the DT works very well for locating objects which do not overlap significantly. Several of the nuclei in the images in figure 3.5 were resolved, no matter which metric was used to calculate the DT. The next task in the segmentation procedure is to define objects by using the UEPs as seeds and labelling the pixels of the mask that surround them. This is done through the application of the watershed

transformation, described next.

3.2 The watershed transformation

The goal of segmenting an image is to partition the pixels of the image into different groups representing each of the objects present in the image and the background. The distance transform converts a binary image to one where local maxima are surrounded by pixels with intermediate values. It gives information about how many objects are present in the image but does not tell us how to assign each of the pixels of the mask to the various objects. This step is performed by the watershed transformation. It transforms an image containing DT peaks and valleys into an image with labelled objects. The watershed transformation takes its name from the field of topography and has been a tool for digital image analysis for more than 20 years [12, 39]. It has also become a useful tool for separating clustered objects during cytometric analysis as well [54, 47].

Consider a curve in one dimension such as the one shown in figure 3.7a. The curve has four local minima separated by hills. Imagine letting a drop of water fall onto this topographic surface. The drop of water will flow downhill until it reaches one of the local minima. Which minimum the water reaches depends on where you drop it initially. Each local minimum defines a catchment basin (the regions between the dotted lines). Any water drops released into a catchment basin will flow to the same local minimum. The places where catchment basins meet are called *watershed* points. In this one-dimensional example, the watershed points are the points marked with triangles.

The description of the watershed above is not suitable to define its behavior in all cases. For example, if a point is part of a local plateau, a drop of water placed there will not flow to any local minimum. The following definition of the algorithm to perform the watershed transformation avoids this problem.

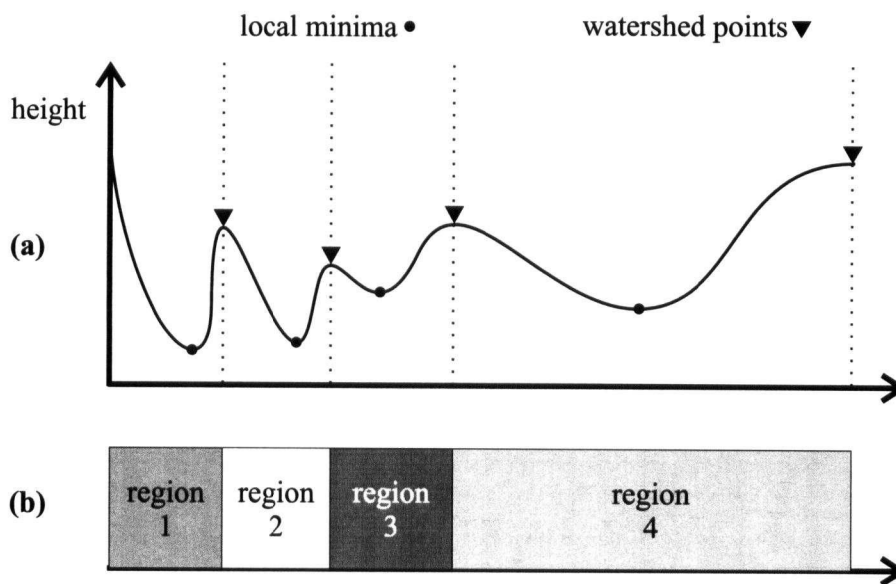


Figure 3.7: A curve in one dimension (a) contains several local minima. Each local minimum defines a catchment basin. A drop of water that falls into a particular catchment basin will flow downhill to the local minimum. The points where catchment basins meet are called *watershed* points (marked with triangles). The watershed algorithm converts the curve into a series of labelled regions, (b), with each region marking a catchment basin.

Consider immersing the one dimensional surface of figure 3.7a into water. Initially the deepest local minimum (on the left of the curve) will be submerged along with the region around it. As the surface is submerged further, other local minima will be submerged as well. Eventually, the water in adjacent catchment basins will meet. This point is called a watershed point. Imagine erecting an infinitely tall thin wall at this point at the moment when the water from two adjacent basins meet. This will keep the water from these basins from merging as the water level is increased (thereby maintaining the separation of the catchment basins and the definition of the respective regions). As the flooding progresses, more catchment basins are defined, and the watershed algorithm gives each point on the curve a label according to its catchment basin.

The final result of applying the watershed transform is the series of regions shown in figure 3.7b. The watershed algorithm thus can be interpreted as a partition function which converts a surface in space into set of regions. Vincent and Soille [67] have published an efficient algorithm for computing the watershed of a digital image in two dimensions. This algorithm was used for the segmentation experiments in this thesis. When the watershed algorithm is applied to two-dimensional images, the watersheds are lines (rather than points) that separate the image into different regions.

When the watershed is used for image segmentation it is usual to apply some kind of transformation to the image prior to applying the watershed. The nature of the transformation applied depends on the kind of image features that we are interested in. Soille [60] calls this transformed image the *segmentation function*. If we are interested in segmenting an image into regions of similar grey tone, we would use the original grey scale image in the watershed algorithm. If we are interested in finding edges, we would use the gradient (section 4.1.1) image. Since we are resolving overlapping objects based on their mask shape, the segmentation function we use is the distance transform of the image mask.

Figure 3.8 shows the application of the watershed algorithm to a two-dimensional image. The original image (figure 3.8a) consists of circles and ellipses with slight contact between the objects. The rounded 5×5 distance transform and the watershed segmentation of the DT are shown in figures 3.8b and 3.8c. Due to the minimal contact between the ellipses, there is a UEP for every ellipse with the exception of one (arrow 1). The contact between this pair of ellipses is sufficient such that they cannot be resolved by the DT—although Meyer [49] describes a local modification of the DT algorithm that may help resolve them.

There is also an instance where an extended ellipse peak (arrow 2) appears as three UEPs. These false peaks are inherent to calculating distances on square lattices and are




Figure 3.8: A binary image (a), the distance transform (b) and the watershed segmentation of the DT (c) are shown.

not a result of any rounding errors in the calculations. False peaks tend to occur when an extended local peak such as the one seen in this instance lies on a diagonal with respect to the lattice.

The circularity of nuclear images makes the watershed well suited to segmenting nuclear clusters. Figure 3.9 shows three images of nuclear clusters from lung tissue section images. In each image there are nuclei with limited contact to the rest of the cluster. These nuclei are all resolved correctly by the DT and watershed. In the cases where there is significant contact or overlap, the results are mixed. The significant overlap between three of the nuclei in the center of panel 3 causes the watershed to find a single peak. In panel 3, the watershed creates extra objects (marked by arrows) which are not actually present in the original image.

So far, it has been discussed how the positions of the UEPs can be used to segment an image. The DT values at these points can also be of use in analysing the binary image. Since the value of each DT local maximum gives the distance to the nearest boundary, a histogram of these values can be used to give an estimate of the distribution of the sizes of objects in an image. Figure 3.10a shows the histogram of the DT maxima for the cluster of circles and ellipses shown in figure 3.8. When segmenting circles, a plot of this type will show the distribution of circular radii. In the case of ellipses, such a plot will show the distribution of the short axes of the ellipses.

Figure 3.10b shows the same histogram for the three nuclear cluster images shown in figure 3.9. Since the nuclei tend to be elliptical, the plot indicates the distribution of the short axes of the nuclei, which tend to be between 7 and 11 pixels. The extreme value of 14 pixels is due to the DT maximum in the center of the cluster in panel 3 of figure 3.9. The presence of an outlier in this histogram gives us evidence that the DT has failed in this region and that the watershed segmentation is not to be trusted for this portion of the nuclear cluster. With only 21 counts in the histogram the evidence is not conclusive,




Figure 3.9: The distance transform and watershed segmentation are shown for three clusters of overlapping nuclei from tissue section images. The arrows show errors in the segmentation. A $20\mu m$ scale bar is shown in the bottom image

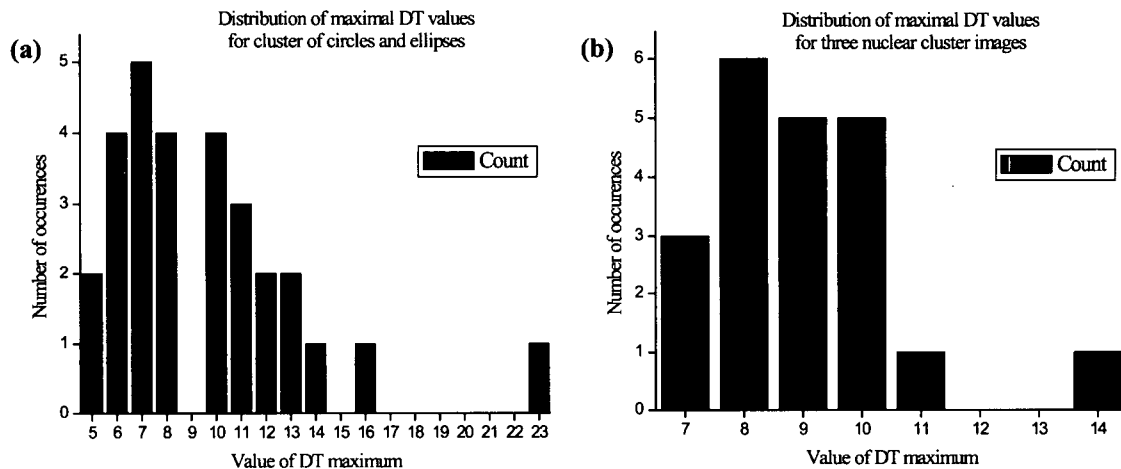


Figure 3.10: The distribution of the maxima of the distance transform for: (a) the cluster of circles and ellipses (figure 3.8), (b) the three nuclear cluster images (figure 3.9).

but when this procedure is applied to an image with hundreds of nuclei, the outliers can be more readily discerned.

3.3 Resolving overlapping objects in digital images

In order to use the watershed of the distance transform as part of an image segmentation system we need to know the limits of the algorithm's ability to correctly resolve overlapping ellipses. This requires answering several questions:

1. How much must two objects be overlapped before the DT returns a single UEP rather than two?
2. How should the overlap of the two objects be measured? As a function of separation, overlap area, or occluded perimeter?
3. When the DT returns multiple UEPs for a clump and the watershed is calculated, how can we trust that it hasn't divided a single object into two or created a false

object (as can be seen in figure 3.9)?

3.3.1 Resolving circles

To answer the first question, it is necessary to understand what causes the merging of UEPs as objects are brought together. The simplest case to analyse is the case of resolving a mask made of two overlapping circles of radii r_1 and r_2 lying on the x -axis, as is shown in figure 3.11a. Let the separation between the centres be s . The height, h , represents the DT value of the point on the x -axis where the circle perimeters meet. It has the smallest DT value of any point on the x -axis between the two circle centres.

If the circles are just touching, the DT will certainly produce two UEPs and resolve the circle centres. As the two circle centres are brought closer and closer, h increases. The condition where the two UEPs merge into one occurs when $h \geq \min(r_1, r_2)$. In the case where the circles are the same size ($r_1 = r_2$), this only occurs when the two circles are identically on top of each other. This means that two circles of the same radius should be resolvable provided that they are not identically on top of one another.

For the exact EDM calculated for circles in digital images, this is nearly the case. For circles with equal radii, it finds two UEPs provided the circles are separated by two or more pixels (dashed curve in figure 3.11b). However, when using the 5×5 DT, the minimum separation required to resolve two equal sized circles increases with the radii of the circles, as is shown by the solid curve in figure 3.11b. The graph plots the minimum separation between circle centres required for the DT to generate two UEPs. For example, when the radii are 20 pixels, a separation of at least 14 pixels between the centres is required to resolve them.

This situation is shown in figure 3.11c. The circles have radii of 20 pixels and have one UEP when separated by 13 pixels and two UEPs when separated by 14 pixels. As the separation of the two circles decreases, the height h in figure 3.11a approaches r . If

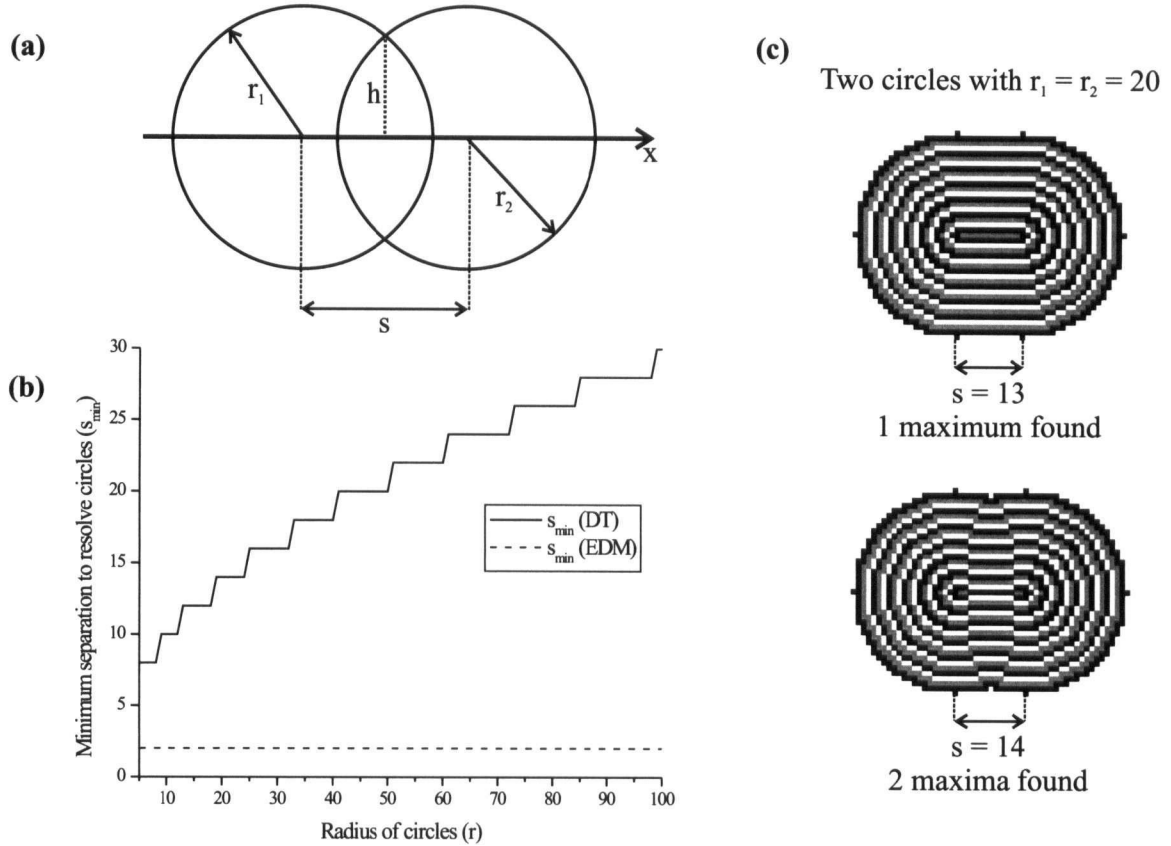


Figure 3.11: Measuring the ability of the DT to resolve overlapping circles. The graph (b) shows the minimum separation required to resolve two circles of equal size. In the case where the radius is 20 pixels (c), the DT cannot resolve two circles separated by 13 pixels, but can resolve them at 14 pixels.

the separation is small enough, the indentation disappears, as is seen in the top part of figure 3.11b. The absence or presence of the indentation determines whether there will be one UEP or two.

It would appear that the approximate DT is significantly inferior to the exact EDM for resolving overlapping circles. However, it turns out that the apparently good results of using the EDM are a result of the objects being exact circles. Removing as little as a single pixel from the boundary of the mask increases the minimum separation required

to resolve the circles. The value of s_{\min} for a pair of circles with randomly removed border pixels varies from a minimum of 2 pixels to a maximum of the DT value for s_{\min} . For example, if the single pixel at the top of either of the two overlapping circles in of figure 3.11c is removed, the result for the EDM jumps from $s_{\min} = 2$ to $s_{\min} = 14$, identical to that of the DT. Thus, there is no advantage to using the EDM in this situation unless we are resolving noiseless circles.

It should be noted that the circles and ellipses generated for these analyses are situated at the center of pixels in the grid. By this, I mean that when referring to a circle centred at (x_0, y_0) with radius r , the algorithm that creates the images generates a circle centred at $(x_0 + 0.5, y_0 + 0.5)$ that extends from $x_0 - r$ to $x_0 + r$ and similarly for y_0 . So strictly speaking, the radius is $r + 0.5$.

The odd staircase effect seen for s_{\min} in figure 3.11b arises from the description of the circle boundary with digital pixels rather than a continuous curve. All distances that would normally be described with a fractional value in a continuous space are truncated in the digital space. This truncation leads to the stepwise behavior of the s_{\min} curve.

It is possible to derive a formula for the minimum separation as is shown in figure 3.12. Let $r - h$ be the difference between the circle radius and the height of the indentation and x be the distance from the circle center to the indentation. The two circles can be resolved if the integer x is large enough that $r - h$ exceeds one pixel. This condition is satisfied when x is large enough so that $x^2 + (r - 1)^2 > r^2$. In terms of x we get $x > \sqrt{2r - 1}$ or $x \geq \sqrt{2r - 1} + 1$. Noting that $s_{\min} = 2x$, we obtain for the DT

$$s_{\min}(r) \geq 2 * (\text{INT}(\sqrt{2r - 1}) + 1). \quad (3.8)$$

This formula produces the step curve seen in figure 3.11b.

As the radii of the circles increase, the minimum separation required to resolve increases as well. However, as r increases, s_{\min} increases only as \sqrt{r} . This means that




Figure 3.12: The construction for calculating the minimum separation of resolvable circles.

s_{\min} becomes a smaller and smaller fraction of the radius. As the radius of a digital circle increases, the pixel errors in the digitization of its border decrease. For example, a digital circle with $r = 3$ looks more like a square than a circle due to digitization. The digitization effects for the border of a circle with $r = 100$, on the other hand, are hardly noticeable. The effective radius of the circles can be increased by using a finer mesh to digitize the circle. This is accomplished by increasing the resolution used to image the circles. Since $\frac{s_{\min}}{r}$ is smaller for circles imaged at higher resolution, this argues that resolving circles of equal radius is best done at the highest resolution available.

The analysis for determining the required separation to resolve two circles of different radii proceeds in a similar fashion to the case of circles of equal radii. As mentioned earlier, two circles such as those in figure 3.13a cannot be resolved by the EDM if $h \geq \min(r_1, r_2)$. In the case where $r_2 < r_1$, they can be resolved if $h < r_2$. This occurs when $y > 0$ and gives the result that the minimum separation is given by $x > \sqrt{r_1^2 - r_2^2}$. Taking the truncation of distances into account, we have $y \geq 1$ and $x \geq \sqrt{r_1^2 - r_2^2} + 1$. Therefore, the minimum separation for circles to be resolved by the EDM is

$$s_{\min}^{\text{EDM}}(r_1, r_2) = \text{INT}(\sqrt{r_1^2 - r_2^2}) + 2 \quad (3.9)$$

When $r_1 = r_2$, equation 3.9 gives $s_{\min} = 2$, which is the result obtained earlier.

As before, the DT requires an indentation to occur where the circle borders meet in




Figure 3.13: The construction (a) for calculating the minimum separation of resolvable circles where $r_1 > r_2$. Distance transforms (b) for a pair of unresolvable and resolvable overlapping circles.

order to be able to resolve them. With reference to figure 3.13, this requires that $r_2 - h$ exceed one pixel (i.e. that $h < r_2 - 1$). This gives $y > \sqrt{2r_2 - 1}$ and $x > \sqrt{r_1^2 - (r_2 - 1)^2}$. Since $s_{\min} = x + y$, we obtain for the DT

$$s_{\min}^{\text{DT}}(r_1, r_2) = \text{INT}(\sqrt{r_1^2 - (r_2 - 1)^2}) + \text{INT}(\sqrt{2r_2 - 1}) + 2 \quad (3.10)$$

Substituting $r_1 = r_2$ into equation 3.10 gives the same result as equation 3.8, which shows that the latter equation is a generalization of the former.

A comparison of $s_{\min}^{\text{EDM}}(r_1, r_2)$ (equation 3.9) to that of $s_{\min}^{\text{DT}}(r_1, r_2)$ (equation 3.10) again shows that the EDM is able to resolve circles with smaller separation. For example, figure 3.13b shows the DTs for two cases of overlapping circles with radii $r_1 = 20$ and $r_2 = 12$. According to equation 3.9, the EDM is able to resolve these circles at a separation of 18 pixels. The circle centres cannot be resolved by the DT when they are

separated by 21 pixels but can be resolved when they are separated by 22 pixels or more.

As before, the superior performance of the EDM over the DT is a result of applying the EDM to noiseless images. Altering the border of the object mask by as little as a single pixel changes s_{\min} so that it ranges from a minimum of s_{\min}^{EDM} to a maximum of s_{\min}^{DT} . Since experimental images will contain noise and objects that are not even perfect ellipses, there is no appreciable performance loss in using the 5×5 DT instead of the EDM.

The form of equation 3.10 shows that (holding r_2 constant) $s_{\min}^{\text{DT}} \propto r_1$. This behavior is different to the case of equal radii circles where $s_{\min}^{\text{DT}} \propto \sqrt{r}$. This difference reflects the fact that if a small circle (radius r_2) overlaps a large one (radius r_1), the circle centres must be separated by at least r_1 or the small circle will be completely obscured by the larger.

If $s_{\min}^{\text{DT}} \propto r_1$, this appears to negate the advantage of imaging the overlapping circles at higher resolutions. Doubling the resolution creates circles that are twice as big, but twice the separation is required to resolve them. So it appears that there is no gain to using the higher resolution. However, the form of equation 3.10 is such that when imaging overlapping circles that have a fixed range of radii, the exact dependency of s_{\min} on the radii is somewhere between square root and linear dependence. If the range of radii is small, $s_{\min}^{\text{DT}} \propto \sqrt{r}$ and we are best off imaging at the highest possible resolution. If the range of radii is very large, we are still constrained to choose the resolution based on being able to image the largest size of object of interest. So we would still try to use the highest resolution that is available.

There are two main conclusions to be drawn from the analysis of distance transforms on binary images of overlapping circles. The EDM is superior to the DT for resolving noiseless circles in digital images. However, when the circles are subject to noise or distortion affects that change their boundaries, the two methods produce equivalent

results. The second conclusion is that there is an advantage to resolving circles using the highest resolution available. If the circles are of equal size, the separation required is proportional to the root of the circle radii. If the circles have different radii, the proportionality lies somewhere between square root and linear, depending on the range of radii being imaged.

3.3.2 Resolving ellipses

The rotational symmetry of circles allows a complete analysis of their separability with DT in terms of three parameters—their center to center separation, s , and their respective radii, r_1 and r_2 . When imaging clustered nuclei, however, we are interested in resolving overlapping ellipses. The simplest instance of resolving two arbitrary ellipses requires adjusting seven parameters. Deriving a closed form expression for the minimum separation in terms of seven parameters is not feasible, so it was necessary to take another approach.

The approach used in this thesis was to perform a Monte Carlo analysis to measure the separability of pairs of overlapping ellipses as a function of overlap. Binary images of overlapping ellipse pairs were generated with random axes lengths and orientations. The positions were also randomly selected such that they overlapped to some degree.

In order to perform the Monte Carlo analysis, it was necessary to select a range of axis lengths for the ellipses. Just as was seen with overlapping circles, the ability of the DT to separate the ellipses depends on their sizes. In order to make the results relevant to the tissue section images segmented in this thesis, the ellipse size range should match the expected sizes of nuclei within these images.

It is possible to obtain an estimate of nuclear parameters using the watershed segmentation as a starting point. For example, figure 3.14a shows the best fit ellipses based on the watershed segmentation of a cluster of nuclei shown in panel 3 of figure 3.9. The

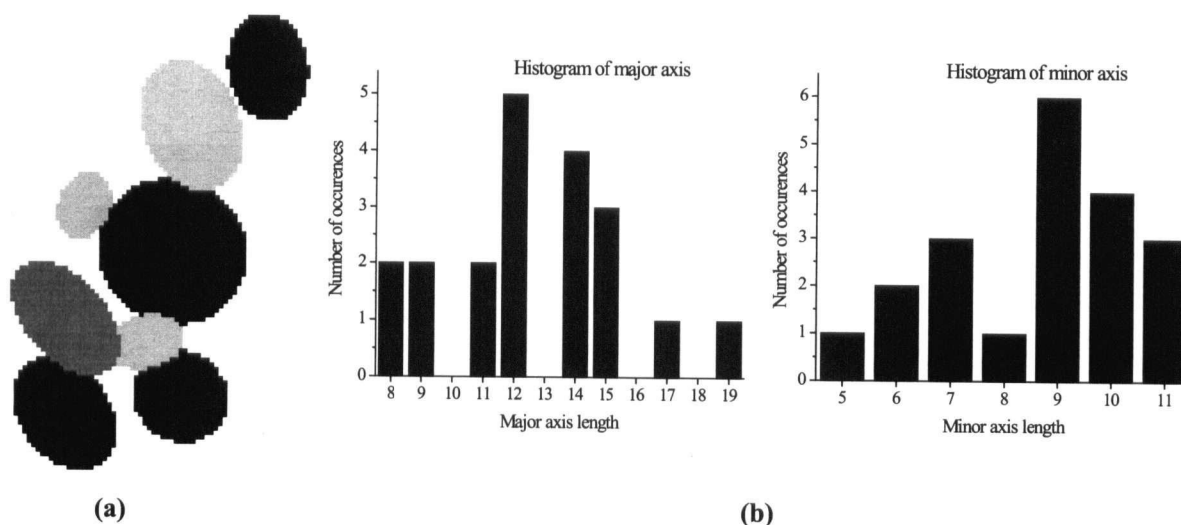


Figure 3.14: The best fit ellipses (a) for the watersheded of the nuclear cluster in panel 3 of figure 3.9, and the histograms (b) of axes lengths for fit ellipses for all 3 clusters in that figure.

ellipses are calculated from the first and second moments of the object center of gravities. The axes lengths of these ellipses provide an estimate of the axes lengths of the nuclei in the image. Figure 3.14b shows the histograms of the major and minor axes lengths for the ellipse fits to all three clusters in figure 3.9.

This histogram procedure was applied to the set of 226 images of lung epithelial tissue sections obtained at $20\times$ magnification with the image acquisition system. Each image is a frame of 1280×1024 pixels containing between 200 and 500 nuclei along with debris, lymphocytes and other objects. The watershed of the distance transform was calculated for the whole frame and ellipse fits to the watershed were calculated. Figure 3.15a shows the histograms of the axes lengths for the 70,570 objects found in the images. Figure 3.15b shows the radii of the best fit circles to the same masks.

The results were obtained without any intervention in the segmentation procedure,

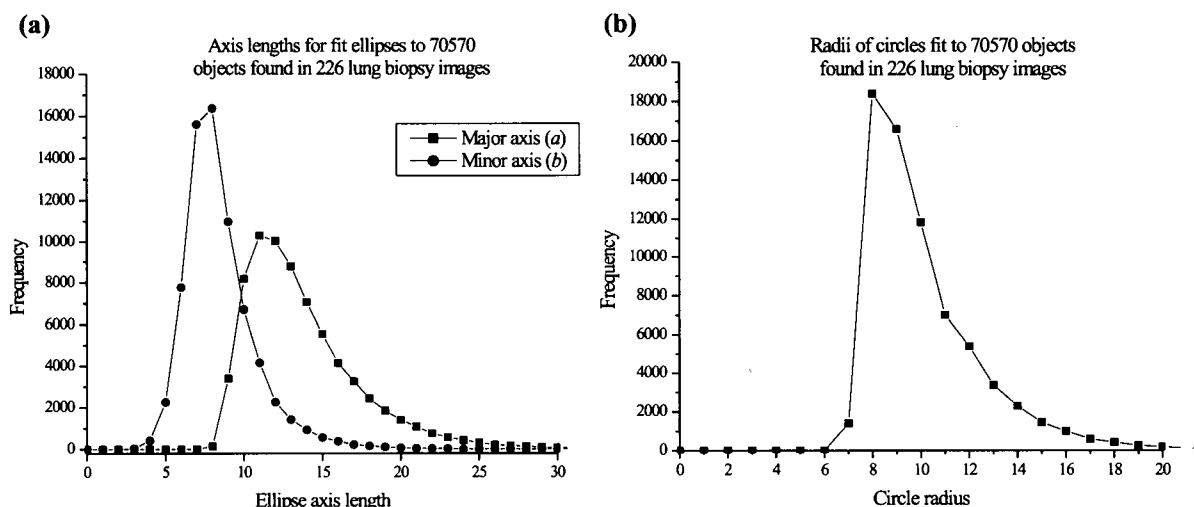


Figure 3.15: The histogram of major and minor axes lengths for the fit ellipses, (a), and radii of circle fits, (b), to the watershed segmentation of 226 lung epithelial tissue sections.

with the exception of automatically rejecting objects with areas less than 150 pixels as being too small to be a nucleus. They therefore contain many instances of improper segmentation, such as the ones seen in figure 3.9. However, the effect of these kinds of errors are to produce excessively large or small objects. By truncating the extremes of the length distributions of the ellipse major axis, (a), and minor axis, (b), we obtain an approximate estimate of ellipse parameters for the Monte Carlo simulations. The portions of the length distributions used were $9 \leq a \leq 19$ and $5 \leq b \leq 13$ pixels.

The two axis length distributions were assumed to be independent for the generation of ellipses for test images. This is not completely true since the coefficient of correlation for a and b for the 70,570 objects was $\rho = .44$. The positive correlation suggests that the objects that were found were slightly more round than oblong. Since the correlation is positive, treating a and b as independent allows the creation of more elongated

objects than were actually measured in the tissue section images. However, the automatic histogram procedure is somewhat flawed because the DT is more inclined to fail for elongated objects on digital grids (as was seen in figure 3.8). So it is likely that many elongated nuclei were missed in this procedure. Treating a and b as independent distributions was unlikely to alter the results, so this assumption was made when generating test images.

As mentioned earlier, the rotational symmetry of circles allowed the derivation of a formula for the minimum separation required to resolve circles in terms of their radii (equation 3.10). When dealing with ellipse pairs, there is no rotational symmetry (although there is axial symmetry). Also, there are four different length parameters, $\{a_1, b_1\}$ and $\{a_2, b_2\}$. This causes more difficulty in describing the overlap between two ellipses. Figure 3.16a shows the parameters for a pair of ellipses while figures 3.16b-d illustrate three possible measures to describe their overlap. The measures are by no means exhaustive, but they are the three most obvious ways to refer to the overlap.

The first measure, MA (figure 3.16b), measures the overlap of the ellipse pair as a function of the ellipse area. The area of overlap, OA, is common to both ellipses, which will likely have different areas. The measure MA is selected to be the maximum of the two percentage overlaps MA_1 and MA_2 . The maximum of MA_1 and MA_2 is used because it makes the measure more sensitive to differences in object separability. For example, when a small ellipse (area= A_1) significantly overlaps a large one (area= A_2), MA_1 will be large while MA_2 will be close to 0. The watershed will be unlikely to separate the ellipses in this case, but MA_2 would give us no clue that this is the case. Hence, for all three of the measures, the maximum of the two possible values is used to give a larger correlation between measure value and separability of ellipses.

The second measure, MP (figure 3.16c), measures the overlap as a function of the perimeters occluded by their overlap. The occluded portion of the ellipse borders are

Three Possible Measures of Overlap for Ellipse Pairs

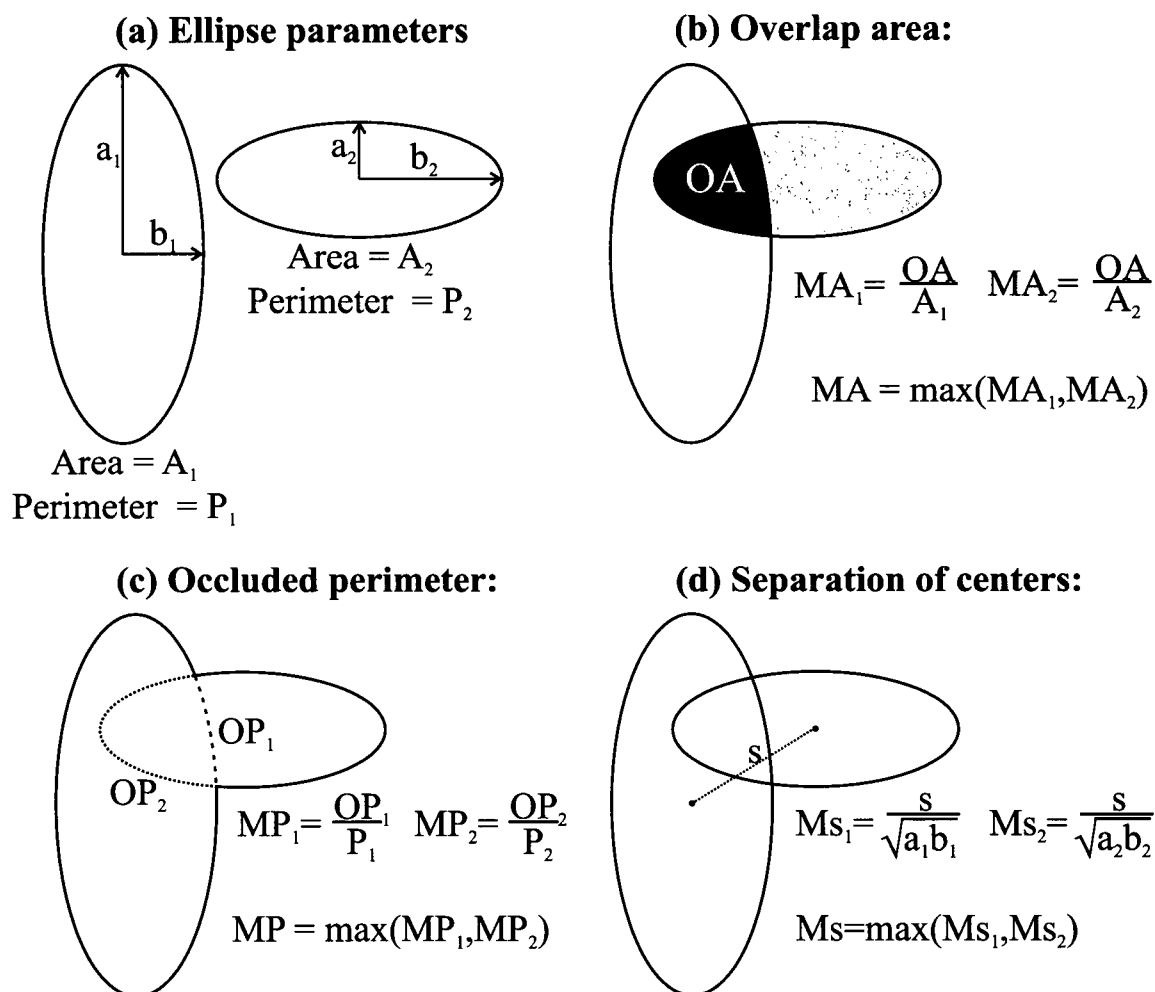


Figure 3.16: Possible measures of overlap for Monte Carlo simulations of watershed separation of overlapping ellipses. (a) shows the parameters of the ellipses, while (b-d) show three possible measures based on the parameters.

denoted in the figure by OP_1 and OP_2 . The measure MP is the maximum of the two percentages MP_1 and MP_2 . Since both MA and MP are percentages, they are bounded between 0–100%.

The third measure, Ms (figure 3.16d), measures overlap as a function of ellipse center separation divided by the “size” of the ellipse. Because it is based on the linear separation of the centres, it bears the closest resemblance to the one derived for circles in equation 3.10. However, in order to remove the dependence of Ms on the scale of the ellipses, it is divided by the geometric mean of the ellipse axes lengths. This renders the measure unitless and eliminates its dependence on the scale of the ellipses.

The area of ellipse with axes $\{a, b\}$ is πab , so the geometric mean turns out to be directly proportional to the square root of the area. The lower extreme for Ms, 0, corresponds to two ellipses with concentric centres. The upper extreme for Ms occurs when a largest possible ellipse just barely contacts a smallest possible ellipse along their major axes. In this situation, the separation s is maximized while \sqrt{ab} is minimized. For ellipses used in this Monte Carlo analysis,

$$Ms_{max} \approx \frac{a_{max} + a_{min}}{\sqrt{a_{min}b_{min}}} = \frac{19 + 9}{\sqrt{9 \cdot 5}} = 4.2$$

Since Ms is a ratio of lengths, it is somewhat more difficult to interpret than the other measures. One way to envision the measure is that a value of $Ms=1.0$ occurs when the center of a smaller ellipse is located very near the border of a larger ellipse. This represents a situation where the watershed is somewhat more likely to fail to resolve the objects than to succeed.

Figure 3.17 shows the calculation of all three measures for two instances of overlapping ellipses. In both cases, ellipse 2 is smaller than ellipse 1, so each of the three overlap measures reflects the larger degree of overlap experienced by ellipse 2. In figure 3.17a, the ellipses are resolved despite 41% area overlap and 31% perimeter occlusion of ellipse

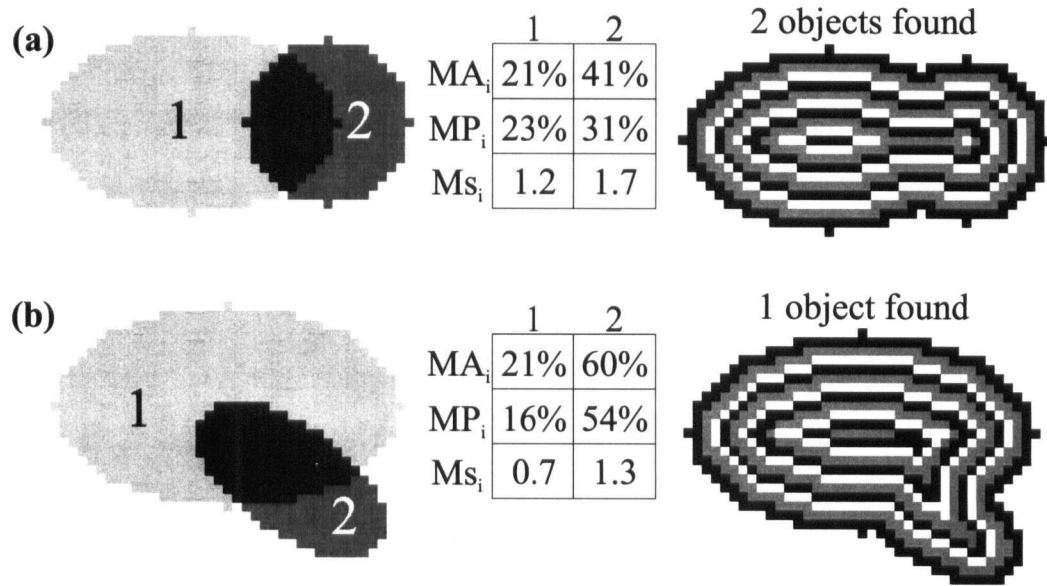


Figure 3.17: The three overlap measures for an instance of two resolvable ellipses (a) and two unresolvable ellipses(b).

2. However, in figure 3.17b the overlap is large enough so that the ellipses cannot be resolved.

In the Monte Carlo analysis, masks for one million randomly generated overlapping ellipse pairs (such as those seen in figure 3.17) were generated. The ellipse axes lengths were selected using the truncated form of the distributions in figure 3.15, as discussed earlier. The integer histograms (without smoothing) were used to generate the axes lengths, so these lengths were constrained to be integers. The xy coordinates were randomly selected so that the overlap ranged from minimal to complete overlap.

For each mask the distance transform was calculated. In many cases, the DT found only one local maximum, as is seen in figure 3.17b. When the DT found more than one maximum, the watershed segmentation was performed (figure 3.9), and ellipse fits to the watershed were generated (figure 3.14a) The separation between each of the ellipse fit centres and the true ellipse centres were calculated. In instances where the separation

was less than 25% of the ellipse axis length the center was considered to be close enough to be called a success. Otherwise, the instance was marked as a failure. The success or failure was tabulated as a function of each of the three measures, and histograms were generated.

Figure 3.18a shows the graph of the percentage of ellipse pairs resolved successfully versus the percentage of area overlap, MA. The solid curve shows the results for the $\{a, b\}$ parameter distributions discussed earlier, which were obtained from images obtained at 20X magnification. The dashed curve shows what the results would be if the scale of the ellipses were doubled (i.e. if we were attempting to segment nuclei in images obtained at 40X magnification). Both curves show a gradual decline as the amount of overlap increases. In contrast, the graph of separability versus percentage of perimeter occluded, MP (figure 3.18b), shows a steep decline for both magnifications. The final graph (figure 3.18c) shows the percentage of ellipses resolved versus the center separation measure, Ms.

For all three measures, the 40X separability results are roughly the same as the 20X results for extreme values of the measures. However, in the middle of the measure ranges, the 40X curve is nearly 20% higher in each of the three cases. This leads to the conclusion that using an increased resolution increases the separability of overlapping ellipses based on shape. Just as with resolving overlapping circles based on shape, the digitization and truncation errors are minimized as their size increases, so this result is expected.

The 40X results for MA and MP (figures 3.18a and 3.18b) remain 3–5% higher than those of the 20X even as both measures approach complete overlap. This is due to errors in calculating the distance transform using $\mathbf{DT}_{5 \times 5}$. As mentioned earlier, the expected error in the $\mathbf{DT}_{5 \times 5}$ value at any pixel is approximately 2.02% times the Euclidean distance to the nearest boundary. The ellipses used in the 40X analysis were twice as large and had a correspondingly larger error in their DT values. These errors are large enough to

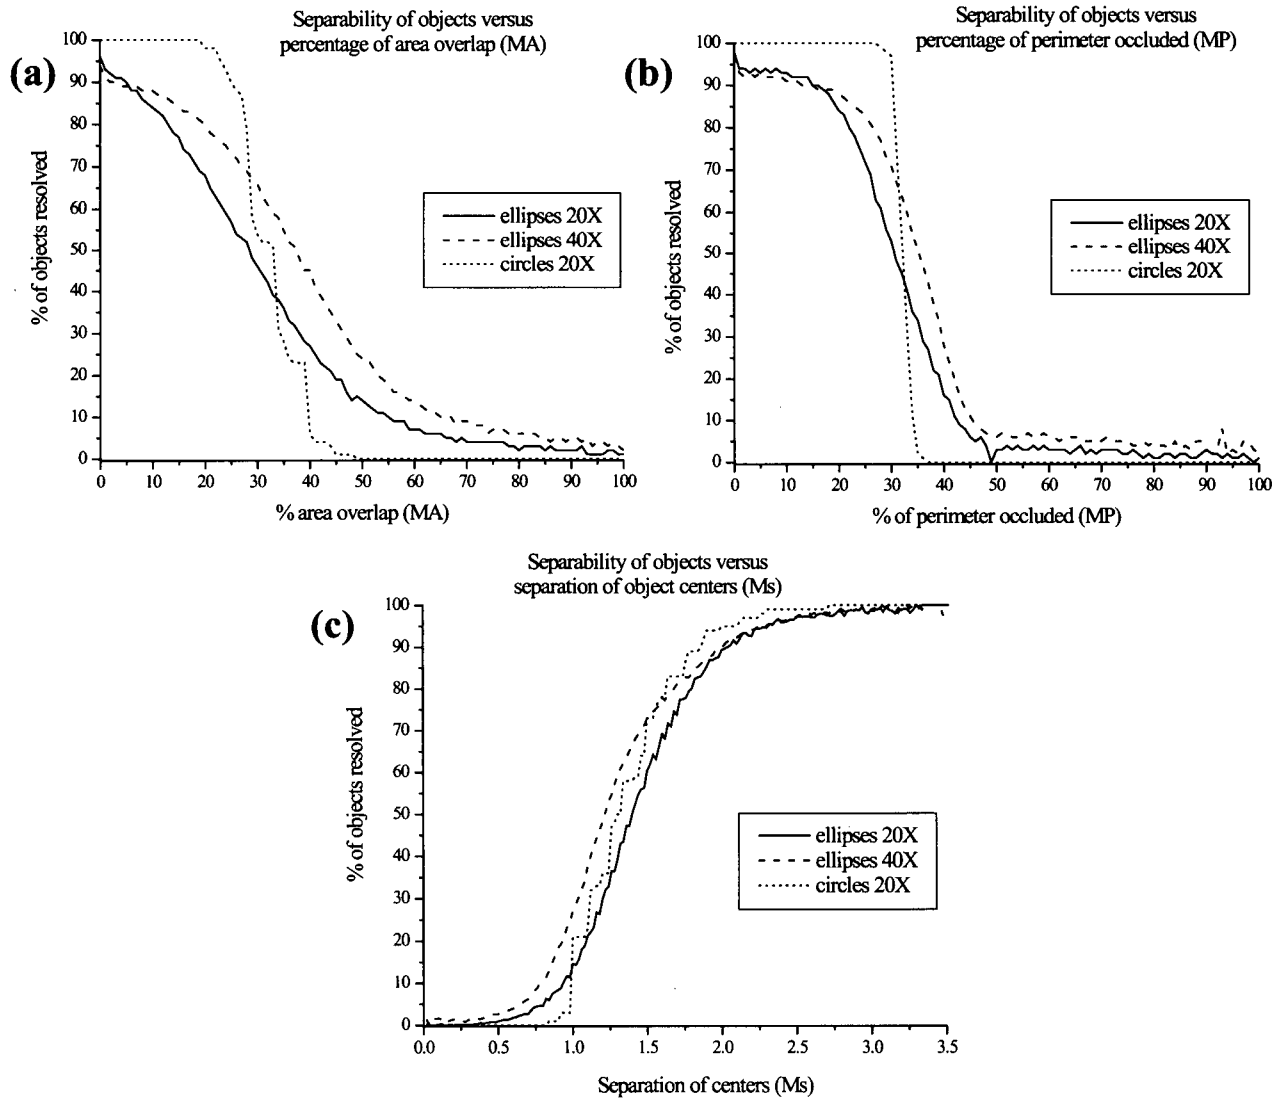


Figure 3.18: The graphs of separability of one million randomly generated overlapping ellipse pairs segmented with the watershed versus each of the three measures of overlap. The theoretical results for circles with the same area distribution are overlaid with dotted lines.

cause false maxima to appear in the DT and create extra watershed objects. These extra objects may happen to fall in the region of overlap, causing the tabulation procedure to incorrectly label them as a successful segmentation.

It is possible to derive closed form expressions for MA, MP and Ms in terms of r_1, r_2 and s for overlapping circles. This is done in Appendix A. Combining the measures with the equation for s_{\min} (equation 3.10) we can tabulate the theoretical separability of circles as a function of each measure. This was done using circles with the radius distribution shown in figure 3.15b. As with the ellipse axes distributions, the extremes were truncated such that the circle radii were limited to $7 \leq r \leq 16$ pixels. The results are shown in the dotted curves in figure 3.18. The curves were calculated by generating a table of measure values as a function of r_1 and r_2 for $s = s_{\min}(r_1, r_2)$. Then for each possible measure value in the histogram, the joint densities for (r_1, r_2) pairs that were separable at this value were summed.

For each measure, the theoretical circle separability curves show the same qualitative behavior as the curves for the Monte Carlo simulation for ellipses. However, the theoretical curves drop off more steeply than experimental curves. The most striking result is the steep drop-off of the circle separability around MP=30%. At MP=31%, 97% of the overlapping circles can be resolved while at MP=35% this figure drops to 2%.

The stepwise nature of the theoretical separability versus MA and Ms curves (figure 3.18a) occurs because of the limited number of radii used. With $r_{\min} = 7$ pixels, $r_{\max} = 16$ and the constraint that $r_1 \leq r_2$, the tables of measure values have only 50 entries ($10 \cdot \frac{10}{2}$). The curves generated from these tables will necessarily contain a series of steps rather than a smooth transition.

When attempting to segment an experimental image we are approaching the task in the reverse of the manner described so far. We start off with an image of a nuclear cluster and threshold to obtain the object mask. We do not know how many nuclei are present

in the mask and apply the DT and watershed to find out. If the watershed returns a single object, we can do no more processing based on shape and must move on to some other method to break up potential clusters.

If the watershed returns more than one object, we would like a measure, M , that tells us what the likelihood is that the watershed has done a good job. It is desirable that the ellipse separability be sensitive to changes in M . An ideal measure would have an operating point such that if M is less than this value, the watershed is known to be 100% successful and if M is greater, then the watershed is known to always fail. Then we can accept or reject the watershed segmentation with confidence. Of the three measures, the slope of the perimeter occlusion measure curve in figure 3.18b, is the steepest and therefore has this property to the greatest degree. Further, we must also have a way to calculate the measure using the information at hand. Each of the measures above relies on knowing the true ellipse parameters, which is something that is not available.

After the watershed has been applied, there is useful information in the form of the contact length between different watershed objects. Figure 3.19a shows a pair of overlapping circles, and the watershed segmentation of the mask of these circles. The MP measure requires knowing the length of the occluded portion of the perimeter of each of the objects, OP_1 and OP_2 . We can obtain an estimate of these quantities by substituting CL, the contact length of the watershed border. This length underestimates OP_1 and OP_2 . For ellipses that overlap only slightly, the difference is not significant. In the worst case, the contact length is close to the diameter of the object, whereas the occluded perimeter is $\frac{\pi}{2}$ times as much. However, since MP is calculated as a percentage of the occluded perimeter over the total perimeter, the total perimeter is also underestimated. The errors partially cancel each other out. Letting P_{ws} be the perimeter of the watershed

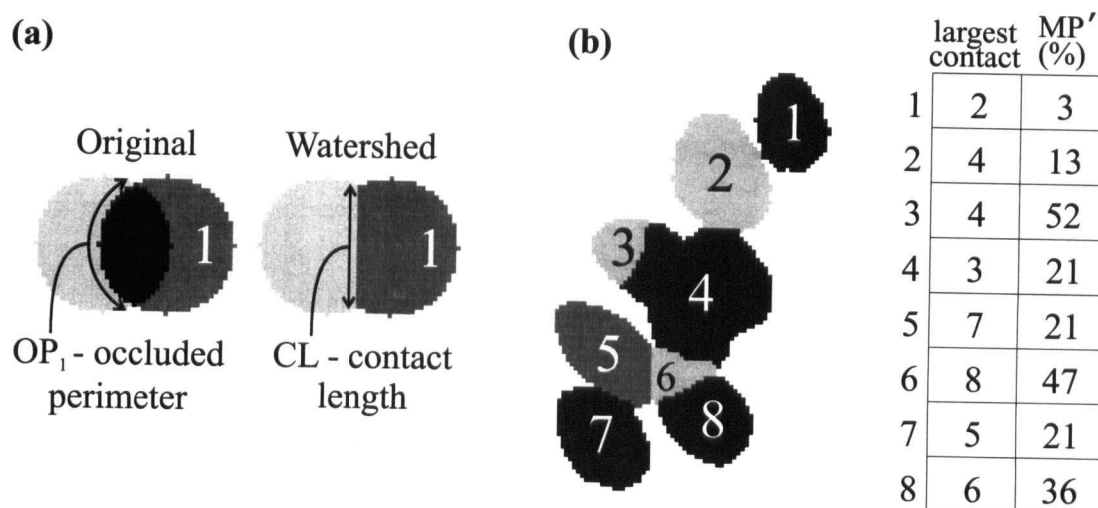


Figure 3.19: A pair of overlapping circles and the watershed segmentation of the mask (a). The MP measure requires knowing OP_i but only the contact length, CL is available. (b), the application of a reformulated version of MP (MP') to the nuclear cluster from figure 3.9.

segmented object, we define MP' as

$$MP' = \frac{CL}{P_{ws}}. \quad (3.11)$$

Reformulating MP as MP' , the ratio of the contact length divided by the visible perimeter, makes it simple to apply to experimental images. After the watershed is applied, MP' is calculated for each pair of touching objects. Figure 3.19b shows this for the nuclear cluster shown in figure 3.9. Each watershed object (labelled 1 through 8) is in contact with one or more other objects. The table shows the labels of the other objects with the longest contact length. The value of MP' for each of the contact pairs is tabulated in the right hand column. By looking for large values of MP' we are looking for likely instances where the watershed has failed. These are situations where the watershed objects should be remerged and an alternate scheme used to segment that portion of the cluster. In this figure, it is clear that objects 3 and 4 should be remerged. Also, object 6 has a long contact length with each of its neighbors. Since the contact length with object

8 is longest, the current scheme remerges objects 6 and 8.

The MP' measure is used as a decision criterion in the initial splitting of nuclear clusters in the segmentation system described in this thesis. Through initial experimentation with tissue section images containing clumps such as those seen in figure 3.9, it was noticed that object pairs with $MP' > 30\%$ were often incorrectly segmented by the watershed. Setting a threshold on MP' at 30% produced an initial segmentation that resolved most mildly touching objects and produced few errors. With figure 3.18b in mind, it is clear how this 30% figure arises. For resolving the perfect ellipses using the exact form of MP , there is a steep drop-off of separability versus MP between $MP=20\%$ and $MP=40\%$. The percentage of ellipses correctly resolved drops from around 90% to just above 10%.

Although the position of the drop-off may shift a few percent, separability versus MP' has the same qualitative behavior. The choice of rejecting watershed segmentations when $MP' > 30\%$ served as an appropriate trade-off between making maximal use of the watershed versus generating too many split objects. This analysis shows how the watershed of the mask can be used reliably for the initial shape segmentation in complex scenes. Objects with a small MP' value are likely single cells (provided that their other features such as area and optical density fall within prescribed single cell limits). Therefore, besides using the value of MP' to decide whether to remerge watershed objects, the value is a useful nuclear feature in itself. It can be used in conjunction with regular nuclear features for creating decision functions. The population of MP' values for a whole image scene could even be used to describe the degree of nuclear contact present in a tissue section image.

Chapter 4

Finding nuclei from grey scale information

When segmenting complex scenes, there are instances where the mask shape of overlapping nuclei is not suitable for correctly separating the nuclei. In these cases we may be able to use the grey scale information to aid in correctly segmenting the scene. Most grey scale segmentation techniques use the presence of edges in the image as an indicator for the limits of an object boundary. This chapter describes how edge filters such as the Sobel filter can be used to locate edges, and how to discern objects from a collection of edges using the Hough transform.

4.1 Edge detection

4.1.1 The Sobel filter

One of the primary operations for segmentation of grey level images is the detection of edges using an edge detection algorithm. Objects in the image can be distinguished from the background by locating discontinuities in the grey level value. For an image whose pixels are given by $f(x, y)$, these discontinuities are found using the gradient vector

$$\nabla f(x, y) = \begin{bmatrix} G_x \\ G_y \end{bmatrix} = \begin{bmatrix} \frac{\partial f}{\partial x} \\ \frac{\partial f}{\partial y} \end{bmatrix} \quad (4.1)$$

where G_x and G_y represent the value of the x and y -direction gradients at the point (x, y) . The magnitude of this vector, referred to as the *gradient* is given by

$$|\nabla f(x, y)| = \sqrt{G_x^2 + G_y^2} \quad (4.2)$$

The exact magnitude of the gradient is not usually as important as finding locations where the gradient value is large. For this reason, it is common to replace the squares and square root with absolute value signs in order to speed up the computation, giving

$$\nabla f(x, y) \approx |G_x| + |G_y| \quad (4.3)$$

The gradient operation is an example of a spatial filtering operator. Filtering operations are done by specifying a *kernel*, which is a matrix of weighting coefficients and convolving the original image with this mask.

The convolution of each image pixel with the kernel $\begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$, takes the pixel to the right, subtracts the pixel to the left and stores the result in the central pixel. This kernel produces large magnitudes anywhere that there is a discontinuity along the horizontal axis (i.e. a vertical line) and implements the central difference definition of $\frac{\partial}{\partial x}$. Similarly, a central difference based definition of $\frac{\partial}{\partial y}$ would be

$$\frac{\partial}{\partial y} = \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} \quad (4.4)$$

Convolution with the kernel in equation 4.4 replaces each image pixel with the difference of the pixel below and the pixel above. The two difference kernels above are sensitive to pixel noise. One way to reduce the effect of noise is to smooth the image through some kind of pixel averaging mechanism. Applying the smoothing kernel $\begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$ to the image after applying the $\frac{\partial}{\partial y}$ kernel has this effect. These two convolution operations are equivalent to using a single pass of the 3×3 kernel

$$\mathbf{D}_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} \quad (4.5)$$

Similarly, the smoothed 3×3 kernel for calculating x-direction derivatives \mathbf{D}_x is given by

$$\mathbf{D}_x = \begin{array}{|c|c|c|} \hline -1 & 0 & 1 \\ \hline -2 & 0 & 2 \\ \hline -1 & 0 & 1 \\ \hline \end{array} \quad (4.6)$$

The convolution of the image $f(x, y)$ with the kernels in equations 4.5 and 4.6 can be used to calculate the two directional gradient components,

$$G_x = f(x, y) \otimes \mathbf{D}_x \quad \text{and} \quad G_y = f(x, y) \otimes \mathbf{D}_y \quad (4.7)$$

where \otimes represents the convolution operation. Combining equations 4.3 and 4.7 defines the Sobel filter, a commonly used operator for edge detection:

$$\nabla f(x, y) \approx |\mathbf{f}(\mathbf{x}, \mathbf{y}) \otimes \mathbf{D}_x| + |\mathbf{f}(\mathbf{x}, \mathbf{y}) \otimes \mathbf{D}_y| \quad (4.8)$$

Figure 4.1a shows a cluster of overlapping nuclei from a Feulgen stained $7\mu m$ thick lung epithelial section. The Sobel filter is effective for highlighting the object from background as is shown by the strong exterior edges in figure 4.1b. The thresholded Sobel image, (c), shows that the largest magnitude edges in (b) are the exterior edges. Figure 4.1a shows that there are also significant interior edges present in the overlap regions of the nuclei. They are easy to detect visually and allow us to discern the elliptically shaped nuclei present in the left half of the image. However, these interior edges are seen only faintly in (b) and do not appear in the thresholded Sobel image. Because the grey level change from background to object is larger than the grey level changes within the nuclei, the exterior edges dominate the Sobel image. Our eye is able to detect the interior edges since it is more sensitive to the local contrast between grey levels.




Figure 4.1: A cluster of overlapping nuclei (a) is shown along with its Sobel image (b), thresholded Sobel image (c), contrast image (d) and thresholded contrast image (e). The contrast operator enhances interior edges of clusters.

4.1.2 Contrast based filtering

Johnson [33] defines a *contrast* based Sobel filter (here referred to as a contrast operator) that is useful for finding interior edges in absorption images. This operator finds edges in low brightness regions by accounting for the local average light level. This is done by modifying the Sobel value at each pixel according to the average grey level value over a specified area. The contrast based Sobel operation $C(x, y)$ is defined

$$C(x, y) = \frac{|\mathbf{f}(\mathbf{x}, \mathbf{y}) \otimes \mathbf{D}_x| + |\mathbf{f}(\mathbf{x}, \mathbf{y}) \otimes \mathbf{D}_y|}{\frac{1}{9}(\mathbf{f}(\mathbf{x}, \mathbf{y}) \otimes \mathbf{A}_{3,3}) + d} \quad (4.9)$$

where d is a DC offset and $\mathbf{A}_{3,3}$ is the 3×3 averaging mask

$$\mathbf{A}_{3,3} = \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$

The final result is then scaled to give the pixel values an acceptable dynamic range. The local averaging term in the denominator penalizes edges in bright regions more than it does edges in dim regions. For the problem of segmenting overlapping nuclei, it enhances the magnitude of interior edges of overlapping objects relative to exterior edges. Figure 4.1d shows the contrast image of the clustered objects. The ratio of interior to exterior edge strengths for the contrast image is higher than for the Sobel image. This can be seen in the two pairs of overlapping nuclei in the left of the image. Figure 4.1e shows how thresholding the contrast image retains the interior edges.

One side effect of enhancing edges in darker regions is that noise is also enhanced in these regions. This effect is seen in figure 4.1d, where there is more noise present than in the Sobel image. The trade-off between noise suppression and dynamic range is controlled by adjusting the offset d in equation 4.9. Increasing the value of d increases the amount of noise suppression in the filtered image. If the value of d were allowed to

become very large compared to the average pixel value, the denominator of equation 4.9 would approach a constant value for all pixels. In this case, the value of $C(x, y)$ would depend only on the numerator. So as the value of d increases the value of $C(x, y)$, to within a multiplicative constant, approaches the Sobel filter (equation 4.8).

There are situations where enhancing local contrast can actually hamper our ability to find the relevant edges and segment an image. If the objects being imaged have significant interior texture, enhancing the local contrast will degrade results. Figure 4.2a shows an image of HL-60 nuclei grown in culture. There is significant nuclear texture visible for this pair of abnormal nuclei. The Sobel and contrast images are shown in figures 4.2b and 4.2d. Both edge images contain a great deal of extraneous edges due to the nuclear texture. In the case of the contrast image these edges are enhanced to the point that their intensities exceed those of the true nuclear edges. This can be seen by comparing the thresholded Sobel image to the thresholded contrast image (figures 4.2c and 4.2e). The Sobel image recovers significant portions of the nuclear edges while scattered interior pixels dominate the thresholded contrast image.

4.2 Hough transform — Finding straight lines

Filtering an image to find its edges provides a starting point to segmenting the objects in it, but we need a way to recognize the existence of shapes from these edge maps. Various techniques exist to make the transition from edges to objects. These can be divided into local techniques such as edge-walking algorithms and global algorithms such as the Hough transform (HT). Hough [29] proposed an algorithm for finding straight lines in digital images. Let (x_1, y_1) be a point in the xy plane and consider the equation of a line passing through that point,

$$y_1 = ax_1 + b \quad (4.10)$$




Figure 4.2: A pair of cultured HL-60 nuclei in (a) are shown along with the Sobel image (b), thresholded Sobel image (c), contrast image (d) and thresholded contrast image (e). Edge recovery fails when the image contains significant nuclear texture.




Figure 4.3: The Hough transform for finding lines in two dimensions. Each point in the xy plane on the left gives rise to a line in the ab parameter space on the right. The intersection of lines in the parameter space represents points that are collinear in the image space for some value of (a, b) . The intersection point, (a', b') , on the right gives the parameters of the dotted line that passes through (x_1, y_1) and (x_2, y_2) .

If we rearrange this equation to solve for the unknown parameters a and b , we get the equation

$$b = -x_1a + y_1 \quad (4.11)$$

We can then plot this line (equation 4.11) in the ab parameter space as is shown in shown in figure 4.3. Each point along the line represents a different (a, b) combination that satisfies equation 4.10. If we repeat this procedure for a second point (x_2, y_2) , we obtain a second line $b = -x_2a + y_2$ in the parameter space. Let the intersection of the two lines in the parameter space be (a', b') . This parameter pair specifies the values of a and b that represent the line in the xy plane that passes through the points (x_1, y_1) and (x_2, y_2) . Thus, using two points selected from a line in the xy plane we obtain a point in the ab plane, whose coordinates specify the parameters of our original line.

This is the basis of the Hough algorithm for finding lines in digital images. We define

an *accumulator* $A(a, b)$, an array of possible values for a and b . The bounds of $A(a, b)$ are set to the expected range of values for slope and intercept, $(a_{min}, a_{max}) \times (b_{min}, b_{max})$. Then, for each feature point (x_i, y_i) in the image space, we increment all the cells of the accumulator along the line $b = -x_i a + y_i$. In essence, each feature point votes for all the lines to which it may belong. The value of each cell $A(a, b)$ is equal to the number of points in the original image that lie along this line. The final step is to scan $A(a, b)$ for peak values. Each peak, (a', b') , represents the likely presence of a line with these parameters the image space.

The power of the HT lies in its ability to detect lines under various conditions. It does not require that a line be continuous, so it will detect lines even if portions of it are obscured. Discontinuities in lines reduce the number of votes cast for a particular peak, but do not cause the method to fail entirely. Noise pixels result in spurious votes being cast in \mathbf{A} ; however unless a large proportion of the noise pixels are collinear, their votes are insignificant to those cast for the true lines present in the image. The winner-takes-all approach to finding peaks means that the position of (a', b') is not shifted by the presence of noise.

The HT of a noisy line image is shown in Figure 4.4. There are two line segments present in the original image along with a field of uniform random noise. Eighty percent of all the pixels are noise pixels. Each pixel of the image space gives rise to a line in the transform space on the right. There are two distinct peaks in the transform space. These peaks correspond to the slopes and intercepts of the two line segments on the left. The votes cast by the noise pixels are visible as the lines of intermediate intensity seen around the peaks.




Figure 4.4: The Hough transform for a noisy line image. The two lines in the image give rise to peaks in the transform space that are clearly visible despite the noise (which constitutes 80% of the pixels in the original image).

4.3 Hough transform — Finding ellipses

Although the original HT was developed to detect lines, the idea of performing a transformation from an image space to a parameter/voting space has been explored for many different shapes [30]. The term Hough transform is generally applied to any of these algorithms that use a transformation to a parameter/voting space to detect objects. To see how it can be applied to finding ellipses, consider a pixel (x, y) from an edge map that is presumed to be from the edge of an ellipse. The equation of all ellipses that pass through this point is given by

$$\frac{[(x - x_0) \cos \theta - (y - y_0) \sin \theta]^2}{a_{maj}^2} + \frac{[(y - y_0) \cos \theta + (x - x_0) \sin \theta]^2}{a_{min}^2} = 1 \quad (4.12)$$

where (x_0, y_0) is the ellipse center, a_{maj} and a_{min} are the semi-major and semi-minor axes lengths, and θ is the ellipse orientation.

Thus the vector $\{x_0, y_0, a_{maj}, a_{min}, \theta\}$ represents the five unknown parameters needed

to describe an ellipse. It is certainly possible to specify that the Hough parameter space be represented by a five dimensional accumulator $A(x_0, y_0, a_{maj}, a_{min}, \theta)$. However, there are several reasons why this approach is not generally used [40]:

1. **Memory requirements.** With five dimensions, the memory space required for this definition of A would be enormous. Two of its dimensions are image height and width, two of them are the range of all reasonable axes lengths, and the final dimension is the discretization of the angles between $0 - 180^\circ$ (into say, 5° increments). For a typical clustered nuclear image, the number of cells in this accumulator are $100 \times 100 \times 15 \times 15 \times 36$ which is more than 80 million cells.
2. **Computational complexity.** The computational complexity grows exponentially with the dimension of the problem. For each possible edge point (x, y) , we would increment all cells of A which satisfy equation 4.12. There are typically a few hundred possible edge pixels in a nuclear cluster image, which would require the casting of votes in $\approx 10^7$ accumulator cells per image.
3. **Parameter extraction.** Locating the maxima in the parameter space requires finding the local maxima of a multidimensional accumulator array. This requires using a multidimensional peak finding algorithm to find the peaks. As the number of dimensions increase the chance of spurious local maxima increases.

Besides requiring more computational effort, using the HT with ellipses is more difficult to visualize. When transforming straight lines, we go from a two-dimensional image space to a two-dimensional parameter space. Since the parameter space can be plotted, it is easy to see how peaks in the parameter space correspond to the lines in the original image. When transforming ellipses, we go from a two dimensional space to a five dimensional space.

It is possible to reduce the dimension of the transform space by using *a priori* information about the objects. Ballard described a generalized Hough transform [8, 9] (GHT) which can be used to find objects of arbitrary shape whose exact shape and orientation is known. This method uses edge orientation information and requires a template of the shape being sought. In the case of finding ellipses, this template is an ellipse with specified axes lengths. Davies [19] adapted the GHT to finding ellipses using this idea and the method has been applied to locating cell nuclei in cervical smears [64]. Since the GHT requires that both ellipse axes be specified the authors restricted themselves to searching for ellipses of a certain size with low eccentricity (essentially circles). This also eliminates the difficulty of having to deal with ellipse orientations at the cost of limited applicability of the algorithm.

Using an ellipse template is one way to reduce the dimensionality of the parameter space. An alternate way is to use pairs of points such as the diameter bisection technique of Tsuji [65]. This method considers edge points in pairs and uses the property that if two points on the edge of an ellipse have antiparallel gradient directions, then their midpoint coincides with the center of an ellipse. By comparing many such pairs of edge points we cast votes for possible ellipse centres. Since the votes are cast in a two-dimensional space identical to the image space the method has the advantage that the results are easy to visualize. To recover the remaining three ellipse parameters, the algorithm is run a second time. On the second pass, edge pixels that vote for a particular center location are recorded. The set of points that voted for a particular center are fitted using a least mean squares technique.

A typical image with multiple objects will contain many points that have antiparallel gradients. Most of these will be spurious and lead to non-useful votes being cast in the parameter space. This is a significant problem with this method. Further, most pairs of points from the edge of an ellipse do not have antiparallel gradients, so none of these




Figure 4.5: The Yuen method for finding the center of an ellipse. P_1 and P_2 are two points on the edge of an ellipse. Their tangents intersect at T and the midpoint of P_1P_2 is M . Due to the geometry of ellipses, the extension of the line TM passes through the ellipse center, C .

edge pairs help in the determination of the ellipse centres.

Yuen [73] maximizes the information contributed by pairs of ellipse edges with his method of finding ellipses based on a geometric property of ellipses. Let $P_1(x, y)$ and $P_2(x, y)$ be two points on the edge of an ellipse as shown in figure 4.5. The tangent lines for these two points meet at T and midpoint of the line connecting them is M . The extension of the line between T and M passes through the center, C , of the ellipse and reaches the other side at point O . As with Tsuji's method, the voting space is congruent with the image space, so votes for the possible center of the circle will be cast along the line specified by TM . Since the point M always lies between C and T , votes need only be cast along the segment OM , because the center must fall on this line segment

The use of all possible pairs of edge points provides significantly more information than is available to Tsuji's method, which can only use those pairs whose gradients are antiparallel. It is clear that the number of voting pairs in Tsuji's method is proportional to

the number of edge pixels m , since each edge point should have exactly one corresponding edge point with an antiparallel gradient. However, Yuen's method may potentially use $\frac{m^2}{2}$ voting pairs, maximizing the amount of information extracted from the image. This means that Yuen's method will be better able to deal with situations where significant portions of the ellipse edge are occluded.

Using points in pairs and their gradients reduces the dimensionality of the parameter space from five dimensions to two dimensions, which is more convenient to search. The price of this convenience is paid in significantly higher computation costs. As mentioned, an image with m edge points requires $\frac{m^2}{2}$ of the geometric constructions described above. A single nucleus will have at least 100 edge pixels, with 300 pixels being a more typical number. This means that there are potentially 45,000 edge pairs to be used in the voting process.

The number of pairs that are considered can be reduced. Point pairs that are too far apart can be eliminated since they must belong to different objects. This sets a limit, d_{max} , of twice the maximum axis length as the cut-off for maximum distance between pairs. Pairs that are too close together should also be eliminated. In a digital image, if P_1 and P_2 are close together, their tangent lines are more likely to be nearly parallel. Since T is calculated as the point where the two tangent lines meet, it will be very sensitive to any errors in the gradient directions at P_1 and P_2 .

A second problem can occur if the tangent lines are close to being parallel. If one considers the triangle P_1P_2T in figure 4.5, the distance from T to the triangle base, M , will be necessarily small if P_1T and P_2T are nearly parallel. Since T and M will be close together, the line TM will not reliably pass through the ellipse center. The particular value of the lower limit, d_{min} will depend on the expected size of the ellipses being sought. Throwing out pairs that are too close or too far amounts to a windowing process on the candidate point pairs. With this computational reduction, the information gain in using




Figure 4.6: The transformation from image space, (a), to vote space (b). The distribution of bright points in (b) represents likely positions of an ellipse center. The voting peaks are shown in (c). Eight of the ten nuclei are located immediately with the Hough transform.

pairwise edges outweighs the computational cost of the method.

The result of applying Yuen's method to nuclear clusters from an epithelial biopsy section is shown in figure 4.6. Figure 4.6a shows several clusters of overlapping nuclei, with three, two and five nuclei in the top, middle and bottom clusters, respectively. The clusters are outlined in black to show the extent of the masks that were determined by thresholding. Figure 4.6b shows the center position accumulators for each of the three clusters. Dark regions represent areas with few or no votes, while lighter regions represent likely positions of ellipse centres. The brightness of each accumulator is scaled so that the maximal peak for that cluster is white.

Applying a peak finding algorithm to figure 4.6b recovers the most significant peaks of the vote space, (figure 4.6c). Since the Hough transform is an evidence accumulation procedure that selects parameter combinations that receive the most votes, we can view these voting peaks as the maximum likelihood estimates of a parametric estimator. For the nuclear clusters in the figure, all five of the nuclear centres are recovered for the

top two clusters. The bottom cluster itself contains five nuclei, of which only three are recovered. The two nuclei that are not identified are fainter than the rest and do not contribute significantly to the vote space.

Once the ellipse center has been located by finding the voting peaks it still remains to find the remaining three parameters a_{maj} , a_{min} , and θ . Yuen's approach attempts to recover all three parameters at once using a three-dimensional accumulator. Another approach [50] is to run a second pass through the image to find the ellipse orientation, and a third pass to recover the axes lengths. However, the disadvantage of using multi-stage algorithms such as this is that systematic errors are propagated and compounded at each stage.

4.4 A proposed method for ellipse detection

This thesis proposes an alternate scheme to recover the remaining parameters by combining the second stage of the Hough transform with a modified form of a newer least squares algorithm for fitting ellipses. From the outset, a grey scale image containing a group of elliptical objects is transformed into an edge map using the contrast based edge operator from equation 4.9. A histogram of the gradient magnitudes is made and all pixels whose gradient magnitude exceeds a threshold are marked as potential edges for use in the Hough transform.

The position of the threshold depends on the nature of the objects being imaged. In the case of imaging absorbing nuclei, interior edges are distinct from background but are not as bright as exterior edges, despite the local contrast enhancement. Since it is believed that they can contribute to the ellipse parameter determination, the gradient threshold is set to a fraction of the median gradient value. For the tissue section images studied in this thesis, threshold values between 60–90% of the median gradient allow the

interior edges to vote and still suppress most of the noise and nuclear texture pixels. A value of 70% was used for the tissue section experiments. The initial stage of the Hough transform follows Yuen's method for finding the ellipse center (x_0, y_0) using pairs of points $P_1(x, y)$, $P_2(x, y)$ that satisfy the condition:

$$d_{min} < |P_1(x, y) - P_2(x, y)| < d_{max}$$

where d_{min} and d_{max} are described above. With each $\{P_1, P_2\}$ pair, the construction shown in figure 4.5 is used to find the line OT , which should pass through the ellipse center. Votes are cast in the accumulator along the line OT . The manner in which they are cast is discussed next.

4.4.1 Sharpening peaks in the accumulator

Different techniques exist to sharpen peaks or suppress false peaks in the accumulator to assist in the interpretation of the parameter space. Many techniques focus on assigning larger weight factors to votes cast by edge pixels that are more certain to be part of the true object. For example, some schemes use the size of the edge gradient to determine the weight factor used in the voting [8, 66]. Another method [72] notes that true lines tend to be thicker than random sequences of edge pixels and satisfy local connectivity criteria that random pixels should not. Presumed edge pixels are given larger weighting factors in the votes that they cast when the criteria are satisfied.

The commonality that ties these types of methods together is the interpretation of the Hough transform as an instance of a maximum likelihood parametric estimator. Stephens [62] demonstrates the relationship between the two by deriving a probabilistic Hough transform using likelihood principles. The result is that the accumulation of votes in the parameter space from a group of independent image features (edges) is equivalent to computing log-likelihood of a cumulative probability distribution. His method shows

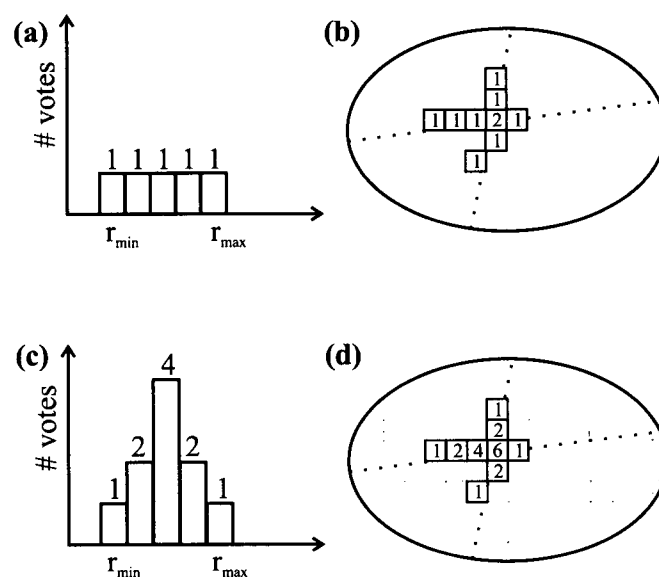


Figure 4.7: Demonstration of uniform (a,b) and weighted (c,d) accumulator voting strategies for the location of the center of an ellipse.

that for problems where a formula for the error distribution for the parameters is known, it is possible to suppress false peaks by spreading votes over a range of cells based on the parameter distributions.

When applying the ellipse Hough transform for finding the nuclear centres we do not have an explicit formula for the errors in the parameters. In the framework derived by Stephens, we can see that the ellipse finding methods that use uniform voting [19, 50] are applying a binary probability density function (PDF) for the parameters and assume *a priori* that the PDF is uniform for all valid parameter combinations. An example of this voting scheme is shown in figures 4.7a and 4.7b. The accumulator is shown after two edge pairs have voted using Yuen's construction (figure 4.5). The votes are cast over a range of valid radii between r_{min} and r_{max} . The accumulation of two votes in the cell in the center of figure 4.7b represents the fact this location is twice as likely as the others to be the ellipse center.

This suggests that it is possible to incorporate prior knowledge of the expected size distribution as weights on the votes cast in the parameter space. An example of this is shown in figures 4.7c and 4.7d. If we knew that the expected distance from border to ellipse center had a distribution such as shown in 4.7c, we could weight the votes cast accordingly leading to a concentration of votes in the accumulator space (figure 4.7d). This can sharpen the peak by making it relatively higher than the surrounding accumulator cells. This kind of weighted voting scheme can be used to account for object size distributions but does not deal with other kinds of errors that affect the voting. Edge orientation errors due to noise in the image or the inherent inaccuracy of the Sobel filter direction (expected to be around 1° [20]) will prevent the line from passing through the center.

The idea of using a weighted voted scheme determined by the known distribution of object sizes was tested for finding the centres of ellipses. For the cytological image set of 212 images described in section 1.4, the best fit ellipses were calculated and the distribution of axis lengths were determined. This distribution was used to generate synthetic images of ellipses such as those shown in figure 4.8a. Both the uniform and weighted voting schemes correctly locate the ellipse center for such perfect ellipses. However, the concentration of votes around the peak differs for the two schemes.

This can be seen by comparing the spread of votes for the uniform voting scheme and the weighted scheme (figures 4.8b and 4.8c). In both of these images the standard deviation for the peak using the weighted scheme is smaller than that of the uniform scheme. This is the case for most of the generated images, although extremely large or small images can show the opposite trend. The standard deviations, $\{\sigma_x, \sigma_y\}$, of the peaks were calculated for both schemes using synthetic images of ellipses generated according to the distributions mentioned. The geometric mean of the average standard deviations was calculated to give a measure of the size of the peak spreads ($\sigma = \sqrt{\sigma_x \sigma_y}$).




Figure 4.8: Two accumulators for two ellipse images, (a), are shown for uniform, (b), and weighted, (c), Hough voting schemes.

The limit of the peaks was determined by using accumulator elements that were at least one third of the peak height. The results were not sensitive to this choice—giving the same trend if a threshold of 25% or 40% had been used. The peak sizes were measured to be $\sigma_{uniform} = 3.05$ pixels and $\sigma_{weighted} = 2.64$ pixels. Thus, under the uniform voting method, the spread of the peaks was roughly 1.15 times larger than using the weighted voting method.

These results are certainly particular to the distribution of the ellipses being processed, but they indicate that there is an advantage to using a weighted voting scheme when the distribution of ellipse sizes is available. The narrower peak distribution prevents the peak from being obscured when there is noise present in the image. In order to obtain a confirmation of this idea for real nuclei, both the uniform and weighted voting strategies were applied to detecting ellipses in 212 images containing 431 overlapping nuclei. The

Classification	Uniform voting	Weighted voting
Good (< 20% error)	315 (73%)	354 (82%)
Marginal (20-40% error)	59 (14%)	45 (10%)
Failure (> 40% error)	57 (13%)	32 (7%)

Table 4.1: Comparison of uniform versus weighted voting for detection of 431 overlapping nuclei using the ellipse Hough transform.

results are shown in table 4.1.

These results will be discussed in more detail in section 4.4.3, where the classification scheme will be explained. At this point it is sufficient to note that the top row contains the nuclei that were considered to be successfully detected with the Hough transforms. Using the weighted voting scheme improved the recognition of the nuclei from 73% to 82%. From this we conclude that using the empirical distribution of object sizes to weight the voting in the Hough accumulator can improve the recognition of ellipses and hence, the recognition of nuclei in images.

4.4.2 Detecting peaks in the accumulator

The combination of discretization errors and edge orientation errors can cause ellipse center peaks to be spread over more than one cell in the accumulator $A(x_0, y_0)$. This can lead to multiple local peaks in the parameter space separated by only a few pixels. This is the case for the two nuclei shown figure 4.9a. The vote space is shown in figure 4.9b using brighter pixels to show accumulator cells with larger numbers of votes. Plotting the vote space as a three dimensional surface, (figure 4.9d), shows the seven local maxima. These false peaks are caused by votes for the true peaks that become spread over several cells and by edge orientation errors.

In order to suppress the false peaks and find the true center of the distribution a

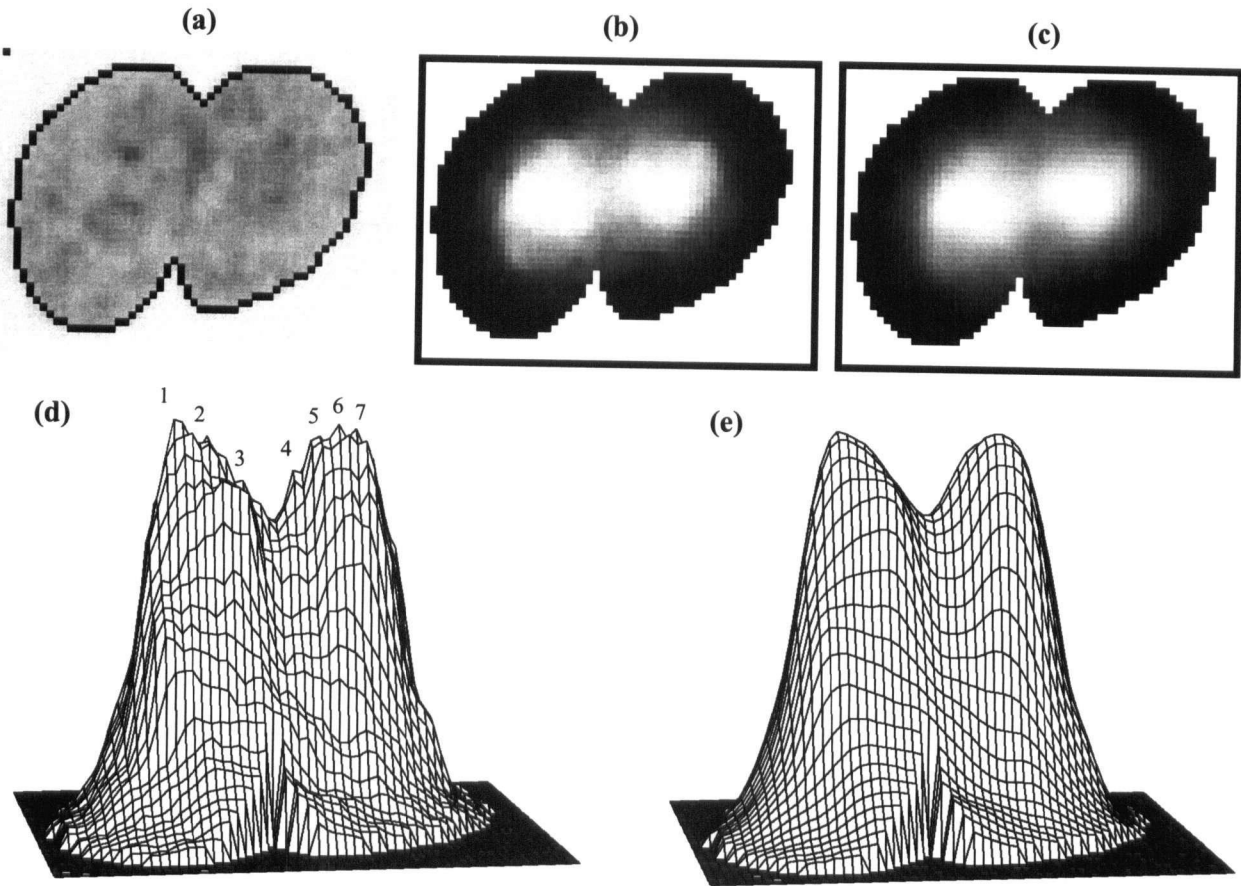


Figure 4.9: An image of overlapping nuclei, (a), the Hough vote space, (b), and the smoothing of the vote space using a Gaussian smoothing operator (c). The vote spaces plotted as surfaces (d,e) show seven and two voting peaks respectively.

Gaussian smoothing filter is applied to A . The size of the kernel standard deviation σ , depends on the expected size of the ellipses being sought. A value of σ that is too small can be easily detected as an excessive number of closely spaced peaks in $A(x_0, y_0)$ and can be adjusted accordingly. The application of a smoothing operator with $\sigma = .625$ to the vote space of figure 4.9b produces the vote spread shown in figure 4.9c. As can be seen from the three dimensional representation of smoothed voting space (4.9e), the seven peaks have been reduced to two.

The ellipse centres are recovered from the peaks of $A(x_0, y_0)$. Let the set of ellipse




Figure 4.10: The voting edge accumulators, B^i for: (a) the cluster of three nuclei in the top of figure 4.6, and (b) the two ellipse centres recovered from figure 4.9c. The brightness of each pixel corresponds to the number of times it voted (in conjunction with another edge) for the respective ellipse center.

center estimates be denoted by (x_0^j, y_0^j) . A second pass of the transform is executed with a new accumulator, $B^j(x, y)$, for each (x_0^j, y_0^j) . During this pass, cells of $B^j(x, y)$ corresponding to pairs of edges that vote for each (x_0^j, y_0^j) are incremented. This produces a set of accumulators that each contain the edges of a different ellipse.

The edge accumulators, B^1 , B^2 and B^3 for the cluster of nuclei in the top of figure 4.6 are shown in figure 4.10a. Similarly, B^1 and B^2 for the two ellipse centres of figure 4.9c are shown in figure 4.10b. The accumulator values are scaled so that the brightness of an edge is proportional to how often it voted for the respective ellipse center.

The thickness of the ellipse edges seen in figure 4.10 is controlled by the threshold on the gradient level of edges which are allowed to vote. Setting the threshold to a low value ensures that all true edges are used but also increases the amount of noise

present in the accumulator. These spurious edges can be reduced by setting the gradient threshold higher. Since the spurious edges tend to cast fewer votes (and appear darker in figure 4.10) they can also be removed by thresholding or applying an erosion operator to the edge accumulator, B^j . What remains is to recover the ellipse parameters from the edges.

4.4.3 Least squares recovery of ellipse parameters

Fitzgibbon [21] recently published an algorithm (B2AC algorithm) for least squares fitting of ellipses. The algorithm improves upon existing ellipse fitting techniques and is adapted for the problem at hand. Let the general equation of a conic be written as the implicit polynomial

$$F(\mathbf{a}, \mathbf{x}) = \mathbf{a} \cdot \mathbf{x} = ax^2 + bxy + cy^2 + dx + ey + f = 0 \quad (4.13)$$

with $\mathbf{a} = [a \ b \ c \ d \ e \ f]^T$ and let

$$\mathbf{x} = [x^2 \ xy \ y^2 \ x \ y \ 1]^T. \quad (4.14)$$

Given a set of N points \mathbf{x}_i , the general procedure for fitting a conic is to find a solution for \mathbf{a} that minimizes the squared error term

$$E(\mathbf{a}) = \sum_{i=1}^N F(\mathbf{x}_i, \mathbf{a})^2 \quad (4.15)$$

In order to obtain a unique, non-trivial solution for \mathbf{a} , it is necessary to apply constraints to the minimization problem. The salient difference between many published conic fitting algorithms is the nature of the constraint applied. The constraint may be linear, of the form $\mathbf{c} \cdot \mathbf{a} = 1$, such as $a + c = 1$, and $f = 1$. It may also be quadratic, of the form $\mathbf{a}^T \mathbf{C} \mathbf{a} = 1$ (where \mathbf{C} is a 6×6 constraint matrix), such as $|\mathbf{a}| = 1$ or $a^2 + \frac{1}{2}b^2 + c^2 = 1$.

Fitzgibbons discusses the properties of existing methods and notes that there is no “direct *specific* ellipse fitting” method. There are general conic fitting methods or iterative

methods that nudge a conic solution towards an ellipse. In principle, it is possible to force the solution vector, \mathbf{a} , of equation 4.15 to be an ellipse by specifying that the discriminant $b^2 - 4ac$ be negative. Problems with arbitrary scaling parameters are overcome by setting this condition as an equality constraint, $4ac - b^2 = 1$. In matrix form this quadratic constraint is written as

$$\mathbf{a}^T \mathbf{C} \mathbf{a} = \mathbf{a}^T \begin{bmatrix} 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \mathbf{a} = 1 \quad (4.16)$$

The ellipse fitting problem may now be written as a generalized eigenvalue problem. Let the matrix of edge point locations be written as

$$\mathbf{D} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \cdots \ \mathbf{x}_N]^T \quad (4.17)$$

Define the *scatter* matrix, \mathbf{S} , as $\mathbf{S} = \mathbf{D}^T \mathbf{D}$. Each element, $\mathbf{S}_{i,j}$, of the scatter matrix will therefore be

$$\mathbf{S}_{i,j} = \sum_k \mathbf{D}_{i,k}^T \mathbf{D}_{k,j} \quad (4.18)$$

Bookstein [15] showed that the problem of minimizing the squared difference error in equation 4.15 subject to the quadratic constraint in equation 4.16 was equivalent to solving a generalized eigenvalue of the form

$$\mathbf{S} \mathbf{a} = \lambda \mathbf{C} \mathbf{a} \quad (4.19)$$

where λ represents a generalized eigenvalue. Fitzgibbon's method uses this result to obtain the least squares estimates of the ellipse parameters from the eigensolutions to equation 4.19 subject to the constraint in equation 4.16.

If the pair $(\lambda_i, \mathbf{u}_i)$ is a solution to equation 4.19, the application of the constraint equation gives the conic parameters, $\hat{\mathbf{a}}_i$ as

$$\hat{\mathbf{a}}_i = \sqrt{\frac{1}{\mathbf{u}_i^T \mathbf{C} \mathbf{u}_i}} \mathbf{u}_i \quad (4.20)$$

Since \mathbf{S} is positive definite and \mathbf{C} is symmetric there will be only one positive generalized eigenvalue λ_i and therefore, only one set of solution parameters $\hat{\mathbf{a}}_i$. This solution is the parameters of the best fit ellipse to the edge points.

Figure 4.11 shows the B2AC algorithm applied to data points sampled from portions of the same ellipse. Each instance represents data subject to different noise conditions. The solid lines show the ellipse fit and the dotted lines show proper fit in the absence of noise. Figure 4.11a shows that the algorithm is able to recover the exact ellipse parameters given only 12 points sampled from a fraction (120°) of the border. Figures 4.11b and 4.11c show the gradual degradation of the fit as the sample points are subjected to increasing Gaussian noise.

In all three cases (figures 4.11a-c), the sampled data is seemingly consistent with either an ellipse or a parabola, given how little of the border is represented in the data. These panels show one of the benefits of this algorithm over other conic fitting techniques for recovering ellipses. It has a low eccentricity bias because the equality constraint 4.16 prevents the discriminant $b^2 - 4ac$ from approaching zero, which is the condition of an ellipse. This constrains the long axis of the fit ellipse from becoming too large even if the data is perfectly consistent with an ellipse. This is desirable for the problem of fitting nuclear images since often only portions of nuclear boundaries are present in the edge map.

Figure 4.11d shows how the algorithm performs given a heavily occluded border. The sample points are scattered in three groups and are subject to small Gaussian noise. Each group represents a roughly 30° portion of the border. Despite only having data from a

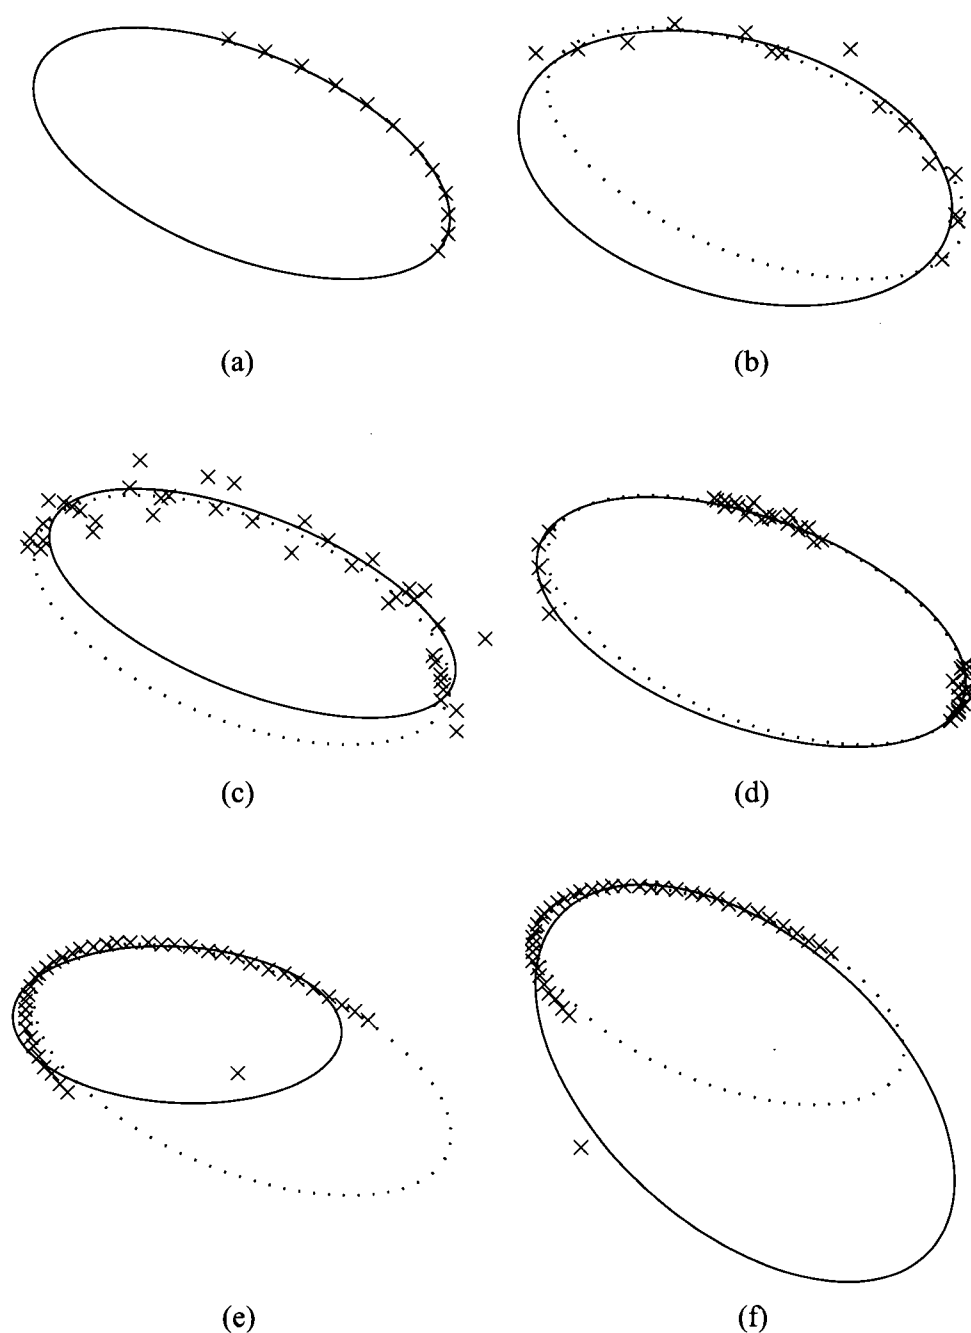


Figure 4.11: Fitzgibbon's B2AC algorithm for fitting ellipses using edge data. The scattered data used to generate the ellipse parameters are shown with \times 's. The points were selected from the same ellipse and subject to various noise levels. The solid lines show the ellipse fit and the dotted lines show the true ellipse in the absence of noise.

quarter of the border, the fit ellipse is nearly identical to the correct ellipse.

Figures 4.11e and 4.11f, demonstrate the limitations of the B2AC approach. Each contains 40 noise free points along the ellipse perimeter and a single outlier. Despite the numerical weight of the 40 points in the calculations, the ellipse fit borders are pulled in, (figure 4.11e), or extended, (figure 4.11f), to accommodate the extra point. The B2AC algorithm is a direct least squares method that gives each data point equal weight. It does not explicitly deal with outliers, which must be removed from the data set.

Removing the outliers is important when applying this algorithm as the final step of the ellipse Hough transform. As is shown in figure 4.10, the edge accumulator for nuclear images is very likely to contain some random false edges which can often be removed by thresholding or erosion of B . However, even after this is done, it is possible for a few stray edge pixels to remain.

We reduce the effect of these pixels by applying a modification to the B2AC algorithm to take the magnitude of edges in B into account. The original algorithm gives all points equal weight in the scatter matrix. It is possible to change the form of the scatter matrix equation 4.18 so that the magnitude of edge in B is used as a weight in a modified form of the equation. Essentially, the weight each pixel contributes to the ellipse parameter calculation is made proportional to the number of times it voted (in conjunction with others) for the ellipse center. The elements of the modified scatter matrix, S' , are given by

$$S'_{i,j} = \sum_k B(x,y) \mathbf{D}_{i,k}^T \mathbf{D}_{k,j} \quad (4.21)$$

where $B(x,y)$ is the magnitude of the edge accumulator for the pixel whose location (x,y) is encoded in \mathbf{D} (according to equations 4.14 and 4.17) as $x = \mathbf{D}_{k,4}$ and $y = \mathbf{D}_{k,5}$.

The modified B2AC algorithm provides the final component to the ellipse Hough transform that was used as part of the nuclear image segmentation system. Figure 4.12




Figure 4.12: The final result for the ellipse Hough transform applied to: (a) the nuclear clusters in figure 4.6a, and (b) the pair of nuclei in figure 4.9a.

shows the application of the ellipse Hough transform to the cluster images of figures 4.6a and 4.9a. The coloured outlines show the ellipses detected by the Hough transform. For one of the nuclei in the center of figure 4.12a, a portion of the red ellipse lies outside the nucleus. This reminds us that the Hough transform uses only the grey scale information in the image and does not use the mask determined by thresholding. The ellipse detected by the HT will often lie partly outside of the object being detected. This is a desirable feature if we are recovering an ellipse from an arc portion. In this application, where we are estimating the shape of a partly occluded nucleus, we can combine the HT ellipse estimate and mask information by performing a logical AND between the two.

The ellipse HT was applied to the cytology image set to test its ability to recognize overlapping nuclei. It was tested using both the uniform and weighted voting schemes described in section 4.4.1. The performance of the ellipse fits was measured by computing the percentage area misfit between the fit ellipse and the true object. The definition of




Figure 4.13: Definition of the measure of difference between the Hough ellipse fit and the true object, (a), and examples showing increasing error levels between the fit ellipses and true objects, (b).

this measure is shown in figure 4.13a. In essence, the error is the sum of the pixels that do not match divided by the total number of pixels for the two objects. If the objects coincide perfectly, the misfit is 0%. If they do not overlap at all the error is 100%.

Examples of error values for circles that overlap by different amounts is shown in figure 4.13b. In general, a value of $Err_A < 10\%$ is a very good fit and values less than 20% represent objects that can usually be recovered by refining their borders afterward. Values greater than 20% are poor fits and beyond 40% are considered outright failures.

Figure 4.14 shows some examples of the successes and failures of the ellipse HT using the weighted voting scheme. The Hough ellipses are superimposed over the grey scale images for the same image gallery that was shown in figure 1.1. The Err_A values are denoted by (r) or (g) to indicate whether they refer to the red or green ellipse. They were calculated using the ellipses and the manually segmented masks that are shown in

figure 1.1. In some of the cases only a single fit ellipse is shown and the error for the other ellipse is listed as 100%. These are occasions where the Hough transform found only a single peak or found a second peak that was too small to be considered a valid peak.

Figure 4.15 shows the graph of the cumulative number of nuclei detected versus the error level for the ellipse fits using the HT with uniform and weighted voting schemes. The difference between the two methods is also shown as the dashed line at the bottom of the graph. The difference is between 20 and 41, depending on the error level. We are interested in the number of nuclei that were adequately determined versus the number of failures. As mentioned earlier, ellipses with a roughly 20% misfit or less are considered to be adequate and those with an error of 40% or greater considered to be failures.

The HT results are summarized using these categories in table 4.1 in section 4.4.1. There are 39 more good fits using the weighted voting scheme—a 9% improvement. As well, the number of failures is 6% lower. The exact results are sensitive to the parameters used in the gradient threshold (section 4.4) and the parameters used to erode the edge accumulator, B .

The sensitivity of the results to different parameter combinations was tested by applying both methods to the image set with different parameter combinations. The gradient threshold percentage was allowed to vary from 50-100% of the gradient mean and the tolerance on the erosion of B was also allowed to vary. A total of 480 runs were performed. The uniform voting scheme had a mean success rate of 78% while the weighted voting scheme had a mean score of 85%. The maximum success rate for the uniform voting scheme was 83% while the maximum success rate for the weighted voting scheme was 90%.




Figure 4.14: The Hough transform applied to some the cultured cell nuclear images. The errors were calculated by comparing the fit ellipses to the manually segmented masks shown in figure 1.1. The (r) and (g) designations refer to the red and green ellipses.




Figure 4.15: The cumulative number of nuclei detected versus the percentage area misfit for the 431 nuclei in the cultured cell data set.

Chapter 5

Refining nuclear boundaries and assigning nuclear optical density

The segmentation procedures described so far give the position and approximate boundaries of the nuclei extracted from images of nuclear clusters. The goal of the segmentation process is to recover useful features from nuclear images. Two important types of features that are used in tissue section analysis are shape features and optical density related features. In order to be able to calculate these features, we need to recover the exact nuclear borders—which may overlap—and to apportion the optical density (OD) in the overlap regions.

This chapter describes how the border refinement can be done using active contour (AC) models and then how to divide the OD once the final borders have been determined. The AC model used in this thesis follows standard AC model techniques, but the novel element is that the model uses the watershed segmented objects and Hough transform ellipses from chapter 4 to simplify the search space for border optimization. This is discussed in section 5.1. In section 5.2, the hypothesis that the OD in overlap regions can be inferred from the OD in the non-overlap regions is studied. In section 5.3 a probabilistic model for reconstructing individual nuclear images from an overlap image is derived. Finally, a sequential procedure for removing nuclei from clusters is described in section 5.4.

5.1 Active contours

Active contour models are a powerful tool for image analysis resulting from the work of Kass, Witkin and Terzopolous [34]. Also called *snakes*, they are energy minimizing splines used to approximate features in an image. They represent a special case of matching a deformable model to an image by defining and minimizing an energy criterion.

This method defines the snake as a curve in space in parametric form as $\mathbf{v}(s) = (x(s), y(s))$ where $(x(s), y(s))$ are the coordinates along the contour and $0 \leq s \leq 1$. The energy of the snake is given by the equation

$$E_{\text{snake}} = \int_0^1 \left[\alpha(s) \left| \frac{d\mathbf{v}}{ds} \right|^2 + \beta(s) \left| \frac{d^2\mathbf{v}}{ds^2} \right|^2 + E_{\text{image}}(\mathbf{v}(s)) + E_{\text{con}}(\mathbf{v}(s)) \right] ds \quad (5.1)$$

where $\alpha(s)$ and $\beta(s)$ are adjustable parameters, $E_{\text{image}}(\mathbf{v}(s))$ is the energy functional of the image and $E_{\text{con}}(\mathbf{v}(s))$ are external constraints that incorporate the user's prior knowledge of snake shape.

The solution to an active contour problem requires adjusting the values of $\mathbf{v}(s)$ so as to minimize E_{snake} . As with other minimization problems the energy space is highly non-linear, and most snake implementations start with an initial configuration that is minimized using some sort of multivariate descent algorithm. A simple example of a snake is shown in figure 5.1a. This shows a snake, $\mathbf{v}(s)$ that is initialized to some random path, and the image data, which is a series of points. The optimal snake for this image is the one that matches the image points as closely as possible subject to the path constraints imposed by $\alpha(s)$ and $\beta(s)$.

The first term in equation 5.1 involves the derivative of the snake with respect to its path in space, $\frac{d\mathbf{v}}{ds}$. This term determines the elasticity of the snake. Making α very large emphasizes the $\frac{d\mathbf{v}}{ds}$ contributions in the total energy. Minimizing these causes the snake to behave as a series of straight segments, much like an elastic band stretched over




Figure 5.1: Fitting a snake to image data. The contour is initialized either randomly or to an estimate of the image data (a). Optimizing the snake results in one that fits the data subject to the relative sizes of elasticity constraint $\alpha(s)$, (b), and the stiffness constraint $\beta(s)$, (c).

a series of pegs. This situation is shown in figure 5.1b. If the ratio of α to β is made very large (perhaps by setting $\beta = 0$), the snake can become second-order discontinuous, having sharp corners.

The second term in equation 5.1 involves the second derivative of the snake with respect to the path, $\frac{d^2 \mathbf{v}}{ds^2}$. This term determines the degree of curvature of the spline fit to the data. Making $\beta(s)$ larger forces the snake to behave more like a stiff rod that can only be bent into smooth curves. This is illustrated in figure 5.1c.

The term $E_{\text{image}}(\mathbf{v}(s))$ contains the information regarding the image and the type of feature in the image that we are trying to fit. Its definition will depend on whether we are looking to trace lines, bright edges, intersect disparate points, etc. For the example shown in figure 5.1a, a suitable definition of E_{image} would be the sum of squares distance between the snake and the image points shown.

The most common task is to find edges of objects in the image, whose grey levels are denoted by $f(x, y)$. To find edges we define E_{image} as some function of the gradient, $\nabla f(x, y)$ (section 4.1.1). Kass et al [34] used the squared gradient with $E_{\text{edge}} = -|\nabla f(x, y)|^2$, but others have used the absolute value [27] with $E_{\text{image}} = -\gamma |\nabla f(x, y)|$. The difference between the two is that the former method puts more emphasis on matching strong edges at specific points rather than the edge as a whole.

The term $E_{\text{con}}(\mathbf{v}(s))$ contains the energy contribution from external constraints imposed by either the user or some other higher process which seeks to guide the snake's evolution into its final configuration. This term would express information about the desired final configuration of the snake. For example, it could encode a requirement such as whether the snake should look for a square or an ellipse, or whether the ends of the snake should be near or far apart. It could also be used to encode interactive forces such as an attempt by a software user to drag a portion of the contour.

Kass' method for generating a snake solution to an image processing problem starts with generating an initial snake approximation to the image feature being fitted. This is done by the user or by some other higher process to put the snake in the region of a solution. Selecting a good initial configuration is important since energy minimization is inherently a local optimization process. Snakes tend to evolve to a nearby contour (i.e. a local energy minimum). Kass' energy minimization method requires deriving the Euler equations for the derivatives of the snake coordinates in terms of the energy functionals. The position of each point of the snake is then evolved by taking a step of a predetermined size along the direction of the gradient. This is essentially a form of steepest descent minimization that nudges the snake with small steps into the final position.

Although the form of the original parameterization of $\mathbf{v}(s)$ and the snake energy (equation 5.1) is in a continuous form, it is more common to implement the snake as a planar

curve of N connected control points (or vertices) on the (x, y) grid, $\mathbf{v}_i = \{v_1, \dots, v_N\}$. The snake-energy is then redefined to replace the integral with a summation over the control points:

$$E_{\text{snake}} = \sum_{i=1}^N \alpha(v_i) |\nabla v_i|^2 + \beta(v_i) |\nabla^2 v_i|^2 + E_{\text{image}}(v_i) + E_{\text{con}}(v_i) \quad (5.2)$$

An important simplification can be made to the space of control points by restricting their locations to fall on a predefined grid. This restricts the nature of shapes that can be represented by the curve, but is a useful step if we have an idea of the shape we are looking for. Figures 5.2a and 5.2b show a continuous snake, and a planar curve representation of the snake lying on such a grid. Each element in the two dimensional grid represents the location of a point in the original image. We discretize the original curve $\mathbf{v}(s)$ into N points, each of which may take on M different values. The grid could be selected to consist of a block of consecutive pixels in the image, but one can also select points spread out in space and connect them with either straight line segments (shown as solid segments in the figure) or with a piecewise spline (dotted segments). The vertices are labelled v_j^k , where $1 \leq j \leq N$ and k is the iteration number of our optimization algorithm.

The snakes described so far have been instances of open contours, where the locations of the ends are not constrained. For imaging problems where we are trying to locate object boundaries, we are interested in closed contours—ones where the first and last point of the contours are connected. This is done by implementing circular boundary conditions on \mathbf{v}_i so that v_1 and v_N are considered to be adjacent when calculating the $|\nabla v_i|^2$ and $|\nabla^2 v_i|^2$ energy terms.

The active contour model used in this thesis follows the one used by Bamford [10] for segmenting individual nuclei in pap smears. In this formulation, $E_{\text{image}} = -|\nabla f(x, y)|$ and $E_{\text{con}} = 0$, and the elasticity factor, α , is set to zero. The image forces are determined




Figure 5.2: A continuous snake (a), and the planar curve representation (b) as a series of vertices on a predefined grid. The vertices may be connected by straight line segments (solid segments) or a piecewise spline (dotted segments).

by both the magnitude and direction of the image gradient by setting $E_{\text{image}} = -\nabla f(x, y)$. The curvature term, v'' , is calculated at each control point v_j using the central difference formula for the second derivative of \mathbf{v} :

$$v_j'' = \left(\frac{v_{j+1} - 2v_j + v_{j-1}}{v_{j+1} - v_{j-1}} \right)^2 \quad (5.3)$$

The overall energy of the snake is now written as

$$E_{\text{snake}} = \sum_{i=1}^N \lambda \left(\frac{v_{j+1} - 2v_j + v_{j-1}}{v_{j+1} - v_{j-1}} \right)^2 - (1 - \lambda) |\nabla f(x, y)| \quad (5.4)$$

where λ is a regularization parameter with $0 \leq \lambda \leq 1$.

Minimizing the snake energy in equation 5.4 forces the snake to follow strong edges in the image, while at the same time maintaining minimum contour curvature. The degree to which these two forces dominate the final contour is controlled by λ . If λ is small, the curvature minimizing term will dominate while as λ approaches one, the snake will follow the brightest points in the image without regard to shape.

In Bamford's experiments, the closed contours were optimized on a circularly symmetric grid centred on the nucleus. For example, figures 5.3a and 5.3b show the image of a nucleus from a tissue section image and a circular sampling grid superimposed over

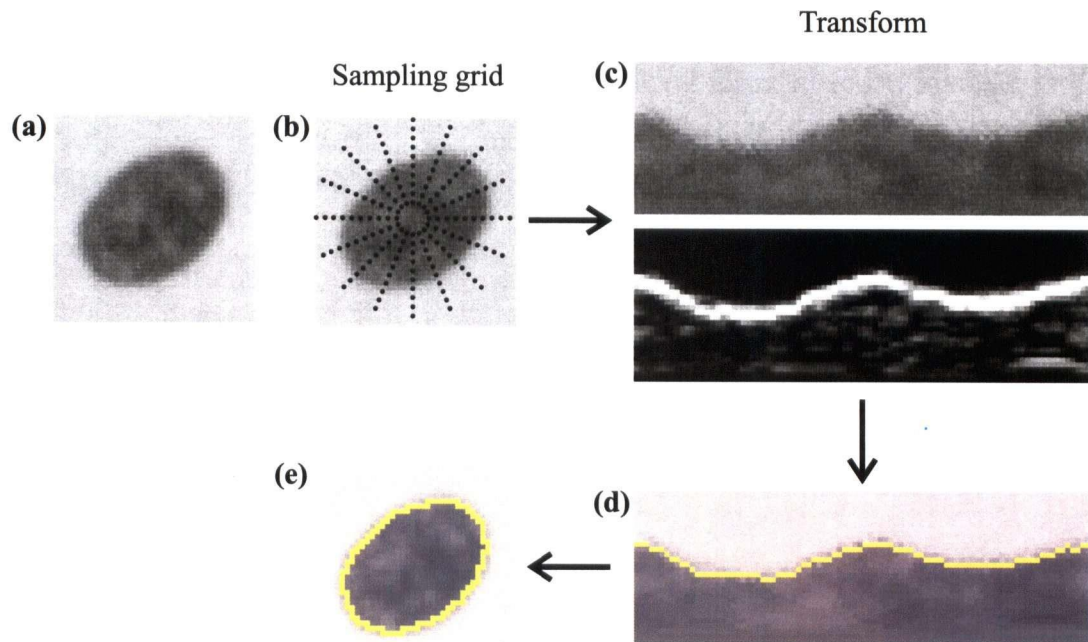


Figure 5.3: Transformation of a nuclear image, (a) using an circular sampling grid (b). The circular transform of the image and the vertical direction contrast based gradient are shown in (c). The nuclear border is a smooth curve in the transform space (c) Transforming back into the image space gives the segmented nucleus (d).

the image. If we sample the grey scale values in this fashion and unwind the grid into a linear strip we obtain the circular transform of the image shown on the right. The transform in figure 5.3c was created by sampling along 90 rays (4° increments) and is 30 pixels tall.

Figure 5.3c also shows the magnitude of the vertical direction contrast based gradient in the transform image. Calculating only the vertical component of the gradient in the transform image is equivalent to calculating the gradient along the grid lines in figure 5.3b. This makes the gradient image sensitive to edges that are perpendicular to the grid lines. If the nucleus is nearly circular, then its entire edge will be perpendicular to the grid lines, and the gradient image values will have maximal magnitude along the true nuclear edges. The benefit of using the directional derivatives will be shown for resolving the borders

of overlapping objects. A suitable snake fit to the transform image and the resulting segmented nucleus are shown in figures 5.3d and 5.3e.

Since nuclei are elliptical in shape using an elliptical sampling grid is more suitable to segmenting such nuclei. There are two main advantages to using an elliptical sampling grid. The first benefit of using the elliptical transform is that for elongated ellipses the directional gradient remains perpendicular to the nuclear edges. The second is that provided the ellipse fit is close to the true nuclear shape, then the border of nucleus will nearly be a straight line in the transform space. This is shown in figures 5.4a and 5.4b, which show an elliptical sampling grid for the same nucleus shown in figure 5.3a and the transform space for this image. In this space, a horizontal line is equivalent to a perfect ellipse in the image space. This means that minimizing the curvature component of the snake energy will favor snakes that are elliptical in shape—a desirable property.

The optimal nuclear border is nearly a straight line (figure 5.4c) in this space rather than a curve (figure 5.3d). When optimizing the border in the circular transform space there is competition between the curvature energy, which favors a circular border, and the gradient energy, which for this image favors an ellipse. In the elliptical transform space, there is less competition between these two forces. In general, the curvature and image gradient terms will compete to force the snake in different directions. Choosing the correct balance between these forces requires adjusting the λ parameter to suit the particular imaging problem.

The most difficult problem for active contour design is how to efficiently optimize the snake configuration to reach the energy minimum. As mentioned earlier, the energy space is highly non-linear, so most techniques use a variation of steepest descent to optimize the contour configuration. This approach does not guarantee the best possible solution since the size of the search space is large. Using the sampling grid approach of figure 5.2, where we have N vertices having M possible values, the number of paths in this space is




Figure 5.4: Transformation of a nuclear image of figure 5.3 using an elliptical sampling grid (a). The transformed image and vertical direction contrast based gradient are shown in (b). Using the elliptical transformation the nuclear border is nearly a straight line in the transform space, (c). The final segmentation of the nucleus is shown in (d).

N^M . For the transform spaces shown in figures 5.3 and 5.4, this amounts to 30^{90} paths.

The approach to contour optimization taken in this thesis is to use the watershed boundaries and Hough ellipse fits as the initial snake configuration. In the transform spaces in figures 5.3c and 5.4b this initial guess is a horizontal line across the transform. The snake vertices are then optimized one at a time so as to minimize the snake energy in equation 5.4 until a local energy minimum is found.

Using the initial border estimate to help direct the snake to a local energy minimum leads to a quick optimization step. However, due to the highly non-linear nature of the energy space it is not likely to lead to the optimal snake configuration. In essence we are making a gradient based descent in the energy space from a single starting point to a local minimum. However, if the Hough ellipses are accurate, the final snake configuration

is likely to be a good one. This trade-off between speed and the chance of obtaining poor segmentations is acceptable since there are many objects to be segmented in a typical tissue section image and an inadequate segmentation for some fraction of the objects is inevitable.

Figure 5.5 shows the active contour segmentation of an overlapping nucleus using the Hough transform ellipse fit as a starting point. The Hough ellipse has axes $a_{maj} = 20$ and $a_{min} = 15$ and orientation $\theta = 82^\circ$. This defines the sampling grid used to create the ellipse transform. In the transform space the ellipse fit appears as a horizontal line. Using the area misfit definition shown in figure 4.13, the ellipse fit error is $Err_A = 8\%$. Most of this error is due to the misfit in the overlap region, which is the dark region in the center of the transform space. The gradient image shows an interior edge that separates the overlap region from the other nucleus. Optimizing the initial contour produces the refined contour and hence the final segmentation shown at the bottom of the image. The misfit for the refined border is $Err_A = 2.4\%$.

When the degree of contact between overlapping nuclei is not too large the watershed can provide a good initial segmentation to which active contour refinement can be applied. Figure 5.6 shows such a scenario. The watershed provides the initial segmentation shown at the top of the figure. From this we calculate the best fit ellipses and use them to construct the elliptical transformation of the object and its initial border. Refining the contour produces the final segmentation shown at the bottom. For both refined nuclei $Err_A \approx 3\%$ compared to the watershed segmentation which had misfits of 6% and 8% for the yellow and red borders respectively.

Expressing the improvement in the segmentation simply as an area misfit percentage hides some of the important benefits of applying active contour refinement to the overlapping nuclear images. For example, while the watershed segmentation overestimates the one nuclear area by 6% and underestimates the other by 8%, the benefit of applying

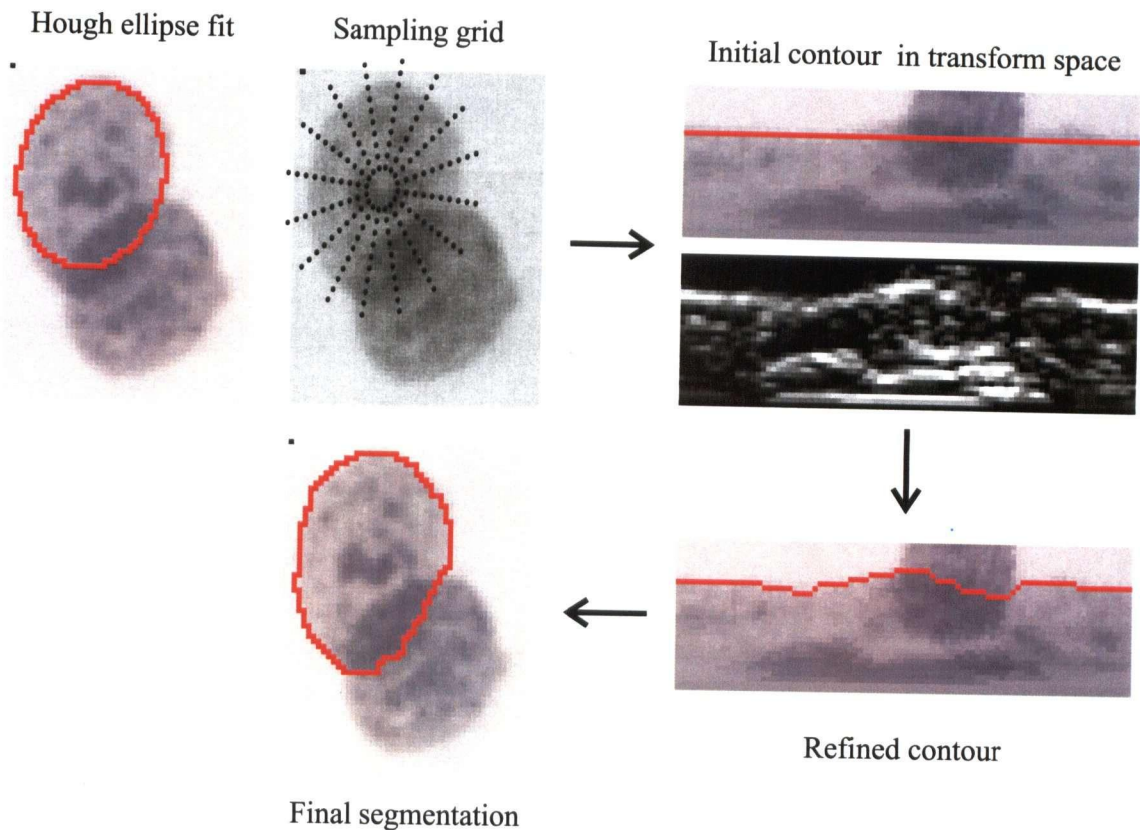


Figure 5.5: Active contour segmentation of an overlapping nucleus based on the Hough transform ellipse fit to the grey scale image.

active contours goes beyond the 3% and 5% improvement in Err_A . Having a segmentation which delineates the overlap region allows other important nuclear features to be calculated. For the segmented nuclei shown at the bottom of figure 5.6, we could also calculate Fourier shape features for the whole nuclei and grey level texture features using only the non-overlapping portions of the nuclei. We could not do this with any confidence for the watershed segmented images, since both the nuclear shape and grey level texture are compromised by the slightly incorrect segmentation.

Active contour refinement provides the final component for recovering nuclear shape for the segmentation system described in this thesis. An experiment was conducted to test




Figure 5.6: Active contour segmentation of two overlapping nuclei based on the watershed segmentation of the thresholded image.

the benefit of applying active contours to all segmented nuclei for the cytological image set. The images were initially segmented using the watershed algorithm. All watershed segmentations for which the occluded perimeter measure (equation 3.11), MP' , exceeded 30% were rejected. These objects were resegmented using the elliptical Hough transform procedure (section 4.4.3), using the combination of Hough transform parameters which identified 89% (384) of the 431 nuclei. Finally, the borders of all objects were refined using active contour refinement.

Figures 5.7a and 5.7b show histograms of the area misfit measure for the segmented nuclei before and after active contour refinement. The number of nuclei that fall into the range $Err_A < 20\%$ (the cut-off for acceptable segmentations) does not change as a result of the border refinement, but the average misfit is improved by refinement. The number of objects segmented correctly to within $Err_A < 5\%$ jumps from 110 to 349 after refinement. The average misfit value (for all nuclei with $Err_A < 20\%$) improves from $8 \pm 4\%$ before refinement to $4 \pm 2\%$ after refinement. Figure 5.7c shows some examples of the segmented images after the border refinement procedure has been applied.

5.2 Testing the optical density hypothesis

The second hypothesis (section 1.3) in this thesis is that in absorbance microscopy images of Thionin stained lung epithelial nuclei, the optical densities of overlapping nuclei are additive. In simpler terms, it is hypothesized that when two nuclei overlap, the overlap region appears darker, and we can predict how dark it should be. It is clear from examination of both cytologic and histologic images that the ODs appear to be additive. Nuclear overlap regions do appear to be darker than the non-overlap areas for the individual nuclei.

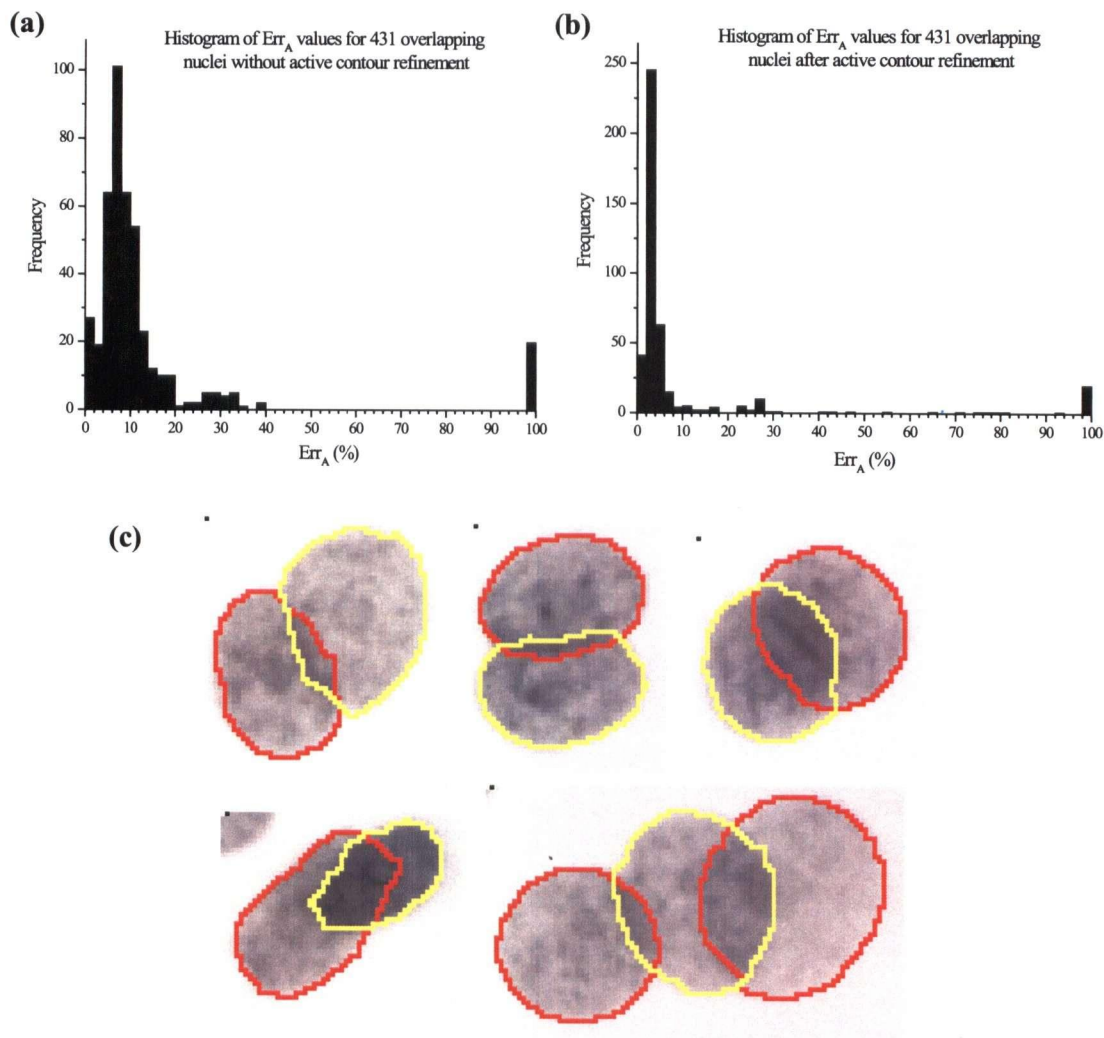


Figure 5.7: Histograms of the area misfit measure (Err_A) for the cultured nuclei (figure 1.1) prior to active contour border refinement, (a), and after border refinement, (b). The refined borders are shown for five of the images (c).

In order to test the hypothesis for cytological images, the optical densities were measured for the cultured nuclei image set (figure 1.1). The optical density was calculated for each pixel (x, y) as

$$\text{OD} = \log \left(\frac{B}{f(x, y)} \right) \quad (5.5)$$

where B is the background intensity for the image and $f(x, y)$ is the pixel intensity. When written this way, equation 5.5 assumes that there is no DC grey level offset that need be subtracted from numerator and denominator. The mean optical densities, $\overline{\text{OD}}$, were calculated for each of the non-overlap and overlap regions (figure 5.8a) by summing the pixel optical densities and dividing by the area of the regions. The entire image set consists of 431 nuclei with 219 regions of overlap. Some of the regions of overlap consisted of only a few pixels and were not considered reliable for analysis purposes. Only nuclear pairs for which the overlap exceeded 20 pixels (roughly 4% of the area of a typical nucleus) were analysed. This left 188 overlap pairs for which the ODs were measured.

Figure 5.8b shows a graph of the mean measured optical density in overlap regions, $\overline{\text{OD}}_C$, versus the sum of the means for the non-overlap regions, $\overline{\text{OD}}_A + \overline{\text{OD}}_B$. The data is scattered, but there is a linear trend visible. The linear regression to the data is plotted with the solid line. With the equation of the line written as $y = mx + b$, the parameters of the regression line are $m = .71 \pm .04$ and $b = .10 \pm .03$. The model prediction, that the optical densities sum linearly, is equivalent to plotting the line $y = x$. This is shown with the dotted line.

The regression slope reveals that the mean optical density in the overlap regions is only 71% of the predicted value—the overlap areas are not as dark as the model predicts them to be. The reason for this is likely related to the size of the nuclei in cytological preparations. As mentioned in section 1.4, the depth of field, d_f , of the imaging system

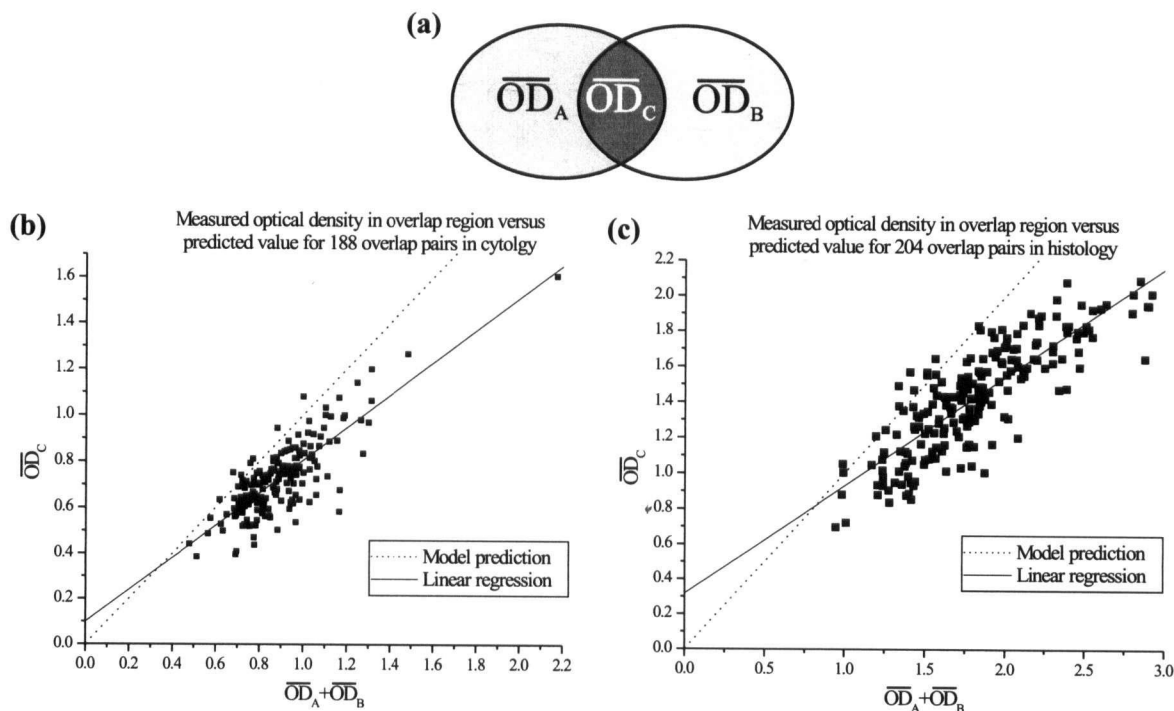


Figure 5.8: Experiments in measuring the mean optical density in overlap regions versus non-overlap regions. The non-overlap regions are designated A and B and the overlap region is designated C, (a). The graphs of measured mean OD in overlap regions versus the sum for non-overlap regions are shown for 188 overlap regions in cytological images, (b), and for 204 regions in biopsy sections, (c). The dotted lines represent the model prediction that $\overline{OD}_A + \overline{OD}_B = \overline{OD}_C$.

is approximately $1.5\mu\text{m}$. This means that when the microscope is adjusted to give the best possible focus of a cell, approximately $3\mu\text{m}$ depth¹ of stained material contributes to the optical density seen at a point. Since the region of overlap will be thicker than the regions containing any individual nuclei, less of the overlap region will be in focus than other portions of the nucleus.

The effect of changing the focal plane position on the measured optical density of an object can be seen in figure 5.9. Figure 5.9a shows an HL-60 nucleus imaged at 15

¹The axial intensity contributions depend on the square of the sinc function, and $2 \times d_f$ gives the distance between the first minima on either side of the central peak [31].

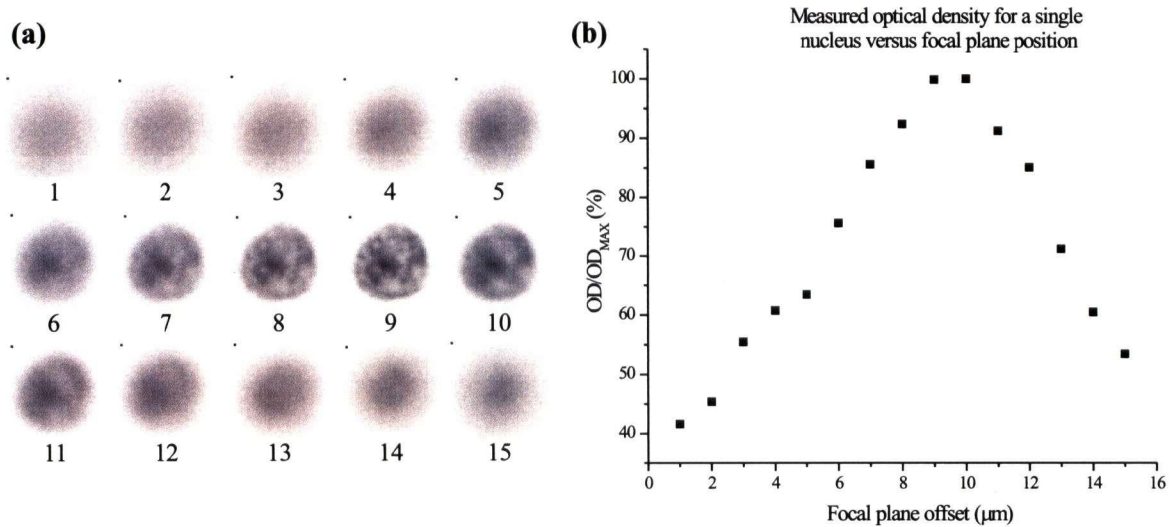


Figure 5.9: Fifteen images of a nucleus (a) captured at $1\mu\text{m}$ increments and the graph of total optical density versus the focal plane position, (b).

different focal positions, each separated by $1\mu\text{m}$ intervals. The best focus position occurs in the middle of the range, around the ninth image. Figure 5.9b shows the total optical density of the nucleus as a function of the focal plane position. The values have been normalized by dividing by the maximum OD measured, so the OD value at each point represents a fraction of the maximum OD.

As expected, the maximum OD measurements occur for the images in the best focus, in the region of image 9. When the focal plane is set at this position, the depth of field is filled with nuclear material that contributes to the measured OD. The OD drops off on either side of the peak and reaches 50% within $6\mu\text{m}$ on either side of the peak. At the extreme, for image 1, which is at least $8\mu\text{m}$ beneath the focal plane and lies beneath the nucleus itself, the OD value is still 40% of the peak value. Here, most of the nuclear material lies outside the depth of field and (due to the drop-off of the axial intensity function) does not contribute significantly to the measured OD. This is the case with the regions of overlap for overlapping nuclei. The portions of the overlap that lie above or




Figure 5.10: Images and manually segmented masks of overlapping nuclei from lung tissue section images. The light grey regions in the masks denote the regions of overlap.

below the depth of field do not make significant contributions to the measured optical density.

The optical density experiment was repeated for images of nuclei taken from biopsies sections. A total of 204 overlap regions were sampled from 41 images of the tissue section image set. Figure 5.10 shows some examples of the overlap regions excised from the images. Manually segmenting nuclei is a more difficult task for biopsy sections than for the cytological images. There is a significant amount of debris present in the image which makes it difficult to determine the best nuclear borders. As well, there are typically multiple overlaps in these images, so in some cases only the portions of nuclei were used for some measurements (as can be seen in the last overlap image in figure 5.10).

Figure 5.8c shows the graph of \overline{OD}_C versus $\overline{OD}_A + \overline{OD}_B$ for the tissue section images. There is more scatter of the data points than was seen in the cytology case, but a somewhat linear trend is present. As with the cytology case, the regression line to the

data has a slope that is less than the model prediction. For the histological segmentation, $m = .61 \pm .03$ and $b = .32 \pm .05$.

The 61% value for the mean OD in overlap regions is again partly due to the depth of field issue. The tissue section samples are all $7\mu\text{m}$ thick, so just under half of the depth of material lies in the region that is in focus. A second, and potentially greater effect, occurs from the cutting of nuclei that occurs in the biopsy sectioning process. Most of the nuclei in the biopsies exceed $7\mu\text{m}$ in diameter, so the act of cutting the $7\mu\text{m}$ slices from the biopsy sample cuts the nuclei into fragments. In areas of overlap, the portions of each of the two nuclei are necessarily thinner than the non-overlapping portions. This may account for the lower OD measured in the overlap portion of the image.

5.3 A model for assigning ODs for overlapping objects

Although there is some uncertainty in the OD measured in nuclear overlap regions, the results shown in figures 5.8a and 5.8b provide a means to estimate the optical density contributions of individual nuclei to regions of overlap. We can use the regression slope to $\overline{\text{OD}}_C$ versus $\overline{\text{OD}}_A + \overline{\text{OD}}_B$ as an empirical weighting factor to apply to OD values measured for overlapping pixels in order reconstruct the true OD contribution of each nucleus.

This is done by using a probability model for the distribution of OD values in each nucleus. When an OD measurement is made for a pixel for the overlap region the OD is apportioned to the two nuclei according to the probability model. In this manner, two individual nuclei are reconstructed with the overlap removed. The grey level values for what was formerly part of the overlap region will be synthetic, however they will represent the most likely values that could explain the OD measured at that pixel location.

The image reconstruction scheme is outlined in figure 5.11. Figures 5.11a and 5.11b

show two overlapping cell culture nuclei and the segmentation of the two objects. The non-overlapping portions of the nuclei are labelled A and B while the overlap is labelled C. Figure 5.11c shows the optical density histograms for regions A and B.

The reconstruction model assumes that pixel OD values are normally distributed, such that $OD_A \sim \mathcal{N}\{\mu_A, \sigma_A^2\}$ and $OD_B \sim \mathcal{N}\{\mu_B, \sigma_B^2\}$. The population means and standard deviations are not known, but they can be approximated by the sample means and standard deviations calculated from the optical density histograms. Assume we measure a pixel in region C to have an optical density equal to x^* . From section 5.2 we know that this OD value is too small by a factor of m (the slope of the graph in figures 5.8b and 5.8c). Let the corrected OD at this pixel be written as $x = \frac{x^*}{m}$.

This optical density value is the sum of the OD contributed by nucleus A and that contributed by nucleus B. Let the OD contributions be assigned as $OD_A = y$ and $OD_B = x - y$. Then the joint probability for this occurrence is written as $p(OD_A = y, OD_B = x - y)$. It is reasonable to assume that the probability distributions for OD_A and OD_B are independent. The darkness of one nucleus is not likely to depend on the darkness of another. In this case,

$$p(OD_A = y, OD_B = x - y) \propto p(OD_A = y) \cdot p(OD_B = x - y). \quad (5.6)$$

It can be shown (appendix B) that this joint probability distribution is of the normal form

$$p(OD_A = y, OD_B = x - y) \sim \mathcal{N}\left\{\frac{\mu_A \sigma_B^2 + x \sigma_A^2 - \mu_B \sigma_A^2}{\sigma_A^2 + \sigma_B^2}, \frac{\sigma_A^2 \sigma_B^2}{\sigma_A^2 + \sigma_B^2}\right\} \quad (5.7)$$

Equation 5.7 gives the probability density for the various ways that the pixel OD can be apportioned into components $OD_A = y$ and $OD_B = x - y$. The maximum likelihood (ML) estimate for OD_A is given by the value of y that maximizes equation 5.7. Since it is a normal distribution, the maximal value occurs when y is equal to the mean. Therefore

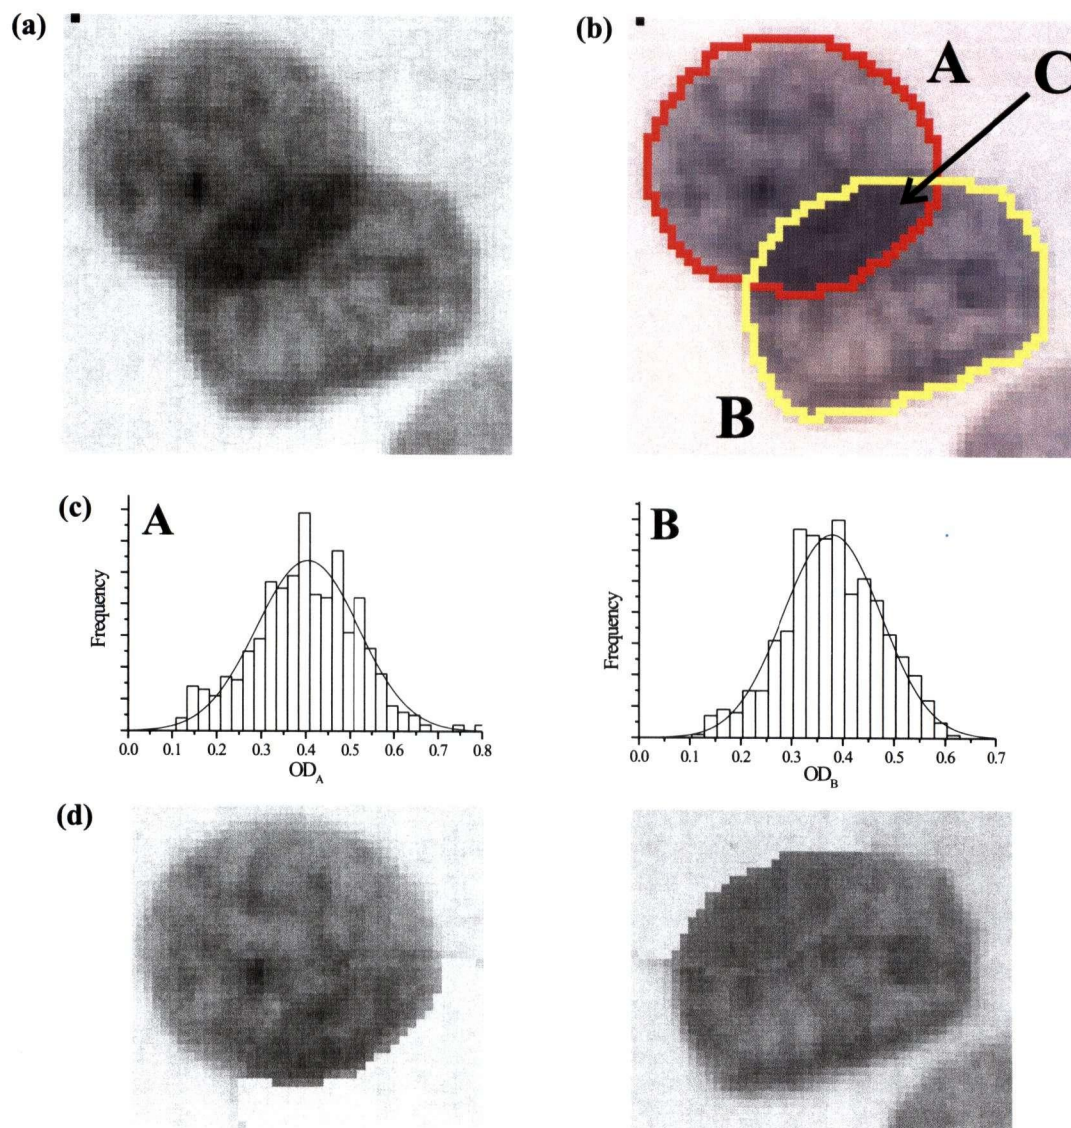


Figure 5.11: Demonstration of the image reconstruction method. The overlapping nuclei (a) are segmented (b), and the optical density histograms are calculated for the non-overlapping regions (c). The model assumes that pixel OD values are normally distributed and performs a Gaussian approximation to the OD histograms. The OD_C are then apportioned using this model giving the reconstructed nuclei (d).

the ML estimate of the reconstructed pixel OD for nucleus A is

$$\hat{y} = \frac{\mu_A \sigma_B^2 + x \sigma_A^2 - \mu_B \sigma_A^2}{\sigma_A^2 + \sigma_B^2}. \quad (5.8)$$

The ML estimate for the reconstructed pixel OD for nucleus B is

$$x - \hat{y} = x - \frac{\mu_A \sigma_B^2 + x \sigma_A^2 - \mu_B \sigma_A^2}{\sigma_A^2 + \sigma_B^2} = \frac{-\mu_A \sigma_B^2 + x \sigma_B^2 + \mu_B \sigma_A^2}{\sigma_A^2 + \sigma_B^2}. \quad (5.9)$$

Equations 5.8 and 5.9 are applied to each pixel in the overlap regions to reconstruct the two nuclei. Figure 5.11d shows the result of using this method to reconstruct the overlapping nuclei from figure 5.11b. The reconstruction was performed using the OD correction factor of $m=0.71$ for culture nuclei from section 5.2.

The reconstructed portions of the nuclei match the rest of the nuclei fairly well, but they exhibit less texture than the non-overlap portions. This occurs due the smaller grey level range seen in the overlap region in figure 5.11a. Figure 5.12 shows the grey level histograms for the non-overlap portions of the two nuclei (A and B) in figure 5.11b and for the overlap region (C). Region A has an intensity range from 85–186, a range of 101 grey levels while region B has a range from 113–187, a range of 74 grey levels. Region C has a smaller range of pixel values, from 80–116, a range of 36 grey levels. The reconstructed regions are calculated by converting an intensity value from region C into an OD and finding the ML estimates from equations 5.8 and 5.9. With only 36 different intensity values entering the equations, only 36 different grey level pairs result.

This means that if the goal of the reconstruction is to obtain the most reliable estimate of the optical density of the nuclei, then the texture must be sacrificed. An alternate plan would be to not use the ML estimates but rather choose a random $\{y, x - y\}$ pair according to the probability density in equation 5.7. This produces reconstructed pixel pairs (for A and B) that are not optimal in the sense that the values do not reflect the most likely assignment of the OD at that pixel. However, the method produces




Figure 5.12: Grey level histograms for the three regions of the nuclear image in figure 5.11b.

images that have more texture in the overlap region. It could be argued since the OD is apportioned for many such pixels, the cases where a pixel is given too much OD will be averaged out by the cases where another pixel is given too little. If this averaging effect were successful it would be possible to reconstruct a nucleus with both optimal total OD and approximate texture.

The possibility of assigning pixels in this fashion was explored, but the images produced were unsatisfactory as they had too large an intensity range in the overlap regions. Figure 5.13 shows the reconstruction of the two nuclei using this pixel assignment scheme. The unreasonably large intensity variations in the overlap portions occur because the normal approximations (seen in figure 5.11c) to the OD histograms are inadequate. Since the data is not normally distributed, the σ values calculated from the sample ODs tended to be large. This caused the tails of the OD probability densities to be overestimated,




Figure 5.13: Reconstructed nuclei from figure 5.11 with OD assignments made according to the joint probability density (equation 5.7) rather than using the ML estimates for that equation.

allowing very dark and very bright pixels to be selected during the image reconstruction. This produced too much texture variation in the overlap regions. A different model for the OD distributions would have to be used in order to be able to reconstruct texture as well as nuclear OD.

The result of performing the OD apportionment is that we now have an estimate of the total optical density of the individual nuclei. For the non-overlap regions in figure 5.11b, $OD_A = 265$ and $OD_B = 251$. The reconstructed regions contribute 85 and 81 units of OD to nuclei A and B respectively. The contributions are nearly equal since the OD means in figure 5.11c are nearly equal. Since the active contour refinement has produced nearly the exact borders for both nuclei the uncertainty in the reconstructed OD values for the nuclei is mostly determined by the uncertainty in the OD correction factor from section 5.2. Since the factor was $.71 \pm .04$, this implies that there is a roughly 6% uncertainty in the reconstructed portions. This gives a final total optical density measurement of $OD_A^{\text{tot}} = 350 \pm 5$ and $OD_B^{\text{tot}} = 332 \pm 5$.

The validity of the total OD value for the nuclei is subject to the applicability of the model of fitting a Gaussian to the non-overlap OD distributions. However, having

only a single image of the overlap, the only alternative to sharing the overlap OD is to ignore the information in the overlap and simply assigning the overlap OD based on OD measurements for the non-overlapping portions. The total OD value also depends on the correctness of the active contour refinement. An overlap region that is not recognized as overlap is treated as a dark portion for one of the nuclei, inflating its OD value. Reconstructing the overlap portion will then assign excess amounts of OD from the overlap region to the already dark nucleus. Despite these difficulties, analysing the OD in nuclear overlaps gives us the potential to extract information from nuclei which is currently not available. Having some nuclei with uncertainty on their feature values may prove to be more useful than not having them at all.

Reconstruction of isolated nuclei from overlapping images has also been studied by Ji and Tucker [32]. Their approach also assumes that the OD in non-overlap regions can be used to predict the OD in overlap regions. Their scheme is more comprehensive in the sense that they used images from multiple focal planes in the specimen to obtain the best focal plane for each portion of the overlap (rather than using a single image as is done in this thesis). However, their scheme simply subtracts a mean OD value from each pixel in the overlap region for the reconstructed nuclei rather than measuring the OD at each pixel and then deciding how to apportion it. In the future, reconstructing nuclei using a multi-focal plane approach is preferable since it can allow us to avoid such things as the empirical OD correction factor calculated in section 5.2.

5.4 A sequential approach to breaking up nuclear clusters

The degree of overlap of nuclei in tissue section images makes them difficult to separate, even with the techniques described so far. For small clusters, the Hough transform may simultaneously tell us the location of all the nuclear centres, as occurs for the top two

clusters in figure 4.6. However, it is more often the case that there are six or eight nuclei in a cluster (after the watershed transform has been applied) and the Hough transform reliably returns the locations of, say, two of them. This poses a problem if we wish to extract as many objects from the cluster as possible. Because this situation occurs often, a sequential approach was taken to the problem of breaking up nuclear clusters.

The sequential nuclear removal procedure is outlined in figure 5.14. When the HT is applied to a cluster, the ellipse with the largest number of votes in the accumulator space is processed first. For the cluster shown in figure 5.14a, nucleus C is processed first. Figure 5.14b shows the elliptical transformation of the Hough ellipse to nucleus C and its gradient. Although, the transformation space looks difficult to interpret, there is an interior edge through the overlap region that the active contour refinement was able to locate. The refined nuclear border is shown in figure 5.14c.

At this point, there is a problem with any attempt to segment what remains of the cluster. Removing one nucleus leaves a remainder containing most of nuclei A and B (figure 5.14d), which is missing some of its material from the overlap region with the nucleus C. The Hough transform is not likely to recover the center ellipse from the remainder, so an attempt was made to recover the overlap portions between the remainder and nucleus C. The two points where these two regions meet are labelled 1 and 2 in the figure. A rectangular grid is superimposed over the image between these two points (figure 5.14e) such that the two points lie on the center of opposite sides of the grid. Figure 5.14f shows the grey levels for the grid. A simple search is done for an open active contour path on this grid that is constrained to start at 1 and end at 2. The yellow line shows the best path for the cluster remainder. This becomes the new border for the clump remainder in this region (figure 5.14g).

We have now recovered the approximate shape for all the nuclei that remain in the cluster, so the procedure can be applied again to the remainder. When this is done for




Figure 5.14: Outline of the sequential procedure for removing nuclei from clumps. A small portion of the border of the clump remainder is refined after each nucleus is removed to allow more nuclei to be retrieved.

this example, we obtain the final border segmentation shown in figure 5.14h. The final borders overlap for nuclei B and C, so we apply the optical density apportioning model to obtain the reconstructed nuclei in figure 5.14i.

Chapter 6

Biopsy section processing experiments

The third hypothesis in this thesis is that it is possible to automatically segment tissue section images by recovering intact nuclei for the purpose of morphometric analysis, and generate results equivalent to those obtained by manual analysis of the image. Morphometric analysis is a standard research procedure for tumour grading and quantitating prognosis in cytopathology and histopathology [6, 7, 70]. When applied to biopsy sections, it requires the segmentation and analysis of the free lying nuclei. Since the nuclei in thin sections are cut, there tends to be a significant variability in the features calculated from the individual nuclei. However, if enough nuclei can be recovered from the tissue section image, useful statistics may be calculated for the sample as a whole.

The usefulness of morphometric analysis for the grading of lung tissue section images is currently being investigated at the BC Cancer Agency Cancer Imaging laboratory [37, 45]. In this study an attempt is made to create a grading system for lung sections that is reproducible and correlates with the patient diagnosis. For each tissue section image, Morphometric analysis results in a score called the morphometric index (MI) which represents the grading assigned to that image. The MI score falls into the range 1–10, where 1 represents a normal diagnosis, 5 represents a diagnosis of severe dysplasia, and a score approaching 10 represents a sample with a large population of abnormal cells, perhaps from a carcinoma *in situ*.

In the current procedure for manually analysing images, the sections are prepared in the manner described in section 1.4, and the following procedure is executed:

1. The area with the most abnormal appearance biopsy section is centred in the field of view under the microscope and an image is captured.
2. A mask is created for the image by setting the threshold at 165. This value has been determined in the past to be suitable for most tissue sections. It is not optimal since staining variations in the biopsies alter the best threshold. However, in the absence of an optimal thresholding method, it was decided to leave the threshold constant for all biopsy sections analysed.
3. A region of interest is drawn around the epithelial layer. For the current analysis scheme, it is necessary to ensure that objects from the subepithelium are excluded from the analysis.
4. Free lying, intact nuclei are selected. The decision as to whether a free lying object is relatively intact is based on operator experience. Consistency in the darkness across the extent of the nucleus is the most obvious feature to indicate that it is acceptable for analysis.
5. Nuclei that have minor contact with others are selected by digitally cutting a portion of the image from the rest. The number of nuclei collected varies widely. Although there are typically a few hundred present in an image, most tend to be clustered. The number collected varies between 20 and approximately 100, with 50 nuclei being the average.
6. The nuclear images are saved and their numerical features are calculated.
7. The nuclei (now represented as an array of feature values) are passed through a decision tree which sorts them into ten different categories from “normal appearing” to “abnormal appearing”.

8. The frequency of objects in these groups is calculated and the morphometric index (MI), a weighted function of the frequencies is calculated.

The decision tree in step 7 consists of a series of ten discriminant functions that sort nuclei into different categories based on increasing irregularity in their shape. The nuclear features that are used in the discriminant function are:

- **area**: the area of the nuclear mask.
- **mean_radius**: the mean distance in pixels from the mask centroid to the edge pixels of the mask.
- **max_radius**: the largest distance from the mask centroid to the edge pixels of the mask.
- **var_radius**: the normalized variance of the distribution of distances from mask centroid to the edge pixels of the mask.
- **sphericity**: the distance from the mask centroid to the nearest edge pixel divided by the distance from the mask centroid to the farthest edge pixel. These two values would be equal for a perfect circle, giving a maximal value of 1.
- **eccentricity**: the ratio of the short axis to the long axis of the best fit ellipse to the mask.
- **inertia_shape**: a roundness measure that is the ratio of the second moment of inertia of the mask to the square of its area.

Step 3 will pose a problem for the automated processing of tissue sections since the routines do not attempt to distinguish epithelium from subepithelium. However as can be seen in figure 2.3, many of the objects present in the subepithelium are either very

small and dark or very elongated. It may be possible to add a decision function to the automated system to exclude such objects.

The automated segmentation system described in chapters 2-5 replaces the manual effort in steps 4 and 5. When an image is segmented manually, the technologist specifically chooses intact nuclei that are in good focus and then segments only these nuclei. So far, the description of the segmentation system has not made any mention of incorporating a classification step of deciding what is a valid nucleus and what constitutes an adequate segmentation of that nucleus. A method was needed to process all the objects generated by the automated segmentation through a decision function. This function would sort them into either "valid" nuclei or junk categories. The "valid" nuclei category would consist of nuclei of all grades that were intact and in good focus. The junk category would be made up of all objects that must be eliminated before the images are processed by the MI decision tree in step 7. This would include fragments, clusters that could not be processed by the Hough transform, leukocytes, subepithelial structures, out-of-focus nuclei, etc.

The approach taken was to create a decision tree to separate nuclei from all other objects in the image frame. Figure 6.1 shows an outline of the complete decision tree system. The objects generated by automated segmentation are classified by a decision tree to separate valid objects from the rest. Only these objects will then be processed by the MI decision tree and eventually contribute to a MI score for the tissue section image.

In order to create the nuclear decision tree, a training set was created by selecting ten tissue section images from the image set of 226 histological images. The images were taken at random with the only stipulation being that no two images come from the same block of biopsy material—some patients will have multiple histological sections created from a single biopsy block. The ten images were then processed automatically and all the segmented objects were classified as valid nuclei or junk objects by a histotechnologist.

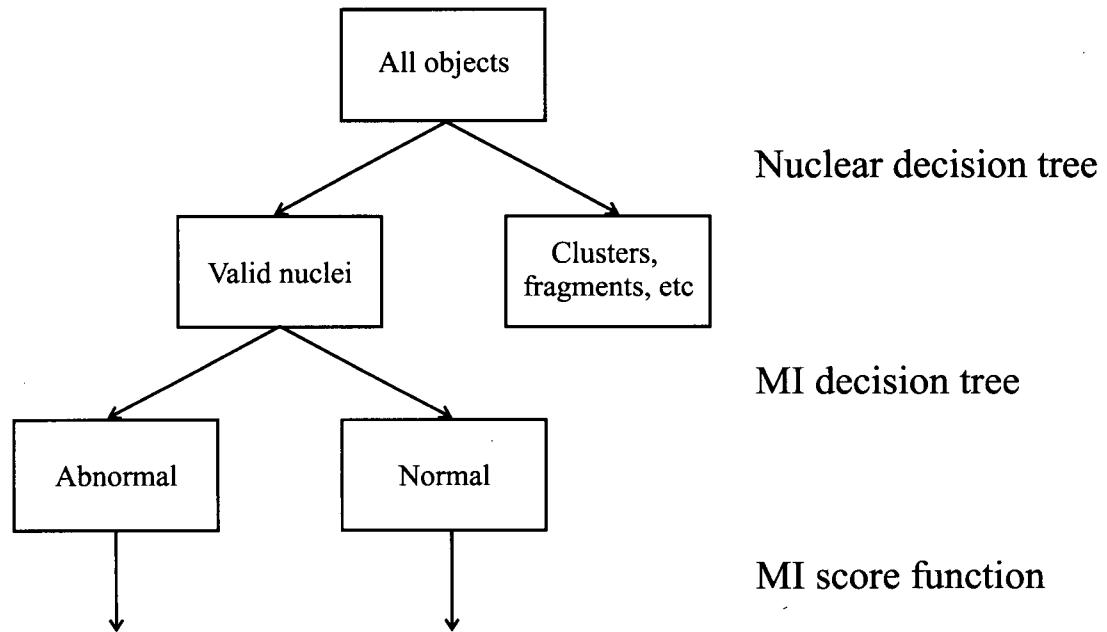


Figure 6.1: Outline of the decision tree steps in the automated MI scoring system.

A decision tree of thresholds and a discriminant function was then created to separate the valid nuclei from the rest.

The question arose as to what parameters should be used for the various components of the segmentation system in order to create the training set. One potential difficulty was that the generation of MI scores from the manual segmentation always used a threshold level of $T = 165$. Since the MI decision tree was created from these objects there was a concern that the objects generated through automatic thresholding (which tends to produce lower thresholds) would be smaller than those obtained with $T = 165$. This might lead to changes in the MI scores for these objects. Since they would tend to be smaller, this would drive the MI score downward. However, changes in other features such as the elongation of the cell with lower thresholds are somewhat unpredictable, so we could not be sure that the MI scores would change systematically.

Because it was not clear what thresholding procedure produces the best results in conjunction with the watershed splitting of nuclear clusters, it was decided to test five different thresholding/watershed schemes:

- T_{165} , WS: Threshold the image at grey level 165, then split clusters using the watershed algorithm
- T_{MG} , WS: Threshold the image at a level set by the mean gradient technique (section 2.2.3), then apply WS.
- T_{GW} , WS: Threshold the image at a level set by the gradient weighted technique (section 2.2.3), then apply WS.
- T_{MG} , WS, D_{165} : Apply T_{MG} , apply WS, then dilate the masks of the resulting objects until the borders reach grey level 165 (or meet another object).
- T_{GW} , WS, D_{165} : Apply T_{GW} , apply WS, then apply D_{165} .

For some objects, the schemes outlined above generate similar segmentations. However, the shape of the nuclear masks is highly dependent on the threshold level, so the watershed segmentation will often give different results for each scheme. For example, figure 6.2 shows the final objects that can result, (b-c), when a grey level image, (a), is segmented in different manners. Thresholding this image at a value of $T = 135$ and applying the watershed algorithm generates two objects (figure 6.2b), while setting $T = 165$ creates a mask which cannot be split with using the watershed (figure 6.2d). Thresholding the image at $T = 135$ and dilating the masks until the object borders reach grey level 165 generates the masks shown in figure 6.2c. Combined, the masks have the same shape as simply using $T = 165$, but the segmentation creates two objects instead of one. Although this example is artificial, it is indicative of the situations that occur when these procedures are applied to tissue sections.

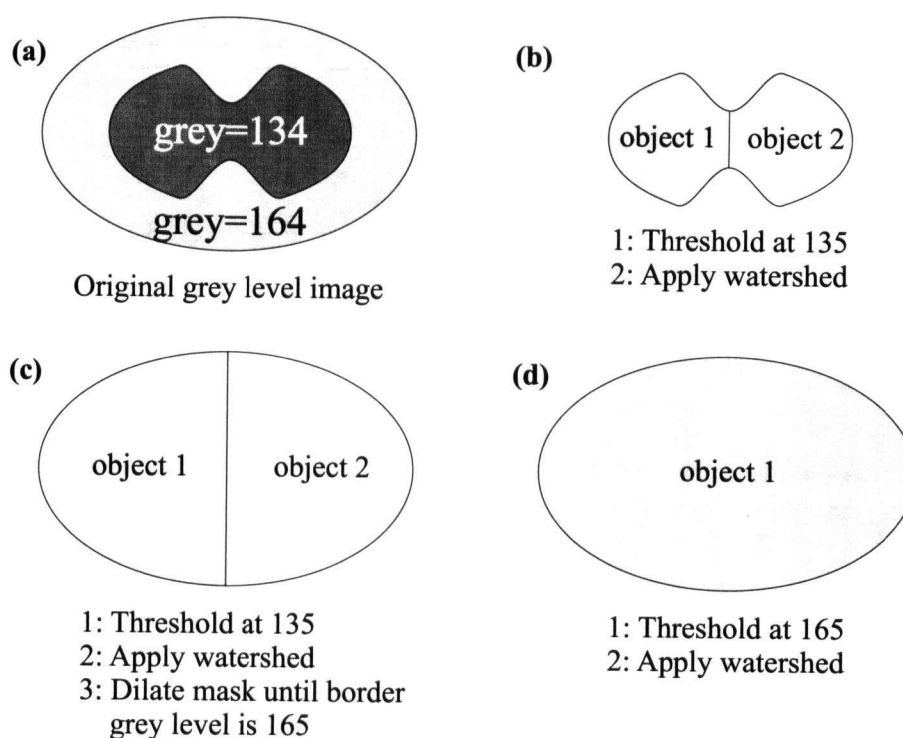


Figure 6.2: Segmenting a grey level image, (a), will lead to different final segmentations depending on the thresholding options used (b-d).

It was decided that to create the training set, two different segmentations of the same ten images would be used. For both segmentations, the mean gradient threshold was used, followed by the watershed splitting of the clusters. One set of images was left in this manner, while the nuclear masks for the other set were dilated until the borders reached grey level 165. This created a dataset of 12882 images. Each object was represented twice in the data with two slightly different masks—one obtained strictly by thresholding and the other by thresholding and dilation. The data was categorized as valid nuclei or junk by a histotechnologist creating a group of 1178 valid nuclei and 11704 junk objects.

A decision tree was created to attempt to reproduce the technologist's results on this data. The tree consisted of a series of feature value thresholds on discriminating features

True group	Called "valid"	Called "invalid"	Total	Percent correct
valid	1009	169	1178	85.6%
invalid	643	11061	11704	94.5%
Total	1652	11230	12882	93.7%

Table 6.1: Confusion matrix for the decision tree created to select valid nuclei for morphometric analysis.

such as

- shape features: to eliminate unsplit clusters and fragments
- texture features: to eliminate out-of-focus objects
- densitometric features: to eliminate segmentations that combined two objects with different ODs into one unit

After the extreme objects were removed, the final step of the tree used a discriminant function with 20 features, created with stepwise selection, to separate the nuclei from what remained. Table 6.1 shows the confusion matrix for the overall classifier on the training set.

The overall performance of the classifier was 93.7% on the 12882 objects. Although this score sounds like good performance for a classifier created on such experimental data, it means that nearly 40% (643 out of 1652) of the objects called "valid" by the classifier were objects that were rejected by the technologist. The technologist's standards for accepting an object were high, however. Many of those 643 incorrectly classified objects were ones that represented valid nuclei, but were rejected because their mask was improper. For these objects, either a portion of the mask had been chopped off by the watershed segmentation or small debris adjacent to the nucleus had caused the mask to be too large.

Because it was not practical to have a technologist manually categorize the nuclei for all the possible segmentation combinations, this classifier was used as the standard of truth by which the results of different segmentation experiments were compared. In addition to having five different thresholding options, four different cluster splitting options were examined:

1. no cluster splitting: This served as a baseline for the segmentation experiments. All clusters are treated as individual nuclei and passed through the decision tree. Under this scheme, free lying nuclei should be categorized as valid nuclei, so all the segmentation methods of this thesis should improve upon these results.
2. WS: The watershed algorithm is applied to all masks and the results are accepted unconditionally.
3. WS($MP' < 30\%$): After the watershed is applied, object pairs for which the contact perimeter measures $MP' \geq 30\%$ tend to be incorrectly segmented. These objects are remerged. This recovers some of the individual elongated nuclei that are often broken up by the watershed algorithm.
4. WS($MP' < 30\%$), HT+AC: After the methods outlined in option 3 have been applied, the Hough transform is used to further break up clusters and the active contour refinement is applied to all segmented objects in the frame. Finally, the optical density assignment model is applied to all segmented objects which have overlapping borders.

This gives a total of twenty different segmentation options. These were all applied to the 10 image training set and segmented objects were categorized using the decision tree. Table 6.2 summarizes all twenty methods and lists the total number of objects generated by the segmentation and the number of “valid” nuclei as given by the classifier. The

runs are split into four groups—one for each cluster segmentation option. The results for run 6 and run 7 are marked with an asterisk because the objects from these two segmentations were the ones used to create the decision tree. Since the decision tree was optimized to detect the nuclei from these two datasets there is the possibility that these two results could appear excessively good. However, the results of runs 6-10, which all use the same cluster segmentation option, are close enough that these two results do not stand out from the others.

Table 6.2 may be used to address many of the segmentation questions posed in this thesis. Treating the 20 runs as four groups, there is a clear increase in the number of “valid” nuclei obtained from the images. In the first group, there are an average of 3389 objects and 548 “valid” nuclei segmented from the images. Since there were ten images, this amounts to 55 nuclei per frame. This represents approximately the number of free lying nuclei that can be collected from a tissue section frame without significant extra effort.

Applying the watershed nearly doubles the number of objects in the dataset and increases the number of “valid” nuclei to an average of 815. If after applying the watershed we remerge objects that contact their neighbors significantly, we reduce the number of objects collected from 6503 to 5657. However, the number of “valid” nuclei rises to 833, on average. This shows that the MP' feature is a useful feature for detecting individual nuclei that have been incorrectly split by the watershed algorithm. Applying the Hough transform and active contour refinement we obtain 1020 “valid” nuclei from the decision tree. This is a 22% increase in the average number of objects obtained from the third group of runs, and nearly double the number of nuclei obtained from the simplest procedures with no cluster separation.

There is certainly an increased computational cost in applying the more sophisticated segmentation algorithms to segment the tissue section image frames. The complete image

Run	Segmentation scheme	Total number of objects	Number of "valid" nuclei
1	T_{MG}, D_{165}	3435	542
2	T_{MG}	3435	588
3	T_{GW}, D_{165}	3369	522
4	T_{GW}	3369	544
5	T_{165}	3341	548
1...5	Mean	3389	548
6	T_{MG}, WS, D_{165}	6442	780*
7	T_{MG}, WS	6442	872*
8	T_{GW}, WS, D_{165}	6525	786
9	T_{GW}, WS	6525	853
10	T_{165}, WS	6584	784
6...10	Mean	6503	815
11	$T_{MG}, WS(MP' < 30\%), D_{165}$	5577	789
12	$T_{MG}, WS(MP' < 30\%)$	5577	900
13	$T_{GW}, WS(MP' < 30\%), D_{165}$	5728	797
14	$T_{GW}, WS(MP' < 30\%)$	5728	891
15	$T_{165}, WS(MP' < 30\%)$	5677	788
11...15	Mean	5657	833
16	$T_{MG}, WS(MP' < 30\%), D_{165}, HT+AC$	5636	997
17	$T_{MG}, WS(MP' < 30\%), HT+AC$	5553	1050
18	$T_{GW}, WS(MP' < 30\%), D_{165}, HT+AC$	5625	1014
19	$T_{GW}, WS(MP' < 30\%), HT+AC$	5697	1032
20	$T_{165}, WS(MP' < 30\%), HT+AC$	5637	1009
16...20	Mean	5630	1020

Table 6.2: Results of the automated segmentation of the 10 tissue section image training set using different segmentation options. The number of valid nuclei is obtained by passing all the segmented objects through the nucleus decision tree.

Symbols:

T_{MG} - Mean gradient threshold selection

T_{GW} - Gradient weighted threshold selection

T_{165} - Set the grey level threshold to 165.

D_{165} - Dilate nuclear masks until border grey level reaches 165

WS - Watershed transformation to break up object clusters

$WS(MP' < 30\%)$ - WS then remerge contacting pairs of objects for which $MP' \geq 30\%$

HT+AC - Apply Hough transform then apply active contour border optimization

analysis system was written in Java and the experiments were run on a desktop computer equipped with an AMD Athlon 1.3Ghz processor. The Java programs were run using the standard Sun Microsystems Java2 version 1.3 runtime environment. Whichever segmentation scheme was used, the object masks required 10-15 seconds to be labelled. For the simplest segmentation schemes (runs 1-5) the thresholding process requires less than a second and the conditional dilation requires 6-8 seconds. The watershed algorithm requires approximately 15 seconds per image and the remerging step (runs 11-20) an additional 10 seconds. The Hough transform and active contour refinement were the computationally intensive operations, requiring roughly 10-12 seconds per object. A whole frame, with 500 or so objects, is processed in approximately ten minutes.

Comparing the results for the two gradient based thresholding schemes, there is virtually no difference between using T_{MG} or T_{GW} . Comparing the results of 1 and 2 versus 3 and 4, etc, T_{MG} generates a few more "valid" nuclei (1-4% more), but not enough to be significant. These experiments did not answer the question as to whether one of these thresholding methods was more suitable than the other. However, with the exception of run 4, both of these thresholding methods outperform the constant threshold of $T = 165$ for each group of runs. As well, the act of dilating the masks after the gradient threshold and watershed had been applied always leads to poorer results than not dilating at all. This can be seen by comparing runs 1 versus 2, 3 versus 4, etc. This suggests that the concern that it is necessary to dilate the masks in order to generate nuclear masks that resemble those obtained manually, may not be justified. The answer to this question can only be obtained by calculating morphometric index scores for an image and comparing to the MI score obtained manually.

The ultimate goal of these experiments was to automatically process section images and generate segmentations that would yield similar morphometric index scores to those obtained manually. The segmentation system and nuclear decision tree were tested using

Image #	Morphometric index
1	2.34 ± 0.26
2	3.90 ± 0.65
3	1.26 ± 0.03
4	2.46 ± 0.14
5	5.45 ± 0.43
6	3.47 ± 0.15
7	5.35 ± 0.26
8	1.24 ± 0.06
9	7.89 ± 0.34

Table 6.3: Morphometric index values for nine tissue section images obtained through manual analysis of the images. The uncertainty values are the standard deviations of the four MI values calculated on two separate manual segmentation attempts by two technologists.

a set of nine tissue section images that have been used as a validation set for morphometry experiments. These images have been manually segmented two times each by two technologists. Each of the segmentations was processed by the MI decision tree to sort them into ten categories of abnormality. Finally, category frequencies were calculated and passed into the MI scoring function to obtain four MI scores. Table 6.3 shows the MI means and standard deviations for the four measurements.

The images in the validation set were selected so that the set contained a wide variety of diagnostic grades. The MI scores in table 6.3 reflect this large diagnosis range. The scores range from a minimum of $MI=1.24$ to a maximum of $MI=7.89$. The standard deviations of the MI values tend to increase with the MI score. Since the scores are dependent on the relative frequency of irregular objects versus more regular objects, the score variations can occur because one technologist collects more irregular objects than another.

The nine images were processed using all twenty of the segmentation schemes outlined

in table 6.2. The “valid” nuclei were obtained using the nuclear decision tree and then the morphometric index tree was applied to measure the MI score for each image. The results for all twenty schemes were graphed for the nine images in figures 6.3, 6.4 and 6.5. For each graph, the mean manual MI score is plotted as a solid horizontal line and the standard deviations are plotted with dashed lines.

The results are mixed, since the automated MI scores lie within uncertainty of the manual ones for only three of the nine images (graphs 1, 3 and 8). For all the others, the automated MI scores are lower. The larger the manual MI score the greater the difference between those and the automatic ones. None of the automated scores exceed 5.0, while the manual scores exceed this value for three of the nine test cases. However, if one takes the uncertainties of the manual MI values into account, the automated MI value does tend to increase with diagnostic category. Figure 6.6 plots the automated MI scores versus manual MI scores. The automated ones form a spread of MI values for each manual one. The automated MI scores increase as the manual MI score increases. This behavior is desirable. The trend does not look monotonic, but some of the manual MI scores have large uncertainties. When these uncertainties are taken into account, the trend is monotonic. Even if the automated MI values are smaller than the manual ones, we would like them to increase monotonically with manual MI value. This would indicate that the automated scores have a useful diagnostic predictive value and only require a rescaling in order to compare them to the manual ones.

The MI score is dependent on the fraction of abnormal looking nuclei in the segmented population. These nuclei tend to be larger and have irregular shape. However, these two properties are also ones that separate the “valid” nuclei from the junk. So the nuclear decision tree is likely to throw out the very objects that the MI decision tree seeks to categorize as abnormal. This may explain why the automated scores seem reasonable for the lower grade images, but then diverge for the higher grade ones. Insufficient

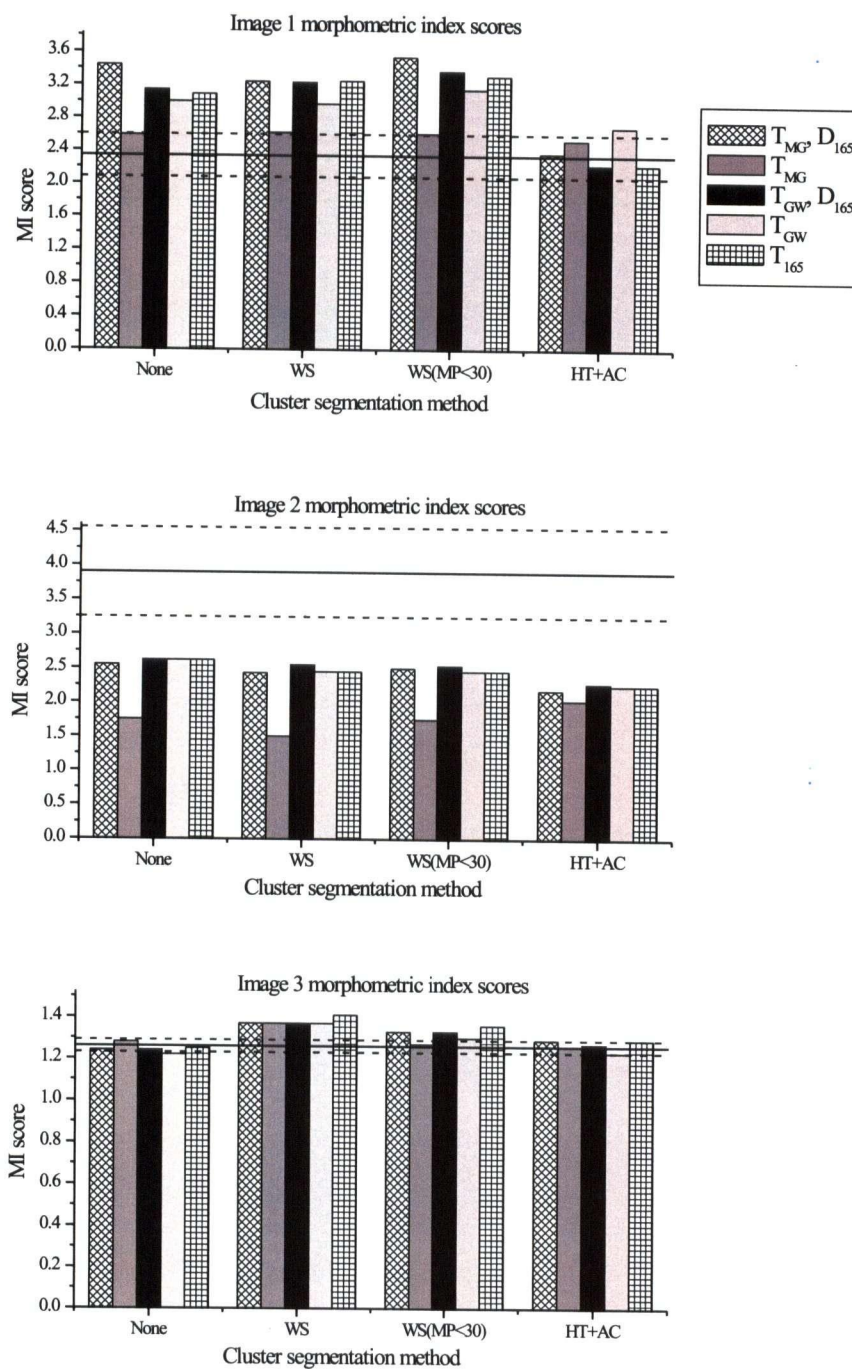


Figure 6.3: Automated morphometric index score calculations for test images 1, 2 and 3. The solid and dashed lines show the mean and standard deviation for the manual analysis of these images.

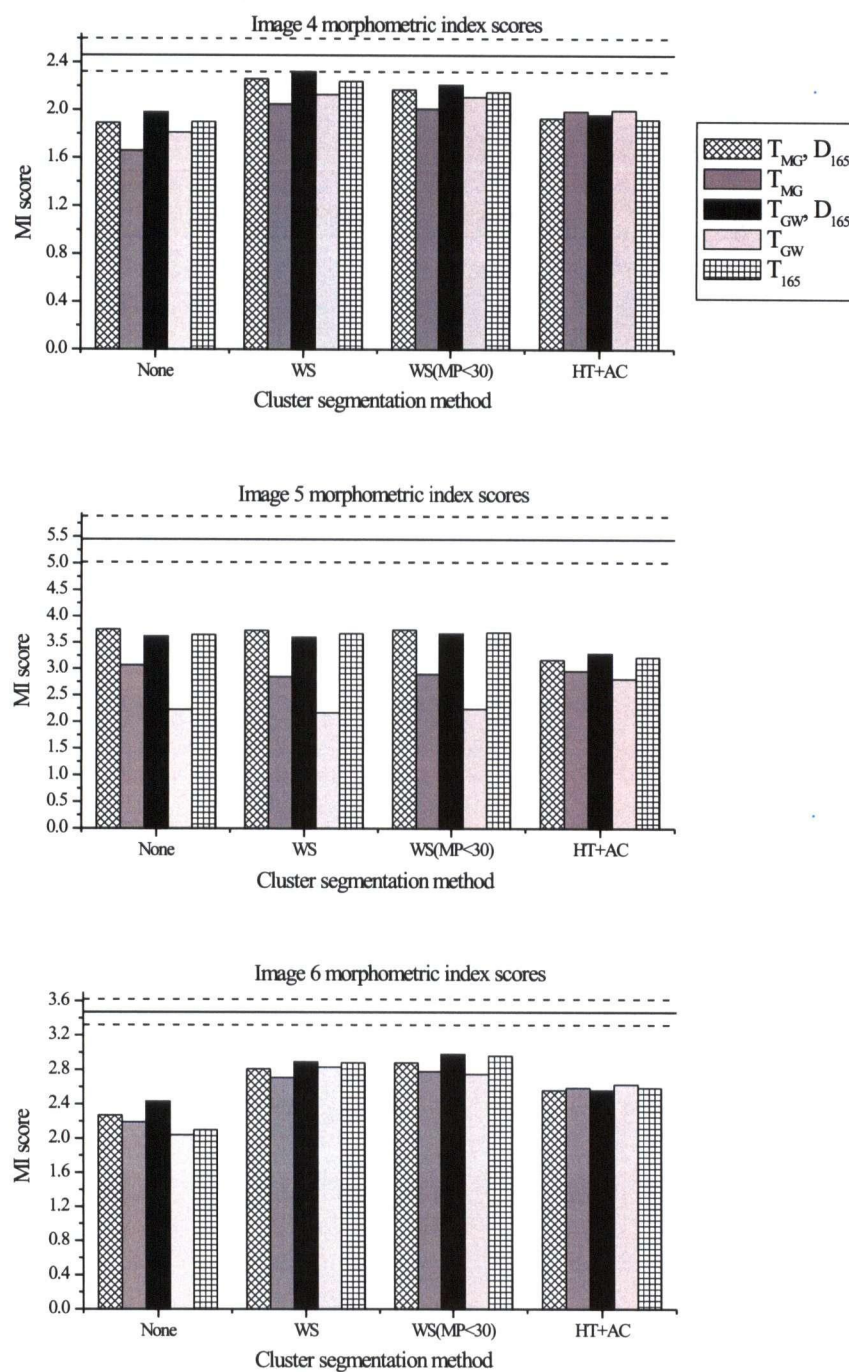


Figure 6.4: Automated morphometric index score calculations for test images 4, 5 and 6. The solid and dashed lines show the mean and standard deviation for the manual analysis of these images.

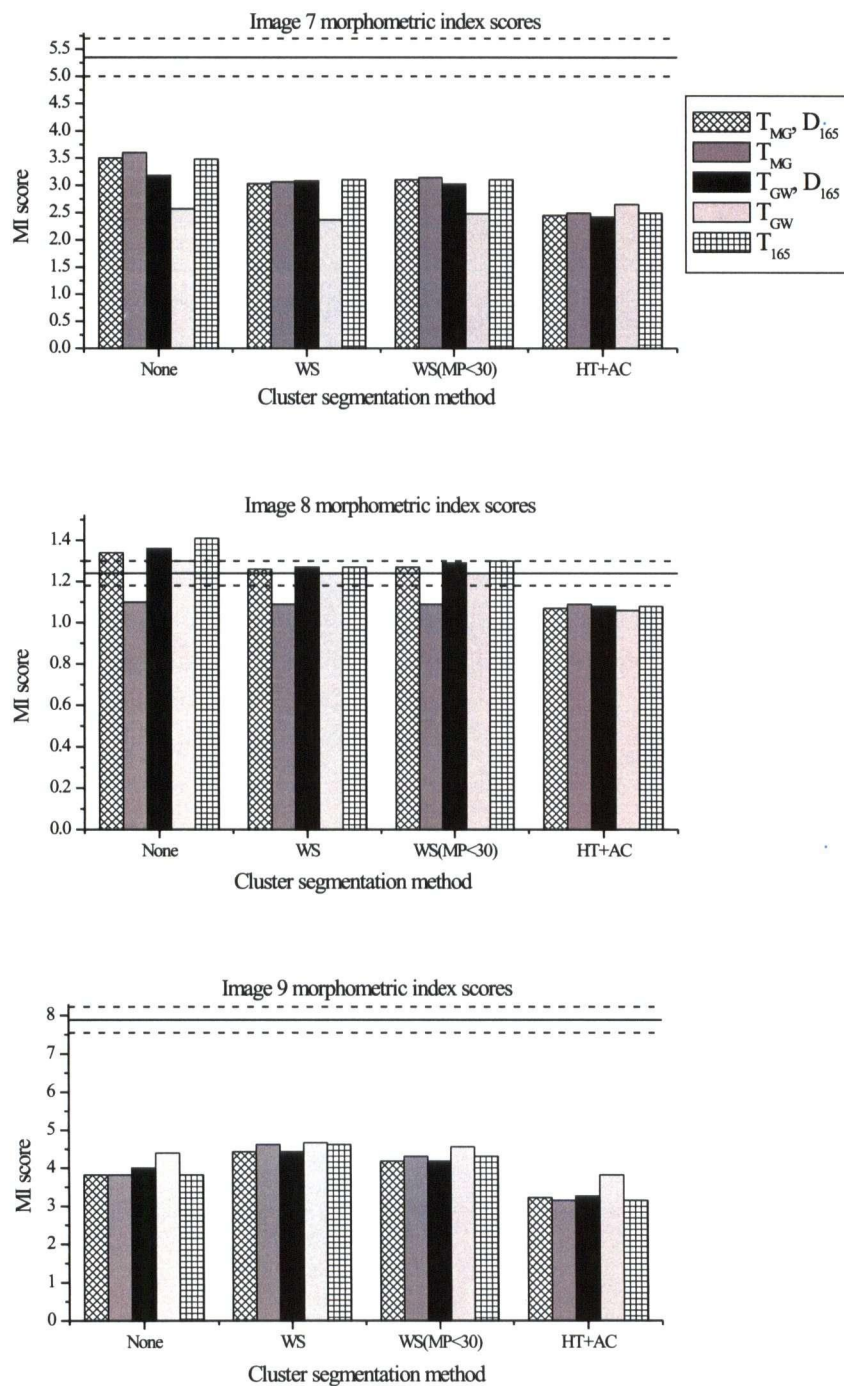


Figure 6.5: Automated morphometric index score calculations for test images 7, 8 and 9. The solid and dashed lines show the mean and standard deviation for the manual analysis of these images.




Figure 6.6: Comparison of the automatic morphometric index calculations to those obtained through manual analysis. The MI score for all 20 different segmentation methods are plotted for each image.

irregular nuclei are being considered by the MI decision tree. This suggests that the task of replacing the operator's discerning ability with a decision tree cannot be completely separated from what is done with the nuclei afterwards. Specifically, it may be necessary to create a decision rule that specifically separates abnormal nuclei from junk so that they are not accidentally thrown out before being analysed.

Another problem that may occur is that the automated segmentation procedure does a better job of segmenting normal nuclei than abnormal ones. Abnormal nuclei tend to have a more irregular shape (which may interfere with the watershed and Hough transform algorithms) and a more irregular optical density distribution (which may affect the active contour refinement routine). Consequently, it is more likely to generate good segmentations for normal nuclei than abnormal ones. These two difficulties both bias the

system to detect more normal nuclei than abnormal ones.

For all nine MI score graphs, there is a downward trend between the MI scores obtained with either watershed scheme versus those obtained by the full active contour segmentation. This occurs because the complete segmentation is finding more nuclei in the image and since it is biased towards finding normal ones, it generates final segmentations with a larger fraction of normal nuclei. This leads to a lower MI score.

The overall value of performing the complete segmentation with active contour refinement versus only selecting the free lying nuclei in an image should not be judged by the results in figures 6.3, 6.4 and 6.5, alone. These graphs appear to suggest no improvement in the MI score between no cluster segmentation versus the full system. However, the graphs are indicative of the performance of the segmentation plus the nuclear decision tree. The results of table 6.2 better reflect the merits of the different segmentation schemes. The number of normal appearing nuclei increases with the segmentation effort applied, with nearly twice as many nuclei being segmented by the complete scheme as compared to doing no cluster splitting.

Chapter 7

Conclusions

This thesis describes the development of an automated image segmentation system for resolving images of nuclei in crowded scenes such as tissue section images where objects tend to overlap. Chapter 2 looks at the problem of automated thresholding of grey level images. Several existing techniques were studied and determined to be unsuitable for automated biopsy analysis. Two possible thresholding schemes (which have been used successfully by others) appear to be well suited for automated biopsy analysis : the mean gradient threshold, T_{MG} , and the gradient weighted threshold, T_{GW} . Application of the two schemes to a set of 226 images revealed significant differences in the thresholds selected by these schemes. T_{GW} was adversely affected by image properties such as out-of-focus regions and overly bright texture. Both schemes were applied to the ten image training set and nine image test set in chapter 6. While they provided equivalent results, it is likely that the image sets of only ten and nine images were not large enough to reveal the differences between the thresholding methods.

The first hypothesis (section 1.3) proposes that it is possible to segment nuclei from complex image scenes by treating them as ellipses and using algorithms that focus on locating ellipses. Chapter 3 describes how mask shape can be used to separate overlapping nuclei by applying the watershed algorithm to the distance transform of the mask. A theoretical analysis was carried out to derive the formula for the minimum distance required to resolve two overlapping circles based on the exact and approximate distance transforms. Because this analysis was not tractable for arbitrary overlapping ellipses, a

Monte Carlo simulation was performed to measure the separability of ellipses (using the DT) as a function of their overlap.

Three different measures of overlap of pairs of ellipses are defined. The first measures overlap as a function of the ratio of overlapping area to the area of the smaller ellipse. The second measures the ratio of the occluded perimeter of an object to its total perimeter, and the third measures the ratio of the separation to the ellipse "size". The Monte Carlo simulations revealed that the perimeter measure performs the best of the three measures. It also showed that for the sizes of ellipses studied, the watershed segmentation should be rejected if the measure exceeded 30-40% for an object pair. Theoretical curves of these measures for overlapping circles were plotted and they displayed the same behavior as the curves for the Monte Carlo results. This confirmed that the perimeter occlusion measure gave the best distinction between correctly and incorrectly segmented object pairs.

The Monte Carlo study results also indicated that the ellipses used would be resolved better at higher magnification. This suggests that if mask shape is used to resolve overlapping elliptical objects, we should use the highest resolution available. Whichever overlap measure was used, overlapping objects were 10-20% more likely to be resolved at a particular overlap level if the magnification were doubled. Thus, the DT + watershed can be reliably used to break large clusters of elliptical objects into smaller ones and resolve small clusters when the overlap is not too large.

Chapter 4 introduces the Hough transform, a grey scale image technique that can be used to find parametric shapes in images. The method was adapted by Yuen to use pairs of edge points and geometry to estimate a line that passes through an ellipse center. His method was modified by implementing a weighted voting scheme based on prior knowledge of the distribution of ellipse sizes. An experiment was performed using ellipses whose size distribution matched those of a cytological image set containing 431 overlapping nuclei. The voting spreads for the uniform voting scheme were measured to

be 1.15 times larger on average than the weighted voting scheme. Larger voting spreads in the Hough accumulator space can lead to false minima occurring for complicated scenes. The uniform and weighted voting schemes were compared for the identification of the location and shape of the 431 cytological nuclei. The uniform voting scheme led to a correct ellipse fitting (of the nuclear shape) of 73% of the 431 nuclei while the weighted voting scheme scored 82%. This suggests that if the distribution of object sizes is known, it is advantageous to encode this information into the Hough voting procedure.

Yuen's center estimating technique was combined with Fitzgibbon's B2AC least squares fitting algorithm to determine ellipse parameters from their edge points. Fitzgibbon's algorithm is biased towards finding ellipses of low eccentricity. consequently, given only a small arc of points, the algorithm will return an ellipse with major and minor axes of comparable lengths rather than a parabola, which other conic fitting methods tend to do. However, since it is a least squares method, it is susceptible to outliers. Fitzgibbon's algorithm was modified to include a weighting factor for the data points in the ellipse parameter calculations. The voting frequency of points in the edge accumulator, B , were used as weighting factors. This modification, combined with a threshold applied to B , reduced the effect that noise had on the recovered ellipse parameters. The method was tested on the image set of 431 nuclei resulting in the 73% and 82% success rates described above. When the method was tested with 480 different combinations of parameters, the mean success rate for the voting scheme was 78%, while the mean success rate for the weighted voting scheme was 85%.

Section 5.1 describes how approximate ellipse fits returned by the Hough transform could be optimized using the active contour fitting model. The snake model used was the one described by Bamford, but the Hough ellipse fit was used in order to create an elliptically shaped sampling space. Using an elliptical grid rather than a circular one prevents competition between curvature and gradient forces when minimizing the snake's

total energy. The active contour border refinement was applied to the cytological image set. An area misfit measure was defined to describe the difference between the true nuclear mask and fits to that mask. This measure was calculated for the Hough ellipse fit to the true mask and for the active contour refinement based on the Hough ellipse. In an experiment, the number of nuclei for which the area misfit measure was less than 5% jumped from 110 to 349 nuclei after active contour refinement. The average area misfit for all nuclei considered to be successfully identified, dropped from $8 \pm 4\%$ to $4 \pm 2\%$ after active contour refinement.

Section 5.2 addresses the second hypothesis of this thesis—that optical densities in absorbance images of overlapping nuclei are additive. The hypothesis was tested by manually segmenting 188 overlapping nuclei pairs in cytological images and 204 overlapping nuclei pairs in histological images. A graph of measured mean OD amount versus predicted mean OD amount was plotted for each case. Both graphs had linear trends with a significant spread of the data. The spread was larger for the histological experiment. The regression parameters (for an equation expressed as $y = mx + b$) were $m = .71 \pm .04$ and $b = .10 \pm .03$ for the cytological measurements, and $m = .61 \pm .03$ and $b = .32 \pm .05$ for the histological measurements. This suggests that only a fraction of the expected optical density is measured in the overlap regions in either case. Several effects contribute to this. The depth of field of the imaging system is approximately $1.4\mu\text{m}$. Thus only part of an object will be in focus and contribute to the measured optical density. Since overlap regions in cytological images are thicker, this effect will be more pronounced in these regions. For histological sections, since the nuclei are cut, the portions of two nuclei that overlap will be thinner in the overlap regions than in the non-overlap regions.

Even though the mean OD in nuclear overlap regions is not a linear sum of the non-overlap regions, the regression slopes can be used as empirical factors for the purpose of reconstructing two individual nuclei from an image where they overlap. In section 5.3,

a probability model was introduced to apportion OD based on the assumption that nuclear OD is normally distributed and that the OD in nuclear overlap regions can be approximated by the OD from non-overlap regions. The joint distribution for the ODs of two individual nuclei was determined to be a Gaussian density that depended on the means and standard deviations of the ODs for each non-overlap region. The method reconstructed the images and apportioned the OD from the overlap region based on the maximum likelihood estimator for the joint distribution. This generates two reconstructed nuclei that give the most likely total OD at the expense of not approximating the nuclear texture as well as possible. The Gaussian model was inadequate for the purposes of approximating the nuclear texture as it produced reconstructed images with excessively large intensity variations.

Chapter 6 discussed the third hypothesis—that it is possible to automatically segment tissue sections and recover intact nuclei for morphometric analysis and generate results equivalent to those obtained by manual segmentation. Part of the hypothesis was successfully addressed in that the automated procedure recovers as many “valid” nuclei as are obtained through manual analysis. The segmentation system described in chapters 2 through 5 was applied to a training set of ten tissue section images with five different thresholding schemes and four levels of sophistication for the cluster splitting routines. Two of the 20 image sets were analysed by a histotechnologist, who separated 12882 segmented objects into 1178 valid nuclei and 11704 invalid objects. This data was used to create a decision tree to classify objects as “valid” or “invalid”. The classifier scored 93.7% on the training data. It was used to judge the success of all 20 segmentation runs on the training set. Extracting free lying nuclei without any cluster separation yielded an average of 55 “valid” nuclei per image frame. Performing the watershed segmentation of the image produced about approximately 83 nuclei on average. Applying the full cluster segmentation system including the Hough transform and active contour refinement (with

OD apportioning) produced 102 nuclei per image frame. This is almost twice the number of free lying nuclei and 22% more nuclei, per frame, than just using the watershed to segment the frame.

The segmentation system and “valid” nuclear decision tree were combined with the existing morphometric index tree being used in the Cancer Imaging department and applied to a set of nine tissue section images that had been manually analysed four times. The HI decision tree was designed to correlate with patient diagnosis, returning a score of 1 for a normal biopsy section and a score of 10 for an extremely abnormal section. The automated MI scores were calculated for all 20 segmentation option combinations and compared to the manual MI scores. The automated scores fell within the standard deviation of the manual scores for only three of the nine images. The automated MI values tended to be close to the manual ones for low MI values. For images with significant numbers of abnormal nuclei the manual MI values increased to nearly 8 diverging significantly from the automated scores. When the uncertainties of the manual scores are taken into account, the automated MI scores increase monotonically, which can be considered a significant success. This would imply that they need only be rescaled to compare them to the manual ones.

The reason for the lack of correlation of the manual and automated scores for high grade dysplasias is twofold. First, the segmentation system was designed to find regular elliptically shaped objects using the watershed and Hough transforms. Abnormal nuclei may have an irregular shape that interferes with the algorithms’ attempt to segment these objects. Abnormal nuclei may also have irregular OD distributions, sometimes having dark clumps inside the nucleus. This may cause the active contour refinement algorithm to incorrectly refine a border to lie along an interior edge inside the nucleus, thereby splitting an object in two.

The second, and likely more important reason, for the lack of correlation of MI scores

for high grade images concerns the interaction of the nuclear decision tree and the MI decision tree. The discerning ability of the technologist is an important element in obtaining good candidate nuclei for morphometric analysis. Ultimately, the reliability of the MI score is determined by the skill of the technologist. The approach taken in this thesis was to attempt to replace the technologist's skill with a decision tree to separate all nuclear shaped objects from the improper segmented objects and debris in the tissue section image frame. This decision tree used features such as the irregularity of the mask shape to separate "valid" nuclei from clusters and fragments. However, the image properties that cause an object to be rejected for further analysis were also the ones that the MI tree used to classify an image as abnormal. The two trees worked in opposition to each other, with the outcome that abnormal nuclei were not being used sufficiently to generate MI scores.

It is concluded that the process of collecting "valid" nuclei cannot be implemented without any reference to what they will be used for. The automated MI scoring procedure used in this thesis first separated "valid" nuclei from junk, and then used the existing MI tree to separate normal from abnormal looking nuclei. It should be possible to create a better automated MI scoring procedure by creating a single decision tree that separates out each category of nucleus from the remainder all the objects. This would prevent abnormal nuclei from being thrown out and improve the diagnostic value of the automated MI scores.

The work done in this thesis can be extended in several ways. The segmentation system, although largely successful, can be improved by making it more robust to the variability of object shapes in tissue section images and account for the effect of debris on the segmentation. This effort should focus on improving the Hough ellipse fits and exploring alternate active contour border refinement schemes. The Hough transform was most likely to fail on large elongated objects, a situation that occurs when two elliptical

objects are lined up along their long axes. The active contour routine was most likely to fail when small debris intersected the nucleus. This caused the method to either include the debris in the object mask, or exclude the portion of the object beneath the debris.

These kinds of errors suggest that it is necessary to have some kind of goodness of fit measure to compare alternate segmentations of the same scene. There are many ways to segment a nuclear cluster image that produce objects whose borders approximate the edges seen in the image. A goodness of fit measure based on the prior knowledge of shapes of nuclei and debris may allow for a better assessment of when a segmentation is to be trusted. This analysis would require further work into characterizing the objects that are found in epithelial biopsy section images.

The morphometric decision tree described in chapter 6 used only the shape of the nuclei to categorize their degree of irregularity. It did not use the optical density information that was also available. A recent study [63] has shown that densitometric information in histologic images is highly correlated with patient prognosis for oral squamous cell carcinomas. However, including densitometric information in quantitative histologic grading systems enters into a controversy [26] around the topic of DNA measurement in histologic sections. The problem is that in thin biopsy sections, most of the nuclei are sliced, so it is improper to speak of measuring its *ploidy*—the measure of DNA complement—distribution. Nonetheless, it is possible to model the sectioning of nuclei during the creation of tissue sections as the slicing of spherical objects [22]. One can then use the distribution of optical density values of the fragmented objects to estimate the true ploidy distribution.

Using the OD assignment method described in section 5.3, it is possible to reconstruct intact nuclei from images in which they overlap. Since we have an estimate of their OD, further work in the development of the MI decision tree may allow the use of OD as a feature to describe the nuclear irregularity. This may lead to the creation of a more

accurate system for grading epithelial biopsies.

Appendix A

Derivation of Overlap Measures for Circles

It is possible to derive explicit formulas for MA, MP and Ms for two overlapping circles of radii r_1 and r_2 . Figure A.1 shows the geometric construction needed to do this. In the figure, the separation of the two circles is $s = x + y$. The circle center locations are labelled C_1 and C_2 and the extent of each circle along their common axis is labelled W_1 and W_2 . The point at which their boundaries meet is P, and the angles are $\angle QC_1P = \alpha$ and $\angle QC_2P = \beta$.

The quantities r_1 , r_2 and s are known, and it is necessary to obtain expressions for x , y , α and β in terms of these. Looking at the triangles, we see that

$$x^2 = r_1^2 - h^2 \quad \text{and} \quad (s - x)^2 = r_2^2 - h^2.$$

After rearranging these equations, the expressions for x and $y = s - x$ are

$$x = \frac{s^2 + r_1^2 - r_2^2}{2s} \quad \text{and} \quad y = \frac{s^2 - r_1^2 + r_2^2}{2s}.$$

The angles α and β are then given by

$$\alpha = \arccos \frac{x}{r_1} \quad \text{and} \quad \beta = \arccos \frac{y}{r_2}$$

Obtaining a closed form expression for percentage of area overlap, MA, requires knowing the area of overlap, OA. The value of OA, is given by

$$OA = 2(A_{PQW_1} + A_{PQW_2})$$

where A_{PQW_1} and A_{PQW_2} are the areas of the two regions PQW_1 and PQW_2 in the figure. The factor of two is required since only the top half of each circle is shown. The area




Figure A.1: Construction for deriving the measures of overlap for two circles of different radii.

A_{PQW_1} is the difference between the area of the sector PC_1W_1 and the triangle PQC_1 .

This is given by

$$A_{PQW_1} = \frac{\alpha}{2} r_1^2 - \frac{1}{2} r_1^2 \cos \alpha \sin \alpha$$

where α is given in radians. Similarly, the area of the region A_{PQW_2} is given by

$$A_{PQW_2} = \frac{\beta}{2} r_2^2 - \frac{1}{2} r_2^2 \cos \beta \sin \beta.$$

Assuming that $r_2 \leq r_1$, the expression MA can now be written as

$$MA(r_1, r_2, s) = \frac{OA}{\pi r_2^2} = \frac{1}{\pi r_2^2} \left[\alpha r_1^2 - r_1^2 \cos \alpha \sin \alpha + \beta r_2^2 - r_2^2 \cos \beta \sin \beta \right]. \quad (A.1)$$

where the terms involving α and β are replaced using

$$\cos \alpha = \frac{s^2 + r_1^2 - r_2^2}{2sr_1} \quad \sin \alpha = \frac{\sqrt{r_1^2 - \left(\frac{s^2 + r_1^2 - r_2^2}{2s}\right)^2}}{r_1} \quad \text{and} \quad \alpha = \arccos \frac{s^2 + r_1^2 - r_2^2}{2sr_1}. \quad (A.2)$$

$$\cos \beta = \frac{s^2 - r_1^2 + r_2^2}{2sr_2} \quad \sin \beta = \frac{\sqrt{r_2^2 - \left(\frac{s^2 - r_1^2 + r_2^2}{2s}\right)^2}}{r_2} \quad \text{and} \quad \beta = \arccos \frac{s^2 - r_1^2 + r_2^2}{2sr_2}. \quad (A.3)$$

The formula for the percentage of perimeter occluded, MP, is more straightforward to calculate. Again, under the assumption that $r_2 \leq r_1$, circle 2 will necessarily have the larger MP_i value. The arc length of the occluded portion of circle 2 will be twice that of the chord PW_2 since only the top half of the circle is shown. Hence, OP_2 is given by

$$OP_2 = 2\beta r_2$$

with β , in radians, is specified in equation A.3. This gives

$$MP(r_1, r_2, s) = \frac{OP_2}{2\pi r_2} = \frac{1}{\pi} \arccos \frac{s^2 - r_1^2 + r_2^2}{2sr_2} \quad (\text{A.4})$$

The normalized measure of center to center separation, Ms, is the simplest to calculate. When $r_1 \leq r_2$,

$$Ms(r_1, r_2, s) = \frac{s}{r_2}. \quad (\text{A.5})$$

Appendix B

Derivation Optical Density Distribution for Overlapping Objects

The probability densities for the ODs in the non-overlap regions of figure 5.11b are given by

$$p(\text{OD}_A = y) = \frac{1}{\sqrt{2\pi}\sigma_A} \exp \left[\frac{-(y - \mu_A)^2}{2\sigma_A^2} \right] \quad (\text{B.6})$$

and

$$p(\text{OD}_B = x - y) = \frac{1}{\sqrt{2\pi}\sigma_B} \exp \left[\frac{-(x - y - \mu_B)^2}{2\sigma_B^2} \right] \quad (\text{B.7})$$

From equation 5.7,

$$p(\text{OD}_A = y, \text{OD}_B = x - y) = kp(\text{OD}_A = y)p(\text{OD}_B = x - y) \quad (\text{B.8})$$

$$= \frac{k}{2\pi\sigma_A\sigma_B} \exp \left[-\frac{(y - \mu_A)^2}{2\sigma_A^2} - \frac{(x - y - \mu_B)^2}{2\sigma_B^2} \right] \quad (\text{B.9})$$

where k is a normalization constant. The values of OD_A and OD_B are linked in that they must sum up to x , so this joint probability function is slightly different than the standard product of Gaussians.

Let $v_A = \sigma_A^2$, $v_B = \sigma_B^2$ the term inside the brackets in equation B.9 be labelled $*$.

Then

$$* = -\frac{1}{2v_A v_B} \left[(y - \mu_A)^2 v_B + (x - y - \mu_B)^2 v_A \right]. \quad (\text{B.10})$$

Expanding equation B.10 gives

$$* = -\frac{1}{2(v_A v_B)} \left[(v_A + v_B)^2 - (2\mu_A v_B + 2x v_A - 2\mu_B v_A)y + \mu_A^2 + x^2 v_A - 2x\mu_B v_A + \mu_B^2 v_A \right] \quad (\text{B.11})$$

Let $k_1 = v_A + v_B$ and $k_2 = \mu_A v_B + x v_A - \mu_B v_A$. Rewriting equation B.11 and removing those terms that can be included in the normalization constant, it can be written as

$$* = \frac{-k_1}{2v_A v_B} \left[y - \frac{k_2}{k_1} \right]^2 \quad (\text{B.12})$$

Equation B.12 is the exponent of a normal probability density with mean

$$\mu = \frac{k_2}{k_1} = \frac{\mu_A v_B + x v_A - \mu_B v_A}{v_A + v_B} \quad (\text{B.13})$$

and standard deviation

$$\sigma = \left[\frac{v_A v_B}{v_A + v_B} \right]^{\frac{1}{2}} = \frac{\sigma_A \sigma_B}{\sqrt{\sigma_A^2 + \sigma_B^2}}. \quad (\text{B.14})$$

These parameters specify the joint density for the OD values in the overlap region in equation 5.7.

Bibliography

- [1] Abmayr, W., L. Abele, J. Kugler, H. Borst(1980) Capabilities of a Nonlinear Gradient and a Thresholding Algorithm for the Segmentation of Papanicolaou-Stained Cervical Cells. *Analytical and Quantitative Cytology*, **2**(3), 221–233.
- [2] Abmayr, W., E. Mannweiler, D. Oesterle and E. Deml(1987) Segmentation of Scenes in Tissue Sections. *Analytical and Quantitative Cytology and Histology*, **9**(3), 190–196.
- [3] Awcock, G.W., and R. Thomas(1996) *Applied Image Processing*, McGraw-Hill, New York.
- [4] American Cancer Society(1996) *Cancer Facts and Figures-1996*, American Cancer Society, Atlanta.
- [5] Auerbach, O., E.C. Hammond, L. Garfinkel(1979) Changes in bronchial epithelium in relation to cigarette smoking. *New England Journal of Medicine*, **300**, 381–385.
- [6] Baak, J.P.A. and J. Oort(1983) *Morphometry in Diagnostic Pathology*, Springer Verlag, Berlin.
- [7] Bacus, J.W., C.W. Boone, J.V. Bacus, M. Follen, G.J. Kelloff, V. Kagan and S. Lippman(1999) Image Mophometric Nuclear Grading of Intraepithelial Neoplastic Lesions with Applications to Cancer Chemoprevention Trials. *Cancer Epidemiology, Biomarkers and Prevention*, **8**, 1087–1094.
- [8] Ballard, D.H.(1981) Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition*, **13**(2), 111–122.
- [9] Ballard, D.H., and C.M. Brown(1982) *Computer Vision*, Prentice-Hall Inc, New Jersey.
- [10] Bamford, P.C.(1999)*Segmentation of Cell Images with Application to Cervical Cancer Screening*, Ph.D. Thesis, University of Queensland, Brisbane, Queensland, Australia.
- [11] Bengtsson, E., O. Eriksson, J. Holmquist, B. Nordin and B. Stenkvist(1979) High Resolution Segmentation of Cervical Cells. *Journal of Histochemistry and Cytochemistry*, **27**(1), 621–628.

- [12] Beucher, S. and C. Lantuéjoul(1979) Use of watersheds in contour detection. *Proc. Int. Workshop Image Processing* Rennes, France.
- [13] Blanz, W.E. and E.R. Reinhardt(1981) Image segmentation by pixel classification. *Pattern Recognition*, **13**(4), 291–298.
- [14] Böcker, W., W.U. Müller, C. Streffer(1995) Comparison of different automatic threshold algorithms for image segmentation in microscope images. *Proc. SPIE - Applications of Digital Image Processing XVIII*, **2564**, 230–241.
- [15] Bookstein, F.L.(1979) Fitting Conic Sections to Scattered Data. *Computer Vision, Graphics, and Image Processing*, **9**, 56–71.
- [16] Borgefors, G.(1986) Distance Transformations in Digital Images. *Computer Vision, Graphics, and Image Processing*, **34**, 344–371.
- [17] Cahn, R.L., R.S. Poulsen and G. Toussant(1977) Segmentation of Cervical Cell Images. *Journal of Histochemistry and Cytochemistry*, **25**(7), 681–688.
- [18] Chow, C.K. and T. Kaneko(1972) Automatic Boundary Detection of the Left Ventricle from Cineangiograms. *Computer and Biomedical Research*, **5**, 388–410.
- [19] Davies, E.R.(1989) Finding ellipses using the generalised Hough transform. *Pattern Recognition Letters*, **9**(2), 87–96. Academic Press Limited, London.
- [20] Davies, E.R.(1990) *Machine Vision*, Academic Press Limited, London.
- [21] Fitzgibbon, A., M. Pilu, and R.B. Fisher(1999) Direct Least Squares Fitting of Ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**(5), 476–480.
- [22] Freed, J.A.(1997) Possibility of Correcting Image Cytometric Nuclear DNA (Ploidy) Measurements in Tissue Sections. *Analytical and Quantitative Cytology and Histology*, **19**(5), 376–386.
- [23] Fu, K.S. and J.K. Mui(1981) A survey on image segmentation. *Pattern Recognition*, **13**, 3–16.
- [24] Garner, D., G. Ferguson and B. Palcic(1994) The Cyto-Savant System. In *Automated Cervical Screening*, HK Grohs and OAN Husain eds., Ch. 23, 305–317, Igaku-Shoin Medical Publishers Inc.
- [25] Gonzalez, R.C.(1992) *Digital Image Processing*, Addison Wesley, New York.

- [26] Gschwendtner, A. and T. Miringner(1997) DNA Measurement on Histologic Slides: Does it Work on Human Tissue? *Analytical and Quantitative Cytology and Histology*, **19**(4), 294-300.
- [27] Gunn, S.R. and M.S. Nixon(1997) A robust Snake Implementation; A dual Active Contour. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**(1), 63-68.
- [28] Harms, H., H.M. Aus, M. Haucke and U.Gunzer(1986) Segmentation of Stained Blood Cell Images Measured at High Scanning Density With High Magnification and High Numerical Aperature Optics. *Cytometry*, **7**, 522-531.
- [29] Hough, P.V.C.(1962) Methods and Means for Recognizing Complex Patterns, US Patent 3,069,654.
- [30] Illingworth, J., and J. Kittler(1988) A survey of the Hough Transform. *Computer Vision, Graphics and Image Processing*, **44**, 87-116.
- [31] Inoué, S. and R. Oldenbourg(1995)Microscopes. in *Handbook of Optics Volume 2*, M.Bass ed., Ch. 17, Mcgraw Hill.
- [32] Liang, J. and J. Tucker(1997) DNA measurement of overlapping cell nuclei in thick tissue sections. *Analytical Cellular Pathology*, **14**, 41-49.
- [33] Johnson, R.P.(1990) Contrast Based Edge Detection. *Pattern Recognition*, **23**(3), 311-318.
- [34] Kass, M., A. Witkin and D. Terzopoulos(1988) Snakes: Active Contour Models. *Internation Journal of Computer Vision*, **1**, 321-331.
- [35] Kohler, R.(1981) A segmentation system based on thresholding. *Computer Graphics and Image Processing*, **15**
- [36] Koss, L.G.(1992) *Diagnostic Cytology and its Histopathological Bases*, J.B. Lippincott Company, Philadelphia.
- [37] Lam, S., and C. MacAulay(1998) Endoscopic localization of preneoplastic lung lesions. In *Clinical and biological basis of lung cancer prevention*, Y. Martinet *et al* eds, Birkhauser Verlag, Basel Switzerland.
- [38] Lam, S., T. Kennedy, M. Unger, Y.E. Miller, D. Gelmont, V. Rusch, B. Gipe, D. Howard, J.C. LeRiche, A. Coldman, A.F. Gazdar(1998) Localization of bronchial intraepithelial neoplastic lesions by fluorescence bronchoscopy. *Chest*, **113**, 696-702.

- [39] Lantuéjoul, C. and S. Beucher(1981) On the use of the geodesic metric in image analysis. *Journal of Microscopy*, **121**, 39.
- [40] Leavers, V.F.(1992) *Shape Detection in Computer Vision Using the Hough Transform*, Springer-Verlag, London.
- [41] Liedtke, C.E., T. Gahm, F. Kappei and B. Aeikens(1987) Segmentation of Microscopic Cell Scenes. *Analytical and Quantitative Cytology and Histology*, **9**(3), 196-211.
- [42] Lin, Y.K. and K.S. Fu(1981) Segmentation of Papanicolaou Smear Images. *Analytical and Quantitative Cytology*, **3**(3), 201-206.
- [43] Lockett, S.J., M. O'Rand, C.A. Rinehart, D.G. Kaufman, B. Herman, K. Jacobson(1991) Automated Fluoresence Image Cytometry - DNA Quantification and Detection of Chlamydial Infections. *Analytical and Quantitative Cytology and Histology*, **13**(1), 27-44.
- [44] MacAulay, C., and B. Palcic(1988) A Comparison of Some Quick and Simple Threshold Selection Methods for Stained Cells. *Analytical and Quantitative Cytology and Histology*, **10**(2), 134-138.
- [45] MacAulay, C., M. Guillaud, J.C. le Riche, C. Dawe, A. Coldman, I. Wistuba, I.W. Park, A. Gazdar, S. Lam(2001) Nuclear Morphometry of Bronchial Intraepithelial Neoplasia. Submitted for publication to *Biomarkers*.
Quick and Simple Threshold Selection Methods for Stained Cells. *Analytical and Quantitative Cytology and Histology*, **10**(2), 134-138.
- [46] Mackin, R.W., L.M. Newton, J.N. Turner, and B. Roysam(1998) Advances in High-Speed, Three-Dimensional Imaging and Automated Segmentation Algorithms for Thick and Overlapped Clusters in Cytologic Preparations. *Analytical and Quantitative Cytology and Histology*, **20**(2), 105-121.
- [47] Malpica, N., C.O. de Solórzano, J.J. Vaquero, A. Santos, I. Vallcorba, J.M. Garcia-Sagredo, F. del Pozo(1997) Applying Watershed Algorithms to the Segmentation of Clustered Nuclei. *Cytometry*, **28**, 289-297.
- [48] McInerney, T. and D. Terzopoulos(1996) Deformable models in medical image analysis: A survey. *Medical Image Analysis*, **1**, 91-108.
- [49] Meyer, F. and S. Beucher(1990) Morphological Segmentation. *Journal of Visual Communication and Image Representation*, **1**(1), 21-46.

- [50] Muammar, H.K. and M. Nixon(1991) Tristage Hough transform for multiple ellipse extraction. *IEEE Proceedings-E*, **138**(1), 27–35.
- [51] National Cancer Insitute of Canada(2001) *Canadian Cancer Statistics 2001*, NCIC, Toronto, Canada.
- [52] Olstad, B. and A.H. Torp(1996) Encoding of a priori information in Active Contour Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**(9), 863–872.
- [53] Pavlids, T.(1977) *Structural Pattern Recognition*, Springer-Verlag, New York.
- [54] Ranefall, P., L. Egevad, B. Nordin and E. Bengtsson(1997) A new method for segmentation of colour images applied to immunohistochemically stained cell nuclei. *Analytical Cellular Pathology*, **15**, 145–156.
- [55] Ranefall, P., K. Wester and E. Bengtsson(1998) Automatic quantification of immunohistochemically stained cell nuclei using unsupervised image analysis. *Analytical Cellular Pathology*, **16**, 29–43.
- [56] Russ, J.C.(1992) *Computer Assisted Microscopy - The Measurement and Analysis of Images*, Plenum Press, New York.
- [57] Russ, J.C.(1992) *The Image Processing Handbook*, CRC Press, Boca Raton, Florida.
- [58] Saccomanno, G., V.E. Archer, O. Auerbach, R.P. Saunders, L.M. Brennan(1974) Development of carcinoma of the lung as reflected in exfoliated cells. *Cancer*, **33**, 256–270.
- [59] Sherman A.B., L.G. Koss, M.C. Abbott and M.W. Liao(1981) A method of boundary determination in digital images of urothelial cells. *Pattern Recognition*, **13**(4), 285–291.
- [60] Soille, P.(1999) *Morphological Image Analysis: Principles and Applications*, Springer Verlag, Berlin.
- [61] St. John T.M. and F.D. Robinett(2000) Lung Cancer: Building Hope through Comprehensive Prevention. *Cancer*, **89**(11), 2377–2386.
- [62] Stephens, R.S.(1991) Probabilistic approach to the Hough transform. *Image and Vision Computing*, **9**(1), 66–71.
- [63] Sudbø, J., W. Kildal, B. Risberg, H.S. Koppang, H.E. Danielsen, A. Reith(2001) DNA Content as a Prognostic Marker in Patients with Oral Luekoplakia, *The New England Journal of Medicine*, **344**(17), 1270–1278.

- [64] Thomas A.D.H., T. Davies, and A.R. Luxmore(1992) The Hough Transform for Locating Cell Nuclei. *Analytical and Quantitative Cytology and Histology*, **14**(4), 347-353.
- [65] Tsuji, S., and M. Matsumoto(1978) Detection of ellipses by modified Hough transformation. *IEEE Transactions on Computers*, **27**(8), 777-781.
- [66] Van Veen, T.M. and F.C.A. Groen(1981) Discretization errors in the Hough transform, *Pattern Recognition*, **14**, 137-145.
- [67] Vincent, L. and P. Soille(1991) Watersheds in Digital Spaces: An Efficient Algorithm Based on Immersion Simulations. *IEEE Trans Pattern Analysis and Machine Intelligence*, **13**,(6), 383-397.
- [68] Watanbe, S.(1974) An automated cancer prescreening apparatus: CYBEST. *Computer Graphics and Image Processing*, **3**.
- [69] Weszka, J.S.(1978) A survey of threshold selection techniques. *Computer Graphics and Image Processing*, **7**, 259-269.
- [70] Weyn, B., G. van de Wouwer, M. Koprowski, A. van Daele, K. Dhaene, P. Scheunders, W. Jacob and E. van Mark(1999) Value of morphometry, texture analysis, densitometry and histometry in the differential diagnosis and prognosis of malignant mesothelioma. *Journal of Pathology*, **189**, 581-589.
- [71] Wright, N. and M. Alison(1984) *The biology of epithelial cell populations*, Clarendon Press, Oxford.
- [72] Yang, M.C., J.S. Lee, C.C. Lien and C.L. Huang(1997) Hough Transform Modified by Line Connectivity and Line Thickness. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**(8), 905-910.
- [73] Yuen, H.K., J. Illingworth and J. Kittler(1989) Detecting partially occluded ellipses using the Hough transform. *Image and Vision Computing*, **7**(1), 31-37.