Average Cell Orientation, Eccentricity and Size Estimated from Tissue Images

by

Peter J.W. Iles

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2005

Abstract

Five image processing algorithms are proposed to measure the average orientation, eccentricity and size of cells in images of biological tissue. These properties, which can be embodied by an elliptical ‘composite cell’, are crucial for biomechanical tissue models. To automatically determine these properties is challenging due to the diverse nature of the image data, with tremendous and unpredictable variability in illumination, cell pigmentation, cell shape and cell boundary visibility. One proposed algorithm estimates the composite cell properties directly from the input tissue image, while four others estimate the properties from frequency domain data. The accuracy and stability of the algorithms are quantitatively compared through application to a wide variety of real images. Based on these results, the best algorithm is selected.

Acknowledgements

Thank you to my supervisors, David Clausi and Wayne Brodland, for their commitment to the quality of my work. Thanks as well to Jim Veldhuis for his practical advice and enthusiasm, and Shannon Puddister for his foundational work on this project. Finally, thanks to Tracey for tolerating my academic leanings.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This thesis addresses a novel computer vision problem: to estimate the average orientation, eccentricity and size of cells in an image of biological tissue (Fig. 1.1a). Recent studies of embryonic tissues have shown that these geometric features are relevant to the mechanics of embryonic development [10] [7] [11].

These features can be embodied by an elliptical 'composite cell' as introduced by Brodland and Veldhuis [9]. This composite cell represents the most typical, or average, cell in the image. The composite cell is defined by its orientation α from the horizontal and major and minor axes L_{major} and L_{minor}, respectively (Fig. 1.1b). The aspect ratio κ and the area A of the ellipse correspond to the eccentricity and size of the composite cell and can be calculated from the axis lengths:

$$\kappa = \frac{L_{major}}{L_{minor}} \tag{1.1a}$$

$$A = \frac{\pi}{4} L_{major} L_{minor}. \tag{1.1b}$$

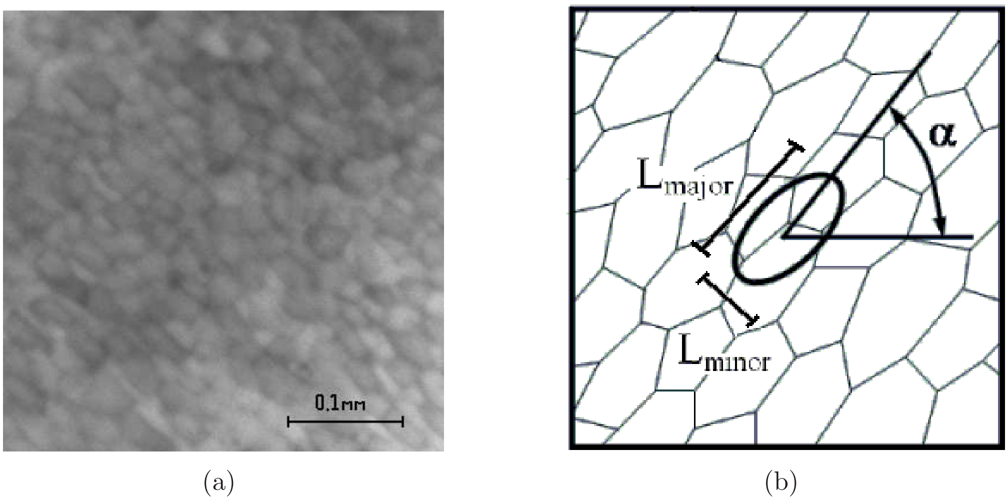

(a) (b)

Figure 1.1: (a) An image of an embryonic epithelium. (b) The composite cell is defined by its orientation α and its major and minor axes, L_{major} and L_{minor}.

The objective of the present work is to develop an algorithm that can estimate α, κ and A and is robust to unpredictable image characteristics, such as variations in illumination, cell pigmentation, cell boundary visibility and cell orientation and shape. This thesis proposes and evaluates five such algorithms: one which estimates the composite cell parameters directly from the raw image, and four others which estimate the parameters from frequency domain data. The frequency domain algorithms are: least squares ellipse fitting, area moments, correlation and axes searching, and Gabor filters. All methods are quantitatively compared through application to a wide variety of real images.

This thesis is organised as follows. Chapter 2 provides background information on this research, including the project context and related research. The five methods are described in Chapter 3 and their performance is evaluated in Chapter 4. Conclusions are made in Chapter 5.

Chapter 2

Background

2.1 Research context

This research is in support of biomechanical modelling of embryonic development [6], the purpose of which is to understand the mechanics of embryogenesis and the mechanical basis of birth defects such as spina bifida and cardiac septum defects. Slight irregularities in the stresses occurring in the embryonic epithelia, of which embryos are largely composed, are believed to cause these conditions. It has been shown that these stresses are related to the orientation, eccentricity and size of the epithelium cells as characterized by α, κ and A respectively [10] [7] [11]. Thus, to understand embryo development, it is crucial to measure these properties.

The tissue images considered in this thesis were taken from intact developing axolotl (amphibian) embryos using a custom microscope setup [35] and from excised pieces of embryonic tissue placed in a novel instrument that can apply a tensile stress to them [37].

The variability of the image characteristics make it difficult to design a robust algorithm. Substantial variability is found in illumination, cell pigmentation, cell boundary visibility and cell orientation and shape (Fig. 2.1). Finally, cell mitosis can cause dark spots to appear between newly forming cells, which can be mistaken as individual cells. An algorithm that is insensitive to these variations is desired.

The need to automate this process rather than relying on human measurement is two-fold. First, the large number of images that must be analyzed makes human measurement impractical. Second, avoiding human bias is preferred, and this can only be guaranteed with an automated method. Algorithmic bias must also be avoided, an issue addressed in Section 4.3.6.

2.2 Related research

To the best of the author's knowledge, the specific problem of estimating the average orientation, eccentricity and size of cells in images of biological tissue has not been addressed in the research literature. However, many image processing techniques have been developed for the general analysis of biological cells. These include methods for cell counting [1] [14] [31], individual cell analysis and tracking [26] [29] and multiple cell analysis and tracking [22] [38] [25]. Of these, the multiple cell analysis and tracking methods are relevant to this research. These methods, however, are either designed for images in which the cells do not touch (as in [25]), images with very high contrast (as in [38]) or images with total cell boundary visibility (as in [22], [38] and [25]). These methods are therefore inappropriate for the images considered in this thesis in which the cells are always touching and often have low-contrast, partly visible boundaries.

The problem considered here could also be approached from a texture analysis standpoint. The embryonic tissue can be thought of as a texture, having a certain average 'texture element' orientation, eccentricity and size.

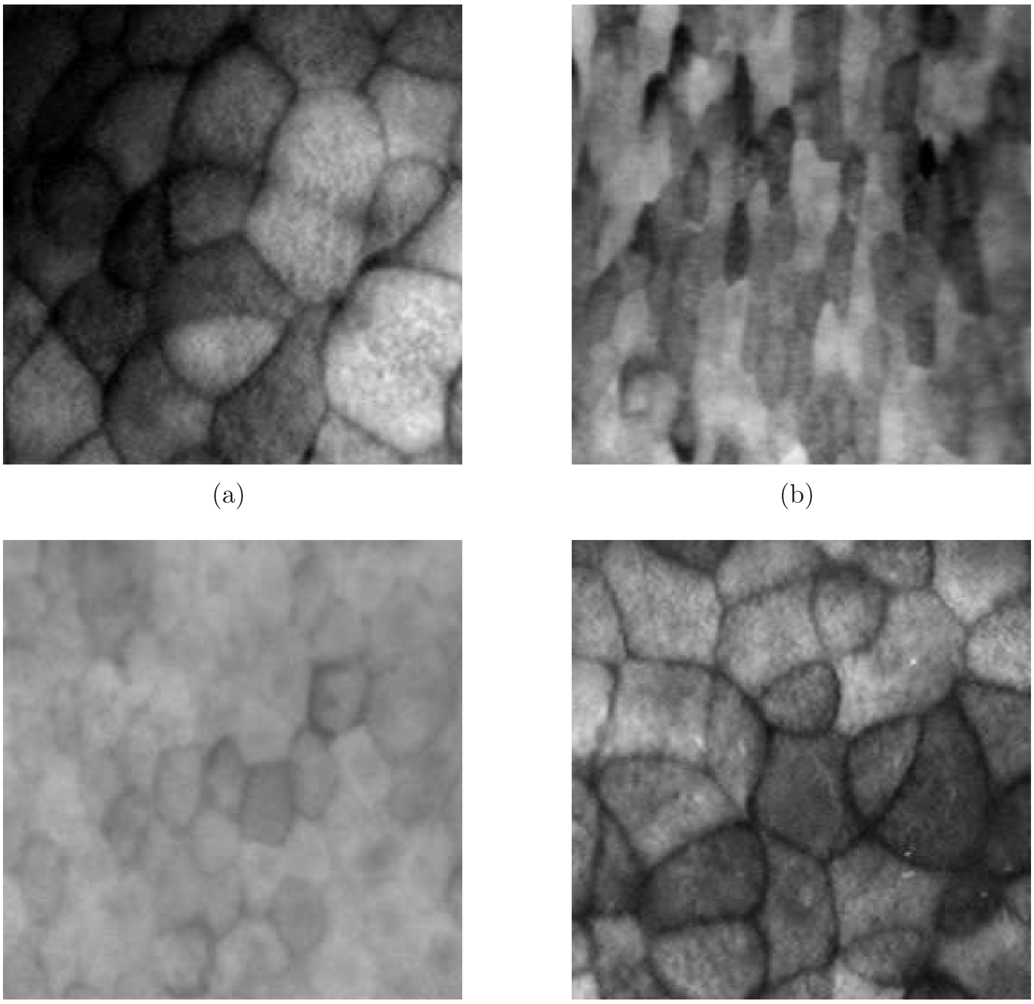
(a) (b)

(c) (d)

Figure 2.1: Variation in image characteristics: (a) illumination, (b) cell pigmentation, (c) cell boundary visibility and (d) cell shape and orientation.

Many methods have been developed to specifically estimate texture orientation. Vaidyanathan and Lynch [34] analyzed the directionality of detected edges, while Sourice et al. [32] proposed computing autocorrelation of pixel intensities to find texture orientation. These two methods were only recognised to be applicable for images containing roughly parallel lines, and hence are not appropriate for complex cellular images.

Others have focused on computing texture orientation by exploiting patterns in the frequency domain. These techniques are based on the concept that the energy in the frequency magnitude response image ($|F|$) is concentrated along an axis perpendicular to the texture orientation. Bigün et al. [2] propose a method to find an axis of symmetry in the $|F|$ image in order to estimate orientation. Kass and Witkin [20] use an orientation-selective bandpass filter and find the orientation yielding the greatest response. These two methods gave inspiration for the orientation estimation components of the frequency methods proposed in this thesis.

A ‘package solution’, however, was desirable: a single algorithm that could elegantly estimate all three composite cell parameters. To this end, methods were investigated that perform a structural analysis of texture. These methods attempt to detect regularly occurring texture elements arranged according to a ‘placement rule’. The placement rule consists of two dominant orientations with corresponding frequencies (Fig. 2.2) from which the three composite cell parameters can be derived. Matsuyama et al. [24] automatically find the placement rule through analysis of the frequency domain while others use co-occurrence matrix features to perform this task [21] [33]. There are three issues with these methods. First, they are meant for use with simple, regularly repeating texture images (as in Fig. 2.2a). The embryonic tissue images do not match this description since they can have large variations in cell size and placement. Second, the embryonic cells are tightly packed which makes it impossible to isolate a single cell in a parallelogram-shaped window

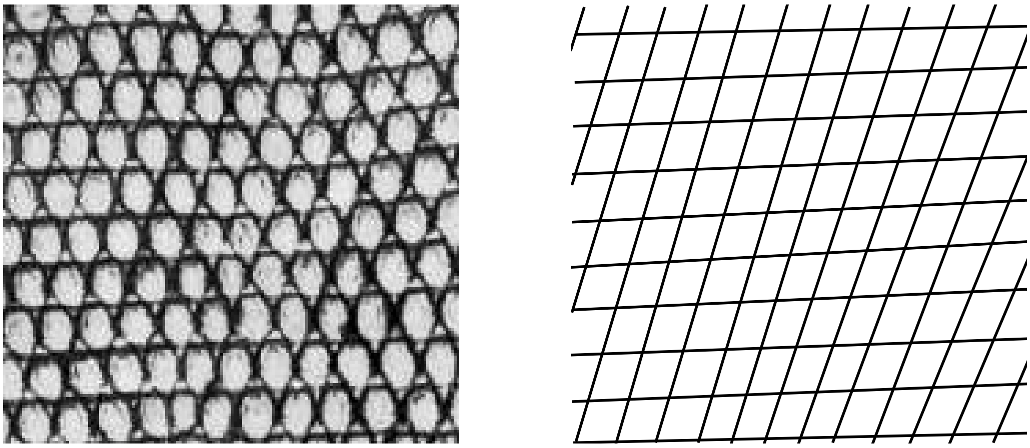

Figure 2.2: Placement rule defining texture element positions. (a) Regularly-repeating texture image [4]. (b) Placement rule grid lines defined by two orientations and two frequencies.

(as in Fig. 2.2b). Finally, in this thesis, the orientation of the cells themselves must be be detected, rather than the orientation of the *placement pattern* of the cells. The vertical orientation of the elements in Fig. 2.2a is not captured in the given placement rule. A placement rule could be used to find κ and A, but as described, these algorithms are not suited to the embryonic tissue images.

Novel methods were then developed because of the inability of the methods described above to estimated all three composite cell parameters. The frequency-based algorithms mentioned above motivated the development of the four frequency domain algorithms proposed in this thesis. The area moments method extends the Bigün et al. method to allow for the estimation of κ and A. The Gabor filters method extends the Kass and Witkin method by searching over a range of frequencies to again allow for the estimation of κ and A. The least squares ellipse fitting and correlation and axes searching methods were developed due to their attractive simplicity. For completeness, and given the unique images considered in this thesis, a spatial domain method was also developed.

Chapter 3

Methods

Five methods to estimate α, κ and A are introduced in this section: one spatial domain method, and four frequency domain methods. Initial versions of the four frequency domain methods were developed by Puddister [28] and are improved upon in this thesis. Please note that these algorithms process images with an intensity range of [0, 1].

3.1 Method #1: Spatial domain (SD)

The spatial domain method attempts to segment the image into cells, and then average the segment measurements to determine α, κ and A. There are three main stages to this method. First, local contrast is enhanced through the application of a new contrast enhancement routine. Next, an iterative watershed segmentation technique is used to detect individual cells. Finally, a segment analysis routine is applied to calculate the composite cell parameters. The spatial domain method is summarised in Tab. 3.1. The optimal algorithm parameters (for the test images introduced in Section 4.1) are given in Table 3.2.

Table 3.1: Spatial domain method summary.

Step	Description	Parameters
1	*Local contrast enhancement*	
1.1	Calculate local averages	r
1.2	Apply local sigmoidal transfer fcns	m
2	*Iterative watershed segmentation*	
2.1	Take image complement	
	loop	
2.2	Apply extended-minima transform	d
2.3	Apply imposed-minima transform	
2.4	Perform watershed segmentation	
2.5	Decrement d	d_{dec}
2.6	exit if min # segments found	n_{min}
	end loop	
2.7	Remove boundary segments	
3	*Segment analysis*	
3.1	Find centroid of each cell	
3.2	Find median 16 ray directions	
3.3	Fit ellipse to find α, L_{major} and L_{minor}	

Table 3.2: Optimal spatial domain method parameter values for test images (introduced in Section 4.1.)

r	m	d	d_{dec}	n_{min}
15	10	0.55	0.04	25

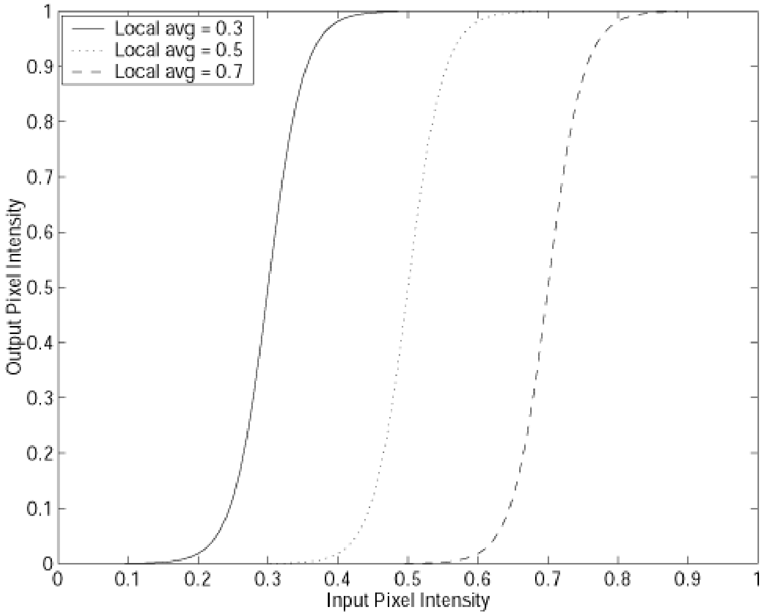

Figure 3.1: Sigmoid transfer functions for average local intensities of 0.3, 0.5 and 0.7 ($m = 10$).

3.1.1 Local contrast enhancement

The purpose of this step is to increase the local contrast of the input image, $f_{in}(x, y)$, to make the image appropriate for watershed segmentation. A local contrast enhancement algorithm was developed that is computationally inexpensive, is easily adjustable and does not saturate the pixel intensity range.

First, the local average is calculated. This is done by filtering the input image, $f_{in}(x, y)$, with a uniform circular local-averaging mask, $h_r(m, n)$, where r is the mask radius in pixels. The resulting image is:

$$f_{avg}(x, y) = h_r(m, n) * f_{in}(x, y). \tag{3.1}$$

Given the local average, a point operation enhances the local contrast using a sliding

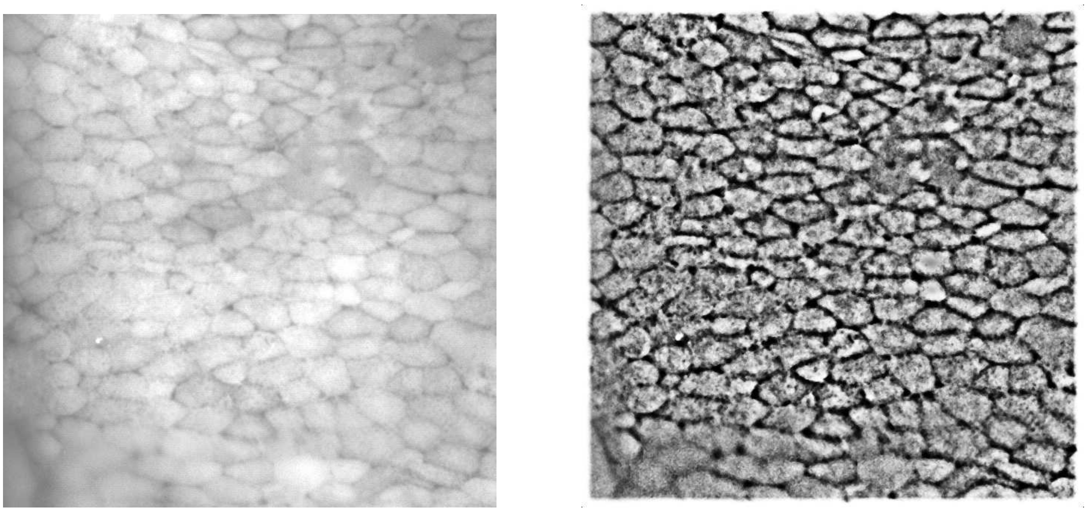

(a) (b)

Figure 3.2: An embryonic tissue image before and after adaptive contrast enhancement ($m = 10$; $r = 15$). (a) Original image (b) Local contrast enhanced image.

sigmoidal function. For each pixel, a sigmoidal transfer function is created which is centred on the local average (Fig. 3.1). When a maximum slope of $m > 1$ is used, this expands the pixel intensity range around the local average, increasing the local contrast. The contrast-enhanced image (Fig. 3.2b) is:

$$f_{out}(x, y) = \frac{1}{1 + e^{4m[f_{avg}(x,y) - f_{in}(x,y)]}}. \tag{3.2}$$

The sigmoidal function was chosen for three reasons. It is monotonically increasing (to provide a sensible pixel intensity mapping), the maximum slope can be set as a single parameter and the output intensity does not saturate the pixel intensity range.

3.1.2 Iterative watershed segmentation

This step attempts to segment the image into cells so that each cell can be analysed in isolation. The proposed iterative watershed method segments the contrast-enhanced image, $f_{out}(x, y)$, into the individual cells wherever possible. The watershed segmentation method is used for this application because it is fast and does not require *a priori* information of the number of segments.

The classic watershed segmentation method treats an image as a three-dimensional landscape (where pixel intensity is elevation) and segments it into its component watershed areas [36]. Every point in a given watershed area can be thought of as draining into a common 'drainage area', just as every point in the Lake Superior watershed drains into Lake Superior. The algorithm developer must choose the most appropriate set of drainage areas for the image. In this context, $f_{out}(x, y)$ contains round hills (the cells) surrounded by a network of narrow valleys (the cell boundaries).

The first step is to find the image complement, (f_{compl}), of $f_{out}(x, y)$ so that the cells become valleys and the boundaries become ridges. This is done so that the watershed segmentation method can recognise the cells as individual watershed areas.

There are three further steps that are iterated over until the number of segments detected exceeds a user-defined threshold, n_{min}. Each step uses an 8-connected (3×3) neighbourhood:

1. *Create binary map of drainage areas*: The extended-minima transform ('imextendedmin' function in MATLAB®) is applied to f_{compl} to produce a binary image where groups of zeros indicate drainage areas. Pixels in a drainage area are a minimum depth d below (darker than) their neighbours [23]. The drainage areas should be located in the interiors of each cell in the image. The minimum depth (d) parameter is used to make this step robust against minor intensity variations.

2. *Transform f_{compl} using drainage area map*: The imposed-minima transform ('imimposemin' function in MATLAB®) then transforms f_{compl} using the binary drainage area map to create a new grey-level image that has local minima located only at the drainage areas [23].

3. *Watershed segmentation*: The classic watershed segmentation method is applied to the transformed f_{compl} from the previous step. This determines the shape of the watershed segments (cells) that 'drain' into the detected drainage areas (local minima).

Each iteration decrements the minimum depth threshold, d, by d_{dec} starting from d_o, resulting in an increasing number of drainage areas. Therefore, the number of segments detected also increases over each iteration (Fig. 3.3). A large d_o is initially chosen which results in only a few cells being detected for most types of input images. When the number of cells detected, n, is greater than the user-defined n_{min}, the iterations stop. An image with well-defined cell boundaries requires few iterations since the intensity depressions are already deep. An image with poorly-defined cell boundaries requires more iterations.

These iterations remove the necessity of setting d as a hard threshold, making the segmentation more robust to different input image types. The threshold n_{min} is a more intuitive parameter and can be more easily set by the user than d. The parameter n_{min} represents the number of cells that must be detected to be confident in the results.

A single additional step is then performed on the segmentation result. Any cell that touches an edge of the image is removed. This is done because the overall shape of these cells is unknown, and hence should not be used in the calculation of the composite cell.


(a) (b) (c)

Figure 3.3: Example iterations of proposed segmentation method show increasing number of detected cells. Boundary cells are removed. The effects of large and small cells are minimised through median measure used in composite cell parameter calculation. (a) $d = 0.65$, (b) $d = 0.57$, (c) $d = 0.51$.

3.1.3 Segment analysis

Average cell measurements are now made using the segmented image. Finding the 'average shape' of the individual cells is not, however, a straightforward task. This task was done by quantifying each cell by taking ray measurements in 16 directions from its centroid to its edge (Fig. 3.4). The medians of these measurements were then found over all n cells. The 16 ray measurements define a sixteen-sided polygon that represents the composite cell. A best-fit ellipse (in the least squares sense) is calculated from this polygon [17] and α, L_{major} and L_{minor} are measured directly. Finally, κ and A can then be calculated using Eqs. (1.1a) and (1.1b).

The median measure was chosen over the mean due to its robust nature. The median measure prevents spuriously detected cells (both large and small) from having a major impact on the final result.

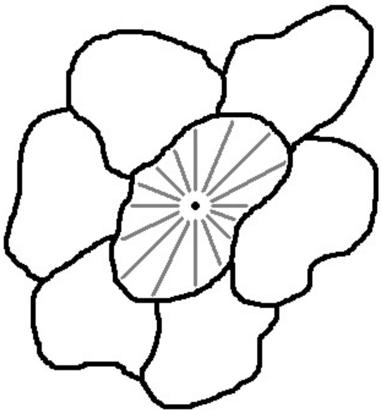

Figure 3.4: 16 ray measurements taken from centroid to edge of cell

3.2 Frequency domain methods

Four frequency domain methods are introduced in this section: least squares ellipse fitting, area moments, correlation and axes searching, and Gabor filters. These methods all take as input the frequency magnitude response image, $|F|$, generated using a two-dimensional fast Fourier transform (FFT) [16].

First, an explanation of how the $|F|$ image captures desired characteristics of the tissue image is given. This is followed by a description of three preprocessing steps that are shared by the four methods. The subsections following this describe the four proposed methods.

3.2.1 Estimating spatial patterns in the frequency domain

An image is converted into the frequency domain using a two-dimensional FFT. The resulting $|F|$ image can then be plotted on the horizontal u and vertical v axes, in pixel units of cycles per image (c.p.i.) with the origin of $|F|$ shifted to the center of the image to follow convention.

To illustrate how the composite cell parameters can be measured in the frequency domain, first consider the synthetic (Voronoi tessellation) cell image (Fig. 3.5a) [13] and its corresponding $|F|$ image (Fig. 3.5b). Note the obvious elliptical shape of the energy distribution.

The composite cell parameters can be derived from an ellipse fit to the energy distribution in $|F|$ (Fig. 3.5b). The strong high-frequency response shown in the vertical direction (along the v axis) in the $|F|$ image corresponds to the minor axis of the composite cell, and can be seen visually in the spatial domain as the high frequency of cell edges encountered in the vertical direction. The relatively weaker low-frequency response in the horizontal direction (along the u axis) corresponds to the major axis of the composite cell, and can be seen in the spatial domain as the lower frequency of cell edges encountered in the horizontal direction. Note that the strongest response in the $|F|$ image corresponds to the minor axis of the composite cell because the number of repetitions of the minor axis in the image is greater than that of the major axis. The orientation, α, is simply perpendicular to the major orientation of the ellipse in the $|F|$ image. The axis lengths of the ellipse in the $|F|$ image can be converted from c.p.i. to spatial pixels, allowing κ and A of the composite cell to be calculated using Eqs. (1.1a) and (1.1b) respectively.

The benefit of using the frequency domain is that the FFT has the effect of averaging spatial variation of cell properties in the tissue. In the context of this research, the strongest responses in the $|F|$ image reveal the *average* cell orientation and axis lengths. This effect makes the frequency domain methods robust to spatial variations in cell shape, orientation and cell edge visibility. The other image characteristic issues, illumination and pigmentation variation are addressed in the preprocessing steps outlined in the next section.

The difficulty in using a frequency approach is achieving an accurate fit of an ellipse to

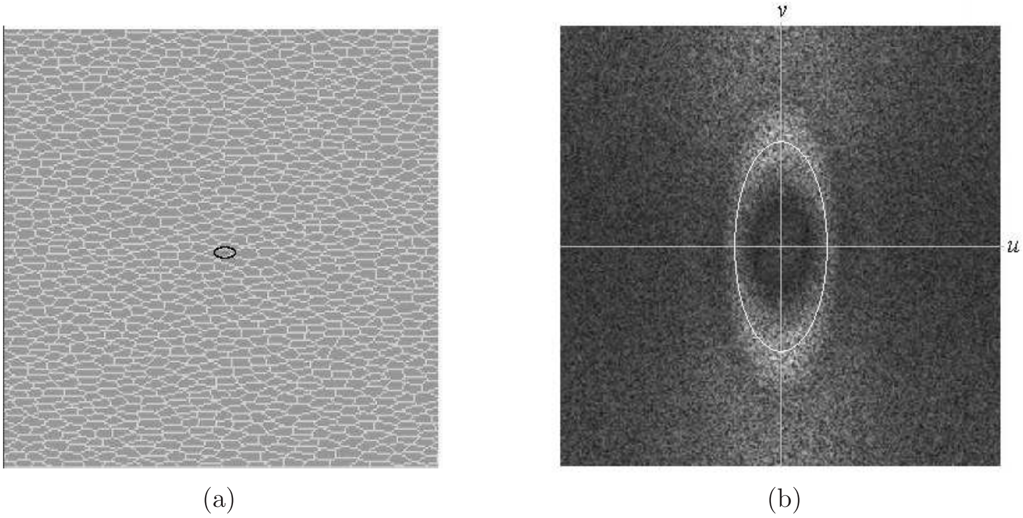

Figure 3.5: A cellular image and its corresponding FFT magnitude response image, $|F|$. (a) A synthetic 350×350 cellular image with $\alpha = 0°$, $\kappa = 1.8$, $A = 24$ pixels. The composite cell is drawn in black. (b) FFT magnitude response image, $|F|$, in pixel units of cycles per image. Transform of composite cell is shown as white ellipse. A square root point operator has been applied to enhance visibility.

the $|F|$ energy distribution. As shown in the next section, images of real axolotl embryonic epithelia do not typically produce $|F|$ images as clean as that shown in Fig. 3.5b. The four methods attempt to achieve this ellipse fitting in different ways. A summary of how the composite cell parameters are found for each of the frequency domain methods is shown in Fig. 3.6.

Stationarity of the geometric statistics of the cells (α, κ and A) is assumed. If these statistics varied from one part of the image to another, this would create multiple patterns in the $|F|$ image. This would not be appropriate input data for the proposed methods.

3.2.2 Preprocessing

Three preprocessing steps are applied before each of the four methods to prepare the raw image for processing:

1. *Find image gradient magnitude*: This is performed on the spatial domain image to detect cell boundaries and remove the effect of cell pigmentation. The combined magnitude image produced by the horizontal and vertical Sobel operators was used [12] but any other fundamental gradient operator could also be used. The FFT is applied after this step.

2. *Set low-frequency components to zero*: Frequency components at or below $F_{low} = 4$ cycles per image are set to zero since these components represent spatial patterns too large to be individual cells. The image data will not be captured at such cell resolutions. This also removes low-frequency illumination variation.

3. *Threshold* $|F|$: The $|F|$ image is first normalised to [0, 1] and then thresholded by setting to zero components below $F_{thresh} = 0.30$ in order to remove low-energy noise. This threshold was a robust choice for all the methods.

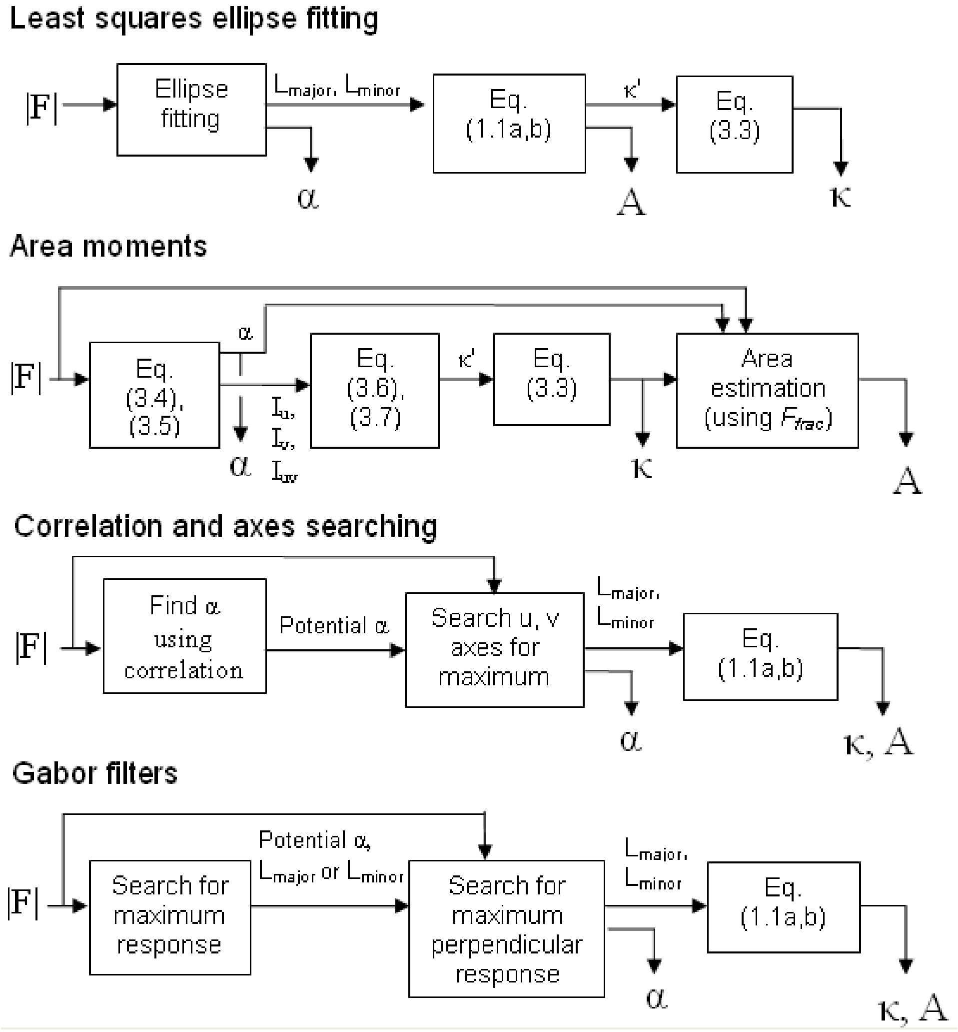

Figure 3.6: Summary of frequency domain methods for finding composite cell parameters α, κ and A from the preprocessed $|F|$ image.

3.2.3 Method #2: Least squares ellipse fitting (LSEF)

A direct way to fit an ellipse to the $|F|$ image data is to find the best-fit ellipse in a least squares sense. Numerous methods have been developed for this purpose, many of which fit data points to a general conic cross-section and use a constraint to force the solution into an ellipse. Fitzgibbon et al. [17] present a direct least squares based ellipse specific method that is robust and computationally inexpensive. As with most other least squares ellipse fitting algorithms, this algorithm uses a form of 'algebraic distance' [30] which approximates the perpendicular distance from a data point to the ellipse and is much simpler to compute. This algorithm is applied to the $|F|$ image.

The $|F|$ image data is converted into a set of data points in the following manner: a data point at pixel column x and row y is created for each 0.1 pixel intensity value. Thus a pixel with an intensity of 0.4 would be converted into four data points. Such a bin value was chosen to create sufficient points to characterise the $|F|$ image. Once the ellipse is fit to this data, α, L_{major} and L_{minor} can be calculated directly, and κ and A can then be calculated using Eqs. (1.1a) and (1.1b), respectively.

The values of κ were found to be biased high for this algorithm. This is because the energy corresponding to L_{minor} in the $|F|$ image was greater than that corresponding to L_{major} (because there are more L_{minor} repetitions in the tissue image). The fitted ellipse would therefore consistently be 'stretched' in the L_{minor} (the long axis in the $|F|$ image) direction, resulting in a greater κ value. A linear corrector function was empirically determined (using test images, as described in Chapter 4) that minimised the least-squared error of κ:

$$\kappa = \kappa' - m_{\kappa}(\kappa' - 1) \tag{3.3}$$

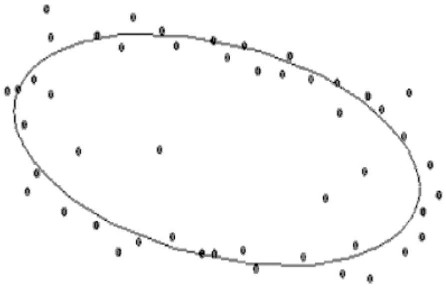

Figure 3.7: Example least squares ellipse fitting [28].

where $m_\kappa = 0.29$ and κ' is the biased estimate of the true aspect ratio. Note that this function does not adjust $\kappa' = 1$, since the aforementioned bias would not have an effect in this case because $L_{minor} = L_{major}$.

3.2.4 Method #3: Area moments (AM)

The second frequency domain method uses the concept of area moments to estimate α and κ. An ellipse, constrained by these two parameters, is then found that contains a certain percentage of energy in the $|F|$ image. The size of this ellipse yields the estimate for A.

Area moments provide a measure of the resistance of an object to angular acceleration about an axis. The ellipse in $|F|$ can be thought of as a two-dimensional object of varying density (pixel intensities) centered at the origin. The idea of using area moments in the $|F|$ image to find texture *orientation* was suggested by Bigün et al. [2]; their method actually solves a matrix eigenvalue problem in the spatial domain.

The proposed method implements this concept completely in the magnitude frequency domain, where all three composite cell parameters can be determined. First, the moments about the u and v axes (I_u, I_v) and product of inertia (I_{uv}) of the $|F|$ image can be found as follows [18]:

$$I_u = \sum_{v=1}^{rows} \sum_{u=1}^{cols} |F(u,v)| \cdot v^2 \tag{3.4a}$$

$$I_v = \sum_{v=1}^{rows} \sum_{u=1}^{cols} |F(u,v)| \cdot u^2 \tag{3.4b}$$

$$I_{uv} = \sum_{v=1}^{rows} \sum_{u=1}^{cols} |F(u,v)| \cdot u \cdot v \tag{3.4c}$$

After computing these inertial measures, α can be computed as:

$$\alpha = \arctan\left(\frac{2I_{uv}}{I_u - I_v}\right) \tag{3.5a}$$

provided $I_u \neq I_v$. In the case that $I_u = I_v$,

$$\alpha = \begin{cases} \pi/4 & \text{if } I_{uv} > 0 \\ 0 & \text{if } I_{uv} = 0 \\ -\pi/4 & \text{if } I_{uv} < 0. \end{cases} \tag{3.5b}$$

Using Mohr's circle [18], $|F|$ can be rotated by α to set I_{uv} to zero, resulting in only two moments:

$$I_1 = \frac{I_u + I_v}{2} + \sqrt{\left(\frac{I_u - I_v}{2}\right)^2 + I_{uv}^2} \tag{3.6a}$$

$$I_2 = I_u + I_v - I_1. \tag{3.6b}$$

κ can then be calculated directly as a ratio of the roots of these moments [5]:

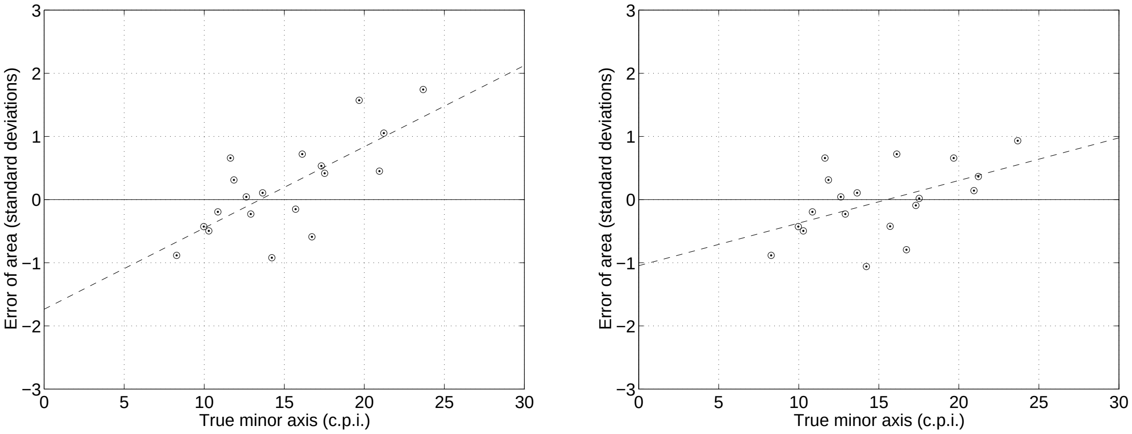

(a) (b)

Figure 3.8: Relation of AM area error, ζ_A (See Chapter 4 for description of error metrics), to true minor axis frequency. (a) Original results show positive correlation. (b) Results using F'_{low} show smaller correlation. Dashed lines show best-fit lines in the least squares sense.

$$\kappa = \frac{max(\sqrt{I_2}, \sqrt{I_1})}{min(\sqrt{I_2}, \sqrt{I_1})}. \tag{3.7}$$

For the same reason as the least squares ellipse fitting method, the κ values for this method were biased high. The corrector function in Eq. (3.3) is applied again here, using $m_\kappa = 0.38$ calculated to minimise the least-squared error.

The composite cell area, A, was estimated by finding an ellipse constrained to the estimated α and κ which contained a fraction F_{frac} of the energy in the $|F|$ image. The fraction that minimised the error for the test images was empirically found to be $F_{frac} = 0.28$ (See Chapter 4 for a description of the test images and error metrics).

The resulting area error, however, was found to be correlated with the number of true minor axis cycles per image (Fig. 3.8a). The reason for this correlation was that as the minor axis frequency became larger, a greater amount of lower frequency noise also

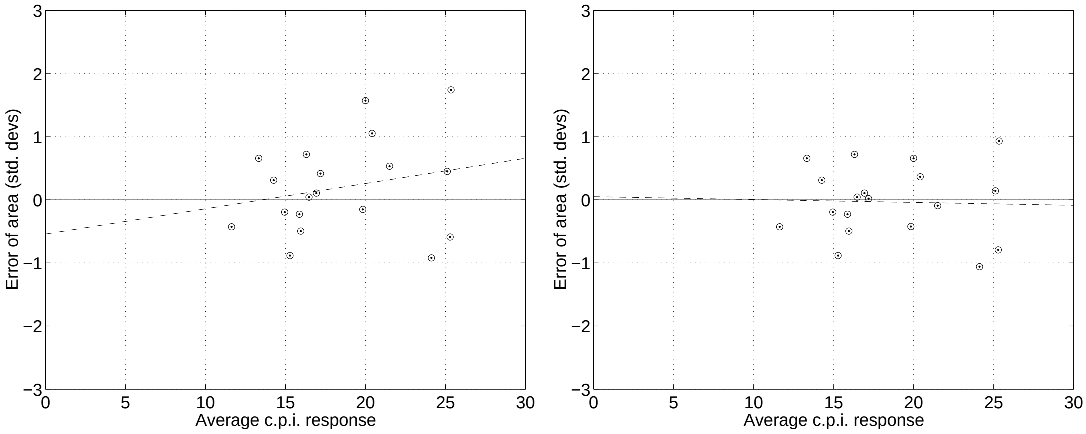

(a)

(b)

Figure 3.9: Relationship between average c.p.i. response and ζ_A. (a) Original results show positive correlation. (b) Correlation is removed by using F'_{low}. Dashed lines show best-fit lines in the least squares sense.

appeared. This caused the detected ellipse to be too small in the $|F|$ image, and hence cause the estimated composite cell to be too large. The opposite effect occurred when the minor axis frequency was too small.

The true minor axis frequency value is unknown from the point of view of the algorithm, and hence could not be used to remove the correlation. However, the *average c.p.i.* in the $|F|$ image could be found which demonstrated a similar correlation with the area error (Fig. 3.9a). This value was used to increase the 'low frequency' threshold defined in Step 2 of the preprocessing. This means that as the average c.p.i. increases, the greater the range of noisy low frequency components that are set to zero. The new low frequency threshold was set to:

$$F'_{low} = \frac{(average\ c.p.i.)}{3} \tag{3.8}$$

which removed the correlation between the average c.p.i. and the area error (Fig. 3.9b).

This, in turn lessened the correlation between the true minor axis frequency and the area error (Fig. 3.8b).

In summary, to estimate A, the AM method first increased the low frequency threshold F_{low} to F'_{low} using the average c.p.i. in the $|F|$ image. Second, an ellipse constrained to the estimated α and κ was found that contained a fraction $F_{frac} = 0.28$ of the energy in the $|F|$ image.

3.2.5 Method #4: Correlation and axes searching (CAS)

The third frequency domain method estimates α by finding the angle of a line of symmetry in the $|F|$ image. The axis lengths are found by searching for maxima along two rays in the $|F|$ image: one parallel and one perpendicular to the line of symmetry.

The correlation coefficient, $-1 < \rho < 1$, is used to find the axis of symmetry, and is defined for two data vectors, a and b as:

$$\rho = \frac{E[(a - \mu_a)(b - \mu_b)]}{\sqrt{var(a)var(b)}} \tag{3.9}$$

where $E[.]$ is the expectation operator and $var(.)$ is the sample variance operator. The sample means for a and b are μ_a and μ_b respectively. If $\rho = 1$, then the elements in a and b have a positive linear relationship. If $\rho = -1$, there is a negative linear relationship. If $\rho = 0$, there is no relationship.

To determine the orientation of the ellipse in the $|F|$ image, a new image is created by rotating $|F|$ by single degrees through $1° < \alpha < 90°$. In creating the rotated $|F|$ images, only a circular area (with a diameter equal to the width and height of the image) in the $|F|$ image is considered. After each rotation, the image is "folded" along the vertical v axis, and ρ is measured between corresponding pixel intensities using Eq. 3.9. This is

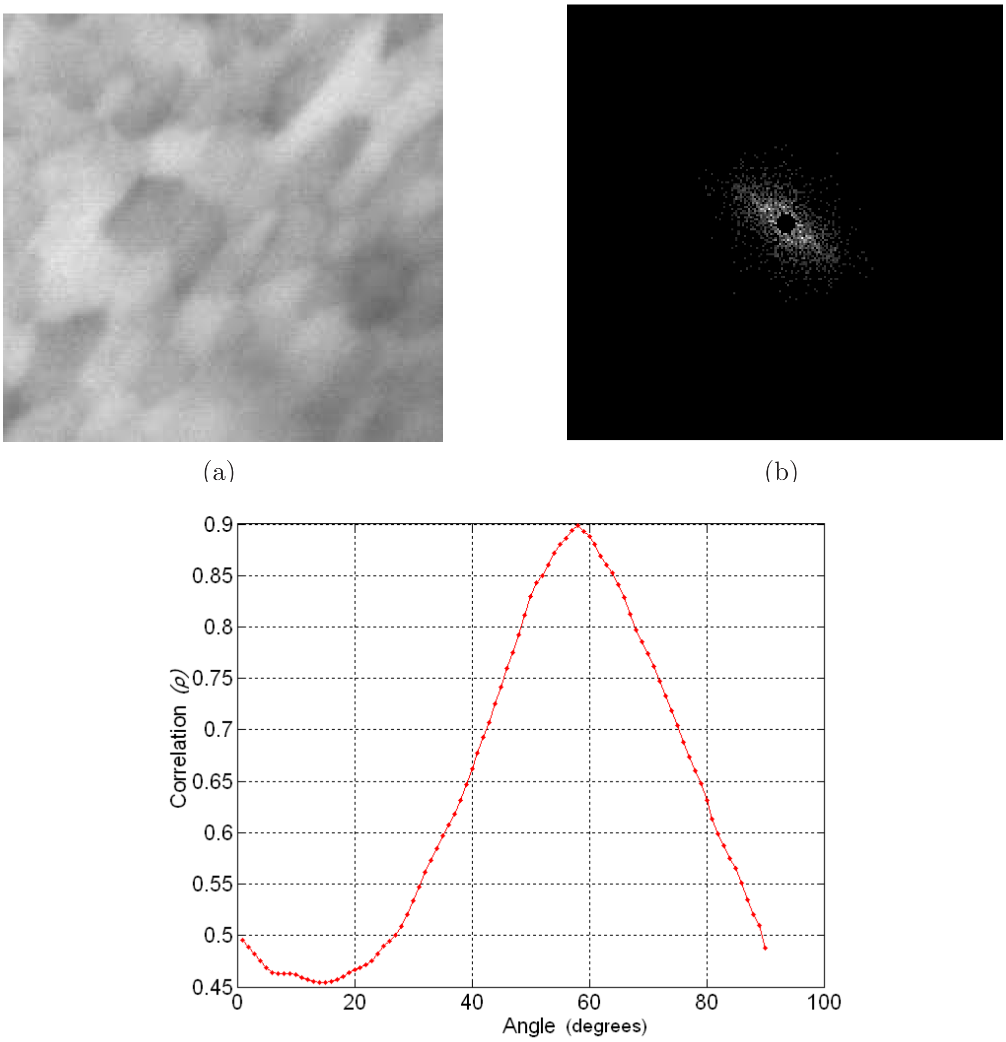

(a)

(b)

(c)

Figure 3.10: Estimating orientation using correlation method. (a) Input image with cells rotated at 60° to the horizontal; (b) corresponding preprocessed $|F|$ image; (c) ρ vs. angular rotation shows a well-defined maximum at 60°. This angle corresponds to *either* the major or minor axis of the ellipse in the $|F|$ image.

done by arranging the pixels in the left and right halves of the image into vectors a and b respectively. When the $|F|$ ellipse is aligned (either horizontally or vertically) with the u and v axes, ρ is at a maximum (Fig. 3.10). The angle of rotation used to produce the maximum ρ therefore corresponds to *either* the major or minor axis of the composite cell.

The maximum responses along the u and v axes are now found. A narrow 3-pixel-wide band along each axis is considered to account for any small errors in the orientation estimation. Of the two detected maxima, the response found at the higher frequency corresponds to L_{minor} while the response found at the lower frequency corresponds to L_{major}. Gaussian smoothing with $\sigma = 1$ is applied along the narrow bands to minimise the effect of noise. Now α is known and κ and A can be calculated from the measured axis lengths using Eqs. (1.1a) and (1.1b), respectively.

3.2.6 Method #5: Gabor filters (GF)

The final frequency domain method searches for the orientation and frequency yielding the maximum Gabor filter response. The maximum response is then found along an axis perpendicular to the axis of the global maximum. The parameters α, κ and A are derived from these two maxima. This is an extension of the Kass and Witkin method [20] described earlier which used an orientation-selective bandpass filter to detect *orientation* only.

Gabor filters are filters that can decompose an image into frequency- and orientation-specific texture features [3] [15]. This means that a Gabor filter can be designed to be sensitive to specific frequencies and orientations. The Gabor filter response at a location (u, v) in the frequency domain is defined as:

$$H(u, v) = \exp\left[-2\pi^2\{(u - F)^2\sigma_x^2 + v^2\sigma_y^2\}\right] \tag{3.10}$$

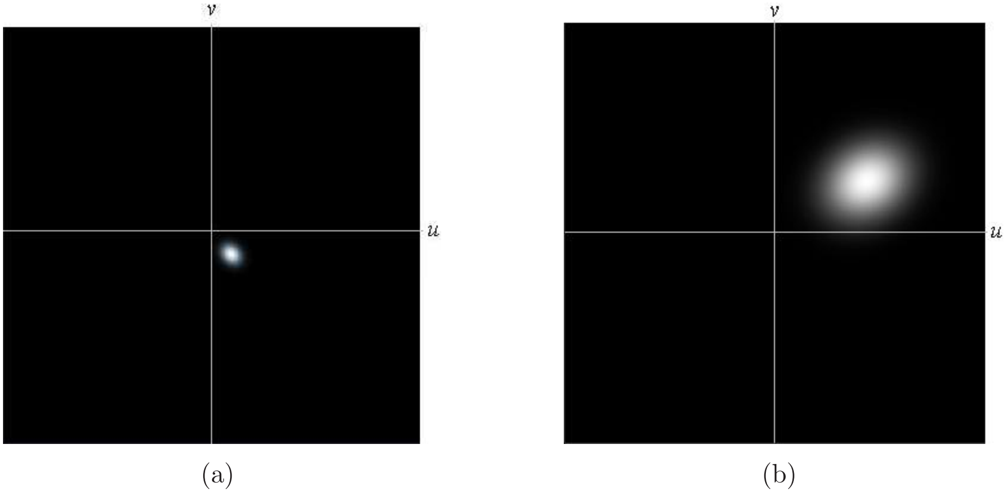

(a) (b)

Figure 3.11: Frequency responses for Gabor filters (B_F = one octave, $B_\theta = 30°$): (a) 16 c.p.i. oriented at $-45°$ and (b) 64 c.p.i. oriented at $30°$.

where F is the frequency, and σ_x and σ_y are the standard deviations in frequency and orientation, respectively [19]:

$$\sigma_x = \frac{\sqrt{\ln 2}(2^{B_F}+1)}{\sqrt{2}\pi F(2^{B_F}-1)} \tag{3.11}$$

$$\sigma_y = \frac{\sqrt{\ln 2}}{\sqrt{2}\pi F \tan(\frac{B_\theta}{2})}. \tag{3.12}$$

Clausi and Jernigan [15] argued that the frequency bandwidth, B_F and the angular bandwidth, B_θ should be set to one octave and 30° respectively in order to best model the human visual system. Two examples of Gabor filters with these bandwidths are presented in Fig. 3.11.

This method first searches for the Gabor filter parameters (frequency and orientation) that produce the strongest filter response in the $|F|$ image. Frequencies are tested at every second cycle per image, and orientations at every third degree. The orientation of the

Gabor filter with the strongest response is associated with the orientation of the composite cell. The frequency of the strongest response corresponds to *either* L_{major} or L_{minor}.

A ray perpendicular to the detected orientation is then searched to find a local maximum. This corresponds to the other axis length of the composite cell. The axis with the higher frequency corresponds to L_{minor}, while the axis with the lower frequency corresponds to L_{major}. The orientation is then associated with the major axis. As before, κ and A can then be found using Eqs. (1.1a) and (1.1b), respectively.

In initial testing, Gabor filters were found to be satisfactory only for very regular celluar images with very strong patterns in the $|F|$ energy distribution [28]. They did not work well for typical cell images and hence are not included in the following section. The maximum Gabor filter responses were consistently being ‘pulled’ to frequencies that were too high, apparently because of the larger filter bandwidth associated with the higher frequency Gabor filters when implemented using a pseudo-wavelet format [15].

Chapter 4

Evaluation of Algorithms

The four methods – spatial domain (SD), least squares ellipse fitting (LSEF), area moments (AM) and correlation and axes searching (CAS) – were evaluated using images of real embryonic epithelia. The Gabor filter method results are not reported due to its inferior performance. This chapter first describes the test images and a technique for measuring the 'true' composite cell parameters. The experimental results are then presented and discussed. Finally, estimation confidence measures are described for the chosen method.

4.1 Test images

The methods were evaluated using 19 individual images of embryonic epithelia and a set of 11 time-lapse images in which a patch of embryo tissue was stretched. These test images and their corresponding gradient magnitude and preprocessed $|F|$ images (used only by the frequency domain methods) are shown in Appendix A (Figs. A.1-3, A.5-7). The images demonstrate a wide variety of image characteristics. An example test image, its gradient magnitude image and preprocessed $|F|$ image are shown in Figs. 4.1a-c.

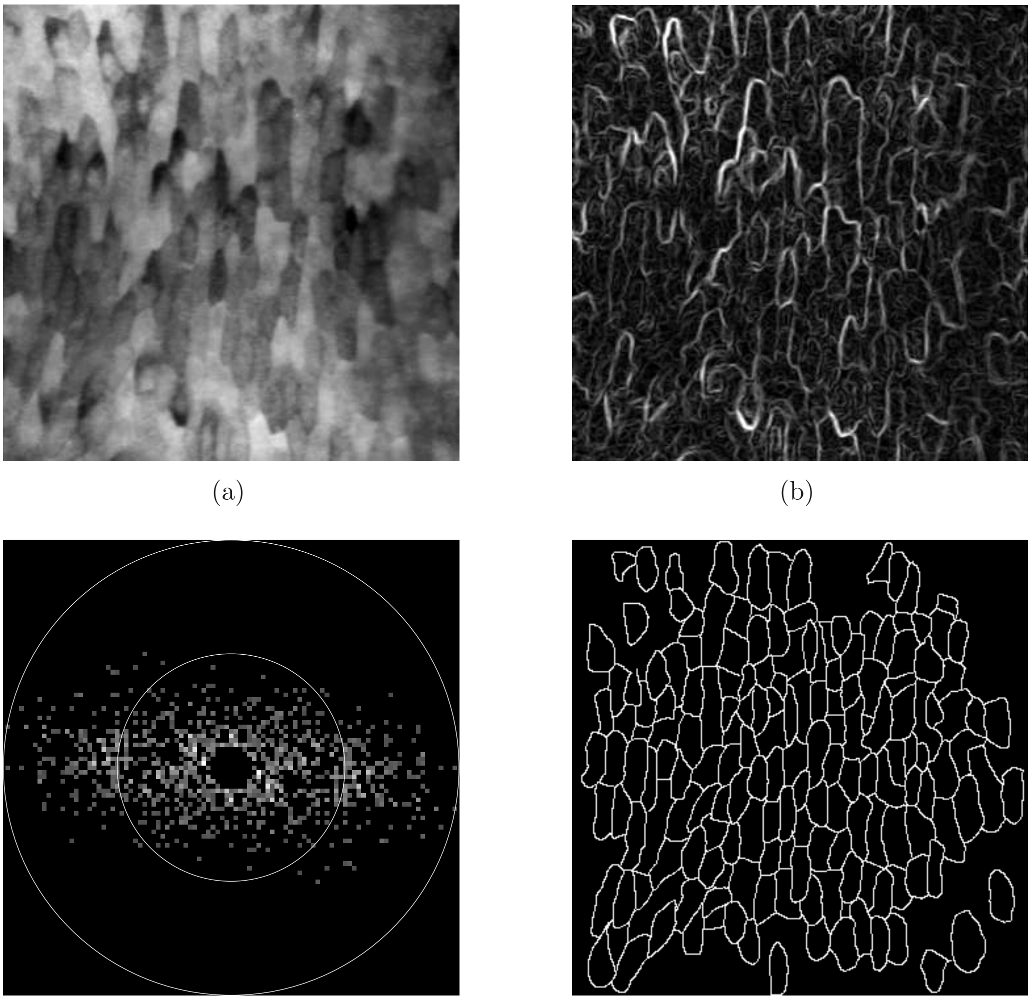

(a) (b)

(c) (d)

Figure 4.1: Example (a) test image and corresponding (b) gradient magnitude image, (c) preprocessed $|F|$ image (concentric circles show 25 and 50 c.p.i.) and (d) hand segmentation (truth) image.

4.2 True measurements

In order to measure the performance of the algorithms, the true values of α, κ and A ($\hat{\alpha}$, $\hat{\kappa}$ and $\hat{A}$ respectively) for each test image must be known. These values were found by hand-segmenting the test images and applying a variation of the AM method. This method, and the associated errors in finding the true measurements are described in the next two subsections. The hand-segmented images are shown in Appendix A (Figs. A.2, A.6). An example hand-segmented image is shown in Fig. 4.1d.

Manually fitting an ellipse in the $|F|$ image was also considered as a way to find the true parameter values. However, the lack of definite peaks in the $|F|$ images made this approach highly subjective.

4.2.1 Methodology

First, the tissue images were hand-segmented. This process was difficult because many cells in the test images did not have completely visible boundaries. To maximise the information captured by the segmentation, cells with roughly 75% visible boundaries were segmented (the position of the invisible edge sections were estimated).

Given the segmented image, $\hat{\alpha}$ and $\hat{\kappa}$ were found by calculating the moments of inertia, I_u and I_v, and product of inertia, I_{uv}, for each cell segment. These values were then averaged over all cells and Eqs. (3.4)-(3.7) were applied. Note that $\hat{\kappa}$ is therefore the average cell aspect ratio *in the direction of* $\hat{\alpha}$. For $\hat{A}$, the segmented cell areas were simply averaged. This method was validated by applying it to synthetic Voronoi tessellations with known $\hat{\alpha}$, $\hat{\kappa}$ and $\hat{A}$ values, and comparing these values to the measured ones.

Standard deviations of $\hat{\alpha}$, $\hat{\kappa}$ and $\hat{A}$ over all cells ($\sigma_{\hat{\alpha}}$, $\sigma_{\hat{\kappa}}$ and $\sigma_{\hat{A}}$ respectively) were also measured. If $\sigma_{\hat{\alpha}}$, $\sigma_{\hat{\kappa}}$ or $\sigma_{\hat{A}}$ were large, the performance of the algorithms was expected to

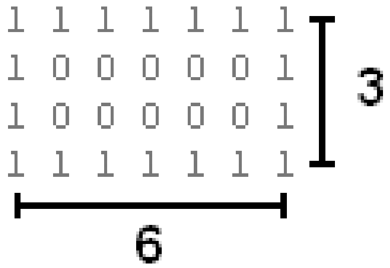

Figure 4.2: This cell, with interior denoted by 0s and edges denoted by 1s, measures 6 pixels by 3 pixels. Each edge is only considered half a pixel thick.

be degraded.

4.2.2 Errors in true measurements

There are three types of errors in the true measurements. These are due to hand segmentation of the input image, the appearance model chosen for the cells and the digitisation of the cells.

The first error was a result of the hand segmentation. The segmentation process is subjective and its error cannot be quantified, but it is a necessary step to find the best estimate of $\hat{\alpha}$, $\hat{\kappa}$ and $\hat{A}$.

The second error was due to the cell appearance model chosen. A cell was defined as being composed of its interior region and *half* the thickness of its edges (Fig. 4.2) to prevent any overlap among cells, since all cell boundaries are shared. This definition also agrees with how the FFT captures the image information. This definition is difficult to follow when hand segmenting cells because the cell edges would have to be made of ‘half pixels’.

Two methods were designed to solve this problem. First, the pixels at the cell edges were given half the weight (intensity) of the interior cell pixels in the area moments calculation.

Second, the cell segment was doubled in size (using bilinear interpolation) and was dilated by one pixel. This doubling has no affect on the resulting $\hat{\alpha}$ and $\hat{\kappa}$ measurements, and the $\hat{A}$ measurement would simply be divided by $2^2 = 4$. The errors for each of the two methods were quantified by applying them to synthetic Voronoi tessellations with known parameters.

The second method was found to produce a more exact solution. The error in the detected $\hat{\kappa}$, for example, lessened as the cell segment grew in size, and was about 1% for a 15×30 pixel cell. Rarely were the segmented cells smaller than this size and hence the resulting errors were considered insignificant.

The third error was due to the digitisation (i.e. into pixels) of the cells. This digitisation affects the calculated I_u and I_v values. However, this effect also lessened with increased cell sizes. This error was measured by hand-calculating the moments for a continuous cell and comparing it to the moments of the digitised cell. For a 14×7 pixel cell the error in both moments was about 1%. This error is less than the error described above and is hence also considered insignificant.

4.3 Experimental results

This section first introduces a performance objective, followed by the presentation and discussion of the experimental results for the 19 test images. A Student's T test was then applied to check for bias in the results, and a time-lapse test with a different set of 11 images was done to measure stability of the results. Based on these results, the most accurate and stable method is chosen. Finally, confidence measures for the chosen method are described.

The results for the four methods when applied to the 19 test images are presented in

Table 4.1: Distance from $\hat{\alpha}$ for each method in degrees ($\epsilon_\alpha = \alpha - \hat{\alpha}$) and number of standard deviations ($\zeta_\alpha = \epsilon_\alpha / \sigma_{\hat{\alpha}}$) for test images. Best results shown in bold.

Image: $\hat{\alpha}$, $\sigma_{\hat{\alpha}}$	SD	LSEF	AM	CAS
	ϵ_α (ζ_α)	ϵ_α (ζ_α)	ϵ_α (ζ_α)	ϵ_α (ζ_α)
a: 36.8, 49.2	−18.1 (−0.37)	+29.2 (+0.59)	+21.4 (+0.43)	+15.2 (+**0.31**)
b:173.1, 16.1	+8.6 (+0.53)	−3.7 (−0.23)	+1.7 (+**0.11**)	+2.9 (+0.18)
c: 85.9, 17.6	+5.1 (+0.29)	−1.1 (−0.06)	+0.1 (+**0.01**)	+0.1 (+**0.01**)
d:141.0, 33.1	+9.8 (+0.30)	−3.6 (−**0.11**)	−3.7 (−0.11)	−5.0 (−0.15)
e:119.4, 27.6	−8.5 (−0.31)	+2.6 (+0.09)	+1.6 (+**0.06**)	−2.4 (−0.09)
f: 87.8, 21.6	+0.9 (+0.04)	−7.0 (−0.32)	−7.5 (−0.35)	−0.8 (−**0.04**)
g: 52.4, 17.3	+5.1 (+0.29)	−1.9 (−**0.11**)	−2.9 (−0.17)	−3.4 (−0.20)
h: 11.3, 34.3	−9.8 (−0.29)	+0.8 (+**0.02**)	+1.1 (+0.03)	+2.7 (+0.08)
i: 75.5, 40.2	−14.2 (−0.35)	−2.8 (−0.07)	−1.3 (−**0.03**)	+5.5 (+0.14)
j:128.4, 45.4	−13.4 (−**0.30**)	−32.2 (−0.71)	−21.5 (−0.47)	−20.4 (−0.45)
k: 8.9, 33.5	−4.0 (−0.12)	+2.3 (+**0.07**)	+2.3 (+**0.07**)	+8.1 (+0.24)
l: 22.7, 25.7	−3.0 (−0.12)	−1.4 (−**0.05**)	−1.7 (−0.07)	−2.7 (−0.11)
m: 91.1, 38.8	+0.6 (+**0.02**)	−23.1 (−0.60)	−14.9 (−0.38)	+21.9 (+0.56)
n: 53.3, 35.2	−2.0 (−0.06)	+3.1 (+0.09)	+0.9 (+**0.03**)	+2.7 (+0.08)
o:155.2, 39.5	+24.6 (+0.62)	+18.9 (+0.48)	+17.4 (+0.44)	+15.8 (+**0.40**)
p: 82.3, 51.2	−13.0 (−**0.25**)	−22.6 (−0.44)	−27.0 (−0.53)	−49.3 (−0.96)
q: 75.2, 44.9	−49.1 (−1.09)	−2.5 (+**0.06**)	−3.2 (−0.07)	+9.8 (+0.22)
r:177.7, 46.1	+4.8 (+**0.10**)	+10.5 (+0.23)	+12.5 (+0.27)	+8.3 (+0.19)
s: 86.4, 26.0	+8.3 (+0.32)	+8.8 (+0.34)	+3.2 (+**0.12**)	+3.6 (+0.14)
Average absolute	10.7 (0.30)	9.4 (0.25)	7.7 (**0.20**)	9.5 (0.24)

Tables 4.1-4.3. The true composite cell values as well as their standard deviations are given. The distance of the measured α from $\hat{\alpha}$ is given in degrees ($\epsilon_\alpha = \alpha - \hat{\alpha}$) and number of true standard deviations ($\zeta_\alpha = \epsilon_\alpha / \sigma_{\hat{\alpha}}$). The distance of the measured κ from $\hat{\kappa}$ is given in dimensionless units ($\epsilon_\kappa = \kappa - \hat{\kappa}$) and also in number of true standard deviations ($\zeta_\kappa = \epsilon_\kappa / \sigma_{\hat{\kappa}}$). Finally, the distance of the measured A from $\hat{A}$ is given in pixels ($\epsilon_A = A - \hat{A}$) and also in number of true standard deviations ($\zeta_A = \epsilon_A / \sigma_{\hat{A}}$). A graphical comparison of the results is shown in Fig. 4.3-4.5.

Table 4.2: Distance from $\hat{\kappa}$; for each method in dimensionless units ($\epsilon_\kappa = \kappa - \hat{\kappa}$) and number of standard deviations ($\zeta_\kappa = \epsilon_\kappa / \sigma_{\hat{\kappa}}$) for test images. Best results shown in bold.

Image: $\hat{\kappa}$, $\sigma_{\hat{\kappa}}$	SD	LSEF	AM	CAS
	ϵ_κ (ζ_κ)	ϵ_κ (ζ_κ)	ϵ_κ (ζ_κ)	ϵ_κ (ζ_κ)
a: 1.05, 0.44	+0.08 (+0.18)	+0.07 (+0.16)	+0.05 (+**0.11**)	+0.38 (+0.86)
b: 1.76, 0.56	+0.11 (+**0.20**)	−0.44 (−0.79)	−0.28 (−0.50)	+0.95 (+1.70)
c: 2.28, 0.67	−0.58 (−0.87)	−0.15 (−**0.22**)	−0.18 (−0.27)	−0.78 (−1.16)
d: 1.50, 0.60	−0.22 (−0.37)	+0.10 (+0.17)	+0.19 (+0.32)	−0.07 (−**0.12**)
e: 1.69, 0.54	−0.19 (−0.35)	+0.00 (−**0.00**)	−0.01 (−0.02)	−0.26 (−0.48)
f: 1.51, 0.36	−0.06 (−0.17)	−0.05 (−0.14)	−0.10 (−0.28)	−0.08 (−**0.04**)
g: 1.85, 0.56	−0.26 (−0.46)	−0.08 (−0.14)	+0.02 (+**0.04**)	−0.02 (−0.26)
h: 1.33, 0.54	+0.04 (+**0.07**)	+0.12 (+0.22)	+0.14 (+0.26)	−0.23 (−0.43)
i: 1.26, 0.44	−0.09 (−**0.20**)	+0.18 (+0.41)	+0.12 (+0.27)	−0.09 (−**0.20**)
j: 1.10, 0.32	+0.13 (+0.41)	+0.04 (+0.12)	+0.00 (+**0.00**)	−0.10 (−0.31)
k: 1.37, 0.42	+0.13 (+0.31)	+0.02 (+**0.05**)	+0.05 (+0.12)	−0.20 (−0.48)
l: 1.64, 0.47	+0.14 (+0.30)	+0.06 (+0.13)	−0.04 (−**0.09**)	−0.53 (−1.13)
m: 1.27, 0.57	−0.09 (−0.16)	+0.09 (+0.16)	−0.01 (−**0.02**)	+0.16 (+0.28)
n: 1.48, 0.67	−0.12 (−**0.18**)	−0.13 (−0.19)	−0.12 (−**0.18**)	+0.35 (+0.52)
o: 1.33, 0.33	−0.04 (−**0.12**)	+0.11 (+0.33)	+0.17 (+0.52)	+0.34 (+1.03)
p: 1.22, 0.58	−0.07 (−**0.12**)	−0.08 (−0.14)	−0.09 (−0.16)	+0.07 (+**0.12**)
q: 1.11, 0.36	+0.20 (+0.56)	−0.01 (−**0.03**)	−0.02 (−0.06)	+0.32 (+0.89)
r: 1.15, 0.41	−0.02 (−**0.05**)	+0.05 (+0.12)	+0.07 (+0.17)	+0.68 (+1.66)
s: 1.46, 0.47	−0.02 (−0.04)	−0.25 (−0.53)	−0.24 (−0.51)	−0.03 (−**0.06**)
Average absolute	0.14 (0.27)	0.11 (0.21)	0.10 (**0.20**)	0.30 (0.62)

Table 4.3: Distance from $\hat{A}$; in pixels ($\epsilon_A = A\text{-}\hat{A}$) and number of standard deviations ($\zeta_A = \epsilon_A/\sigma_{\hat{A}}$) for test images. Best results shown in bold.

Image: $\hat{A}$, $\sigma_{\hat{A}}$	SD	LSEF	AM	CAS
	ϵ_A (ζ_A)	ϵ_A (ζ_A)	ϵ_A (ζ_A)	ϵ_A (ζ_A)
a: 2744, 1267	−627 (−0.49)	−2135 (−1.69)	−1119 (−0.88)	−180 (**−0.14**)
b: 2027, 708	−344 (−0.49)	−871 (−1.23)	+15 (**+0.02**)	+624 (+0.88)
c: 479, 200	−129 (−0.65)	−206 (−1.03)	+76 (**+0.38**)	+1272 (+6.36)
d: 619, 222	+298 (+1.34)	−440 (−1.98)	−176 (**−0.79**)	+1027 (+4.63)
e: 504, 224	+317 (+1.42)	−220 (−0.98)	+161 (**+0.72**)	+602 (+2.69)
f: 1010, 395	−290 (−0.73)	+461 (−1.17)	+262 (**+0.66**)	+2700 (+6.84)
g: 658, 264	−148 (−0.56)	−352 (−1.33)	−113 (**−0.43**)	−123 (−0.47)
h: 297, 144	+228 (+1.58)	−160 (−1.11)	+16 (**+0.11**)	+81 (+0.56)
i: 472, 259	+250 (+0.97)	−289 (−1.12)	−51 (**−0.20**)	+578 (+2.23)
j: 452, 243	+308 (+1.27)	−308 (−1.27)	−120 (**−0.49**)	+433 (+1.78)
k: 502, 205	+444 (+2.17)	−292 (−1.42)	−18 (**−0.09**)	+2114 (+10.31)
l: 476, 189	+389 (+2.06)	−297 (−1.57)	−79 (**−0.42**)	+320 (+1.69)
m: 412, 204	+348 (+1.71)	−223 (−1.09)	+63 (**+0.31**)	+240 (+1.18)
n: 373, 158	+215 (+1.36)	−219 (−1.39)	−36 (**−0.23**)	+262 (+1.66)
o: 203, 79	+70 (+0.89)	−92 (−1.16)	+52 (**+0.66**)	+141 (+1.78)
p: 507, 276	+308 (+1.12)	−304 (−1.10)	+12 (**+0.04**)	+541 (+1.96)
q: 208, 82	+527 (+6.43)	−76 (**−0.93**)	+77 (+0.94)	+1287 (+15.70)
r: 211, 74	+583 (+7.88)	−101 (−1.36)	+11 (**+0.15**)	+1008 (+13.62)
s: 1243, 560	−19 (**−0.03**)	−1004 (−1.79	−594 (−1.06)	+1225 (+2.19)
Average absolute	307 (1.74)	424 (1.30)	161 (**0.45**)	777 (4.03)

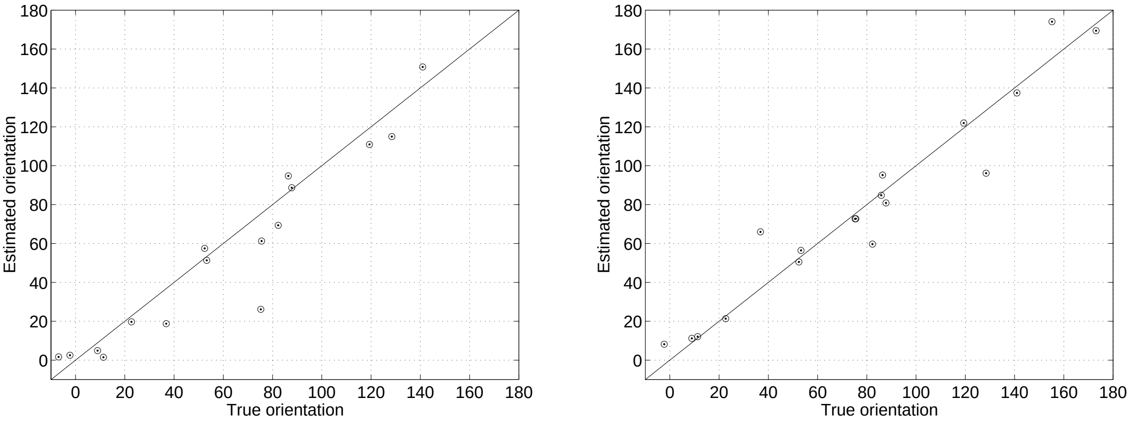

(a) SD (b) LSEF

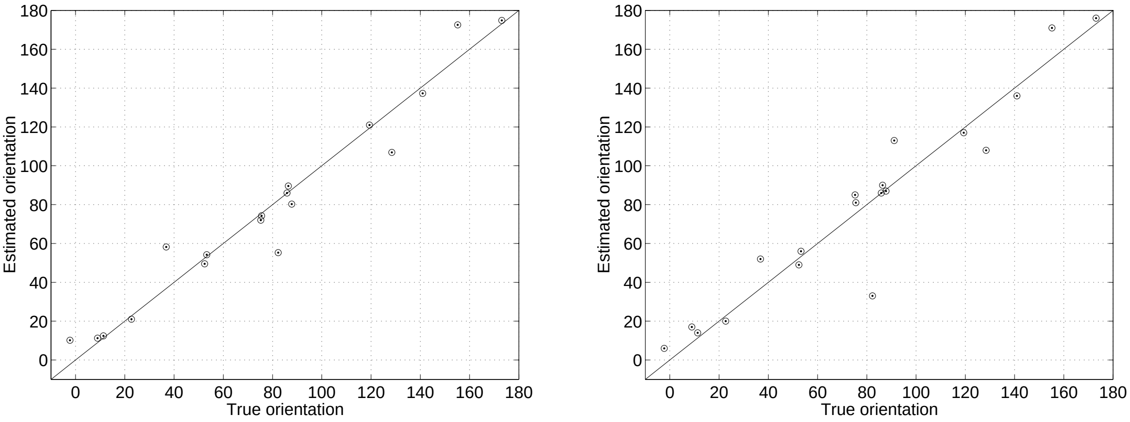

(c) AM (d) CAS

Figure 4.3: α estimation results for 19 test images for each of the four methods.

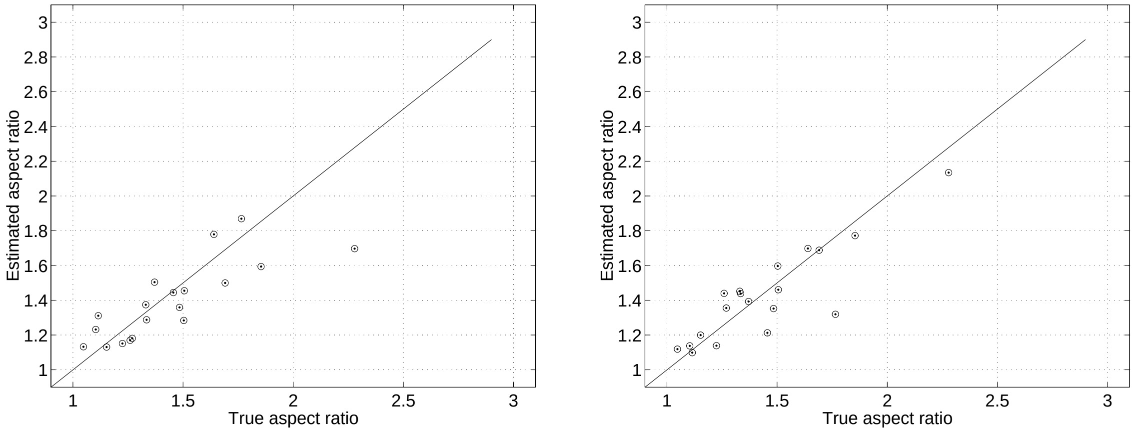

(a) SD (b) LSEF

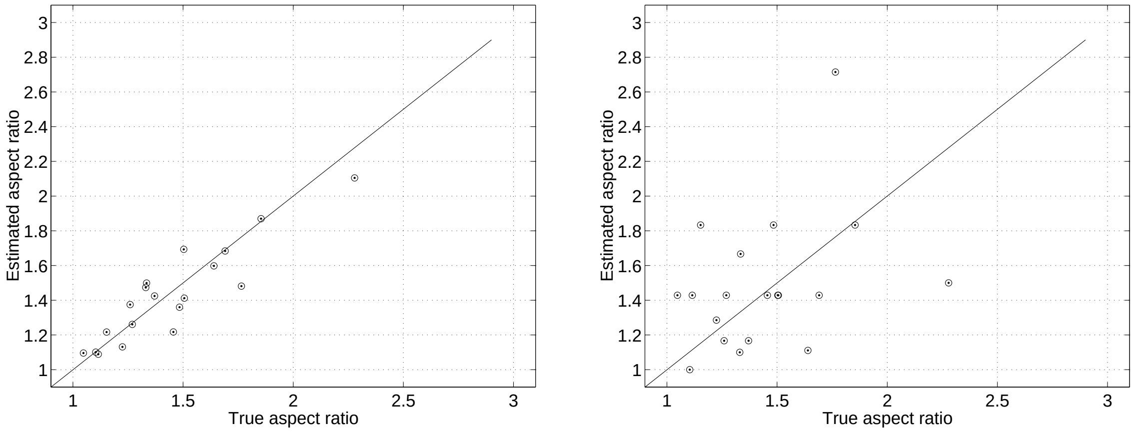

(c) AM (d) CAS

Figure 4.4: κ estimation results for 19 test images for each of the four methods.

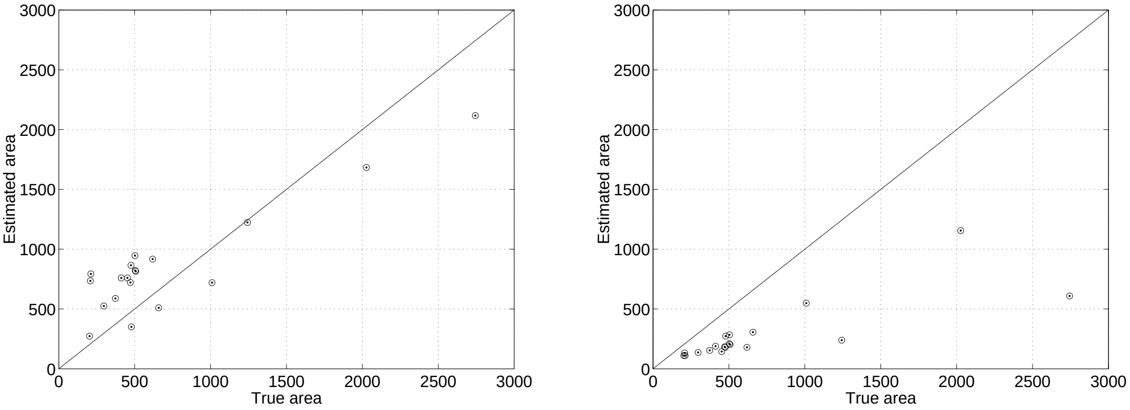

(a) SD

(b) LSEF

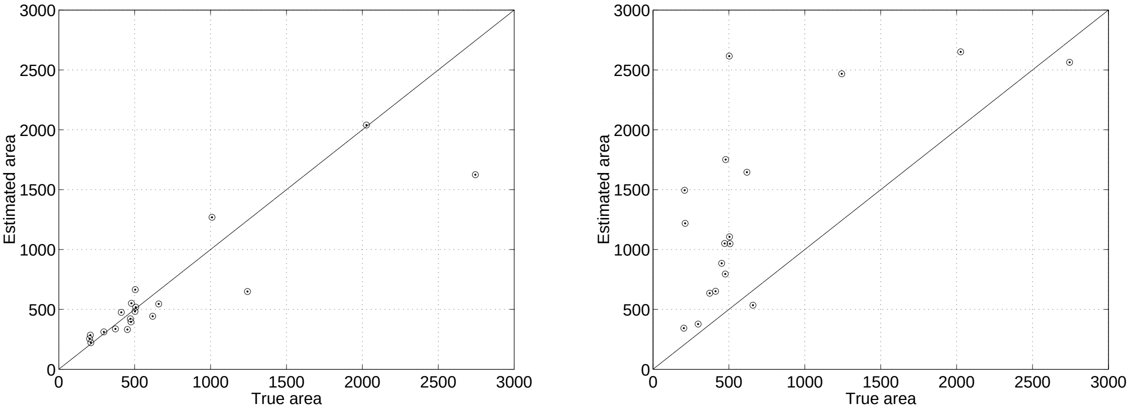

(c) AM

(d) CAS

Figure 4.5: A estimation results for test 19 images for each of the four methods.

4.3.1 Performance objective

When the errors are measured in standard deviations, the natural variability of the cells are taken into account. These measures therefore provide the most complete information on the performance of a method. The performance objective decided by the author is for the estimated α, κ and A values to consistently lie within a distance of one true standard deviation from their true values (i.e. $|\zeta_\alpha| < 1$, $|\zeta_\kappa| < 1$ and $|\zeta_A| < 1$) for all image types.

4.3.2 Orientation discussion

All methods estimated $\hat{\alpha}$ with similar accuracy and achieved the performance objective of $|\zeta_\alpha| < 1$ for all 19 test images. The AM method was most accurate, with the lowest average $|\zeta_\alpha|$ (0.20) and the lowest maximum $|\zeta_\alpha|$ (0.53 for image (p)). All methods had some difficulty with images (a) and (j) because for both cases $\hat{\kappa} \leq 1.1$; meaning that the composite cell is almost circular and hence $\hat{\alpha}$ is difficult to detect, even visually. In contrast, the orientations in images (c) and (g) are easily detected by all the methods because in both cases $\hat{\kappa} \geq 1.8$. The SD method yielded the least accurate overall results due to the difficulty in correctly segmenting the individual cells because of inconsistent cell boundary visibility.

4.3.3 Aspect ratio discussion

The AM and LSEF methods were the most accurate in estimating $\hat{\kappa}$. The AM method had the lowest average $|\zeta_\kappa|$ (0.20) and the lowest maximum $|\zeta_\kappa|$ (0.52 for image (o)). The LSEF method performed almost identically as it also achieved the performance objective and had an average $|\zeta_\kappa| = 0.21$. This was not surprising given the similarity of these methods: both these methods use every response value in the $|F|$ image to estimate $\hat{\kappa}$.

The SD and CAS methods were less accurate. The SD method achieved $|\zeta_\kappa| = 0.27$, but was again affected by inconsistent cell boundary visibility. The CAS method had large errors in estimating $\hat{\kappa}$ for a number of the images because the $|F|$ images often did not have well-defined peaks in the major and minor axes directions.

Each method underestimated $\hat{\kappa}$ for images with inconsistent cell boundary visibility such as (f) and (s). This can be explained by noting that the methods would estimate $\kappa = 1$ (the lowest possible value) if cell boundaries in an image were completely obscured.

4.3.4 Area discussion

The AM method was most accurate for estimating $\hat{A}$. This method had an average $|\zeta_A| = 0.45$ and also achieved the performance objective for all test images except image (s). This error was due to the inconsistent boundary visibility in this image. As mentioned earlier, however, there remained a correlation between the error in A and the true minor axis frequency (Fig. 3.8b). This was addressed by enforcing a minor axis range (to the user) of 10-22 c.p.i., in which the average $|\zeta_\kappa| < 0.5$. This excludes images (a) and (q) which both yielded poor results. This minor axis frequency range is large enough to be easily followed by the user during manual selection of the analysis window.

The SD method performed less accurately and tended to overestimate the cell area. This is because the watershed-detected cells were often bigger than the actual cells due to incomplete cell boundary visibility, such as in images (k) and (l). This caused the cell segments to 'spill over' into other cell segments. The large errors in images (p) and (q) are due to another problem: because of the large number of cells in these images, the minimum-cells threshold ($n_{min} = 25$) was not large enough to sufficiently characterise the image. This demonstrates a weakness in the SD method: the need to define five individual parameters.

Both the LSEF and CAS methods performed much more poorly. The LSEF method consistently underestimated $\hat{A}$ due to its squared distance measure. This caused high-frequency energy to overly affect the fitted ellipse and pull it outwards. The CAS on average overestimated $\hat{A}$. This method would often mistakenly identify peaks in the $|F|$ image that were too low in frequency that were present because of low-frequency noise.

4.3.5 Algorithmic complexity

Each of the four tested algorithms has similar complexity: $O(N^2)$ where $N \times N$ are the dimensions of the input image. For the SD method, each step has an $O(N^2)$ complexity: the local contrast enhancement step performs a fixed task on each of the N^2 pixels, the watershed method does the same [36] and the segment analysis step processes each N^2 pixel to find the centroid of each cell segment. The LSEF method requires $O(N^2)$ steps to build the $N^2 \times 6$ 'design matrix' [17] and the AM method requires $O(N^2)$ steps to calculate I_u, I_u and I_{uv}. The CAS method requires $O(N^2)$ steps to calculate the correlation at each image rotation. Interestingly, the Gabor filter method, which showed the poorest accuracy, also has the highest complexity: $O(N^3)$ since N^2 multiplications have to be done to calculate the filter response at each $O(N)$ frequencies.

The average computation times for the 19 test images (in MATLAB® 6 on a Pentium 4 @2.4GHz) were also similar (Table 4.4). The LSEF method had the smallest average computation time at 3.3 seconds.

Table 4.4: Average computation times for test images for each method in MATLAB® 6 on a Pentium 4 computer at 2.4 GHz. Preprocessing time (including performing FFT) is included for frequency domain methods.

	SD	LSEF	AM	CAS
Time (s)	4.9	3.3	6.1	6.6

Table 4.5: Paired T-test results for test images for each method at a 5% confidence level. Associated probability given in parentheses.

Parameter	SD	LSEF	AM	CAS
α	Pass (0.32)	Pass (0.68)	Pass (0.68)	Pass (0.85)
κ	Pass (0.26)	Pass (0.59)	Pass (0.63)	Pass (0.63)
A	Pass (0.07)	FAIL (<0.01)	Pass (0.25)	FAIL (<0.01)

4.3.6 Bias test

A paired Student's T-test was applied to the results to determine whether any method was biased in its calculation of any of the composite cell parameters. The employed test determined whether the means of the calculated and true measurements were statistically different at a confidence level of 5% (Tab. 4.5).

As shown, all measurements were unbiased except the $\hat{A}$ estimations for the LSEF and CAS methods. As discussed earlier, the LSEF method consistently underestimated A because the fitted ellipse was being pulled outwards by high-frequency energy in the $|F|$ image due to its squared-distance measure. The CAS method would often mistakenly identify peaks in the $|F|$ image that were too low in frequency that were present because of low-frequency noise. This resulted in an average overestimation of $\hat{A}$.

4.3.7 Time-lapse testing

To further test the accuracy while also testing the stability of the four methods, each was applied to set of 11 time-lapse images in which a patch of embryo tissue was stretched. Two example images are shown here (Fig. 4.6), and the entire set is shown in Appendix A (Fig. A.5). Plots were produced that show the true and estimated composite cell parameters for each method over all 11 images (Fig. 4.7).

The time-lapse images used for this test were sub-images cut manually from a set of

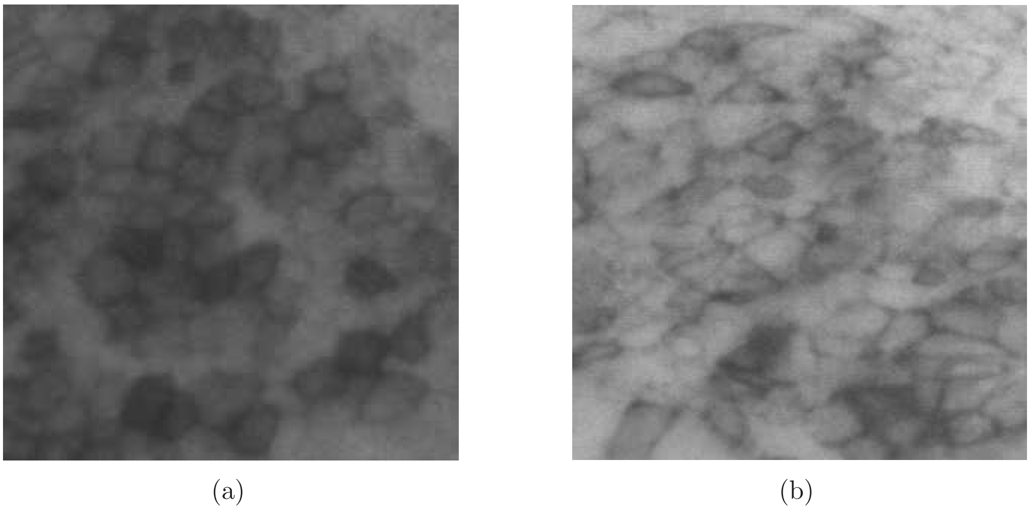

(a) (b)

Figure 4.6: (a) First and (b) last images in the time-lapse image series.

larger images. The sub-images were selected to track roughly the same patch of cells. An automated window-selection method would make future time-lapse tests more meaningful. A method has already been developed to track a patch of cells [8], but it is recommended that this be expanded to allow for the automatic selection of a square window.

The orientation results show that each method performed similarly accurately. All methods had difficulty estimating the orientation for the first six images. This was expected, however, given the low $\hat{\kappa}$ for the first six images (Fig. 4.7b), and the instability of the true orientation. Each method became stable and accurate as $\hat{\kappa}$ increased over the last few images.

The aspect ratio results demonstrate the stability and accuracy of the LSEF and AM methods and the instability of the SD and CAS methods. The similarity in the LSEF and AM $\hat{\kappa}$ estimations is expected given the similarity of these methods.

The area results show that the AM method alone yields stability and accuracy. The CAS method method was unstable and overestimated $\hat{A}$, because it mistakenly identified

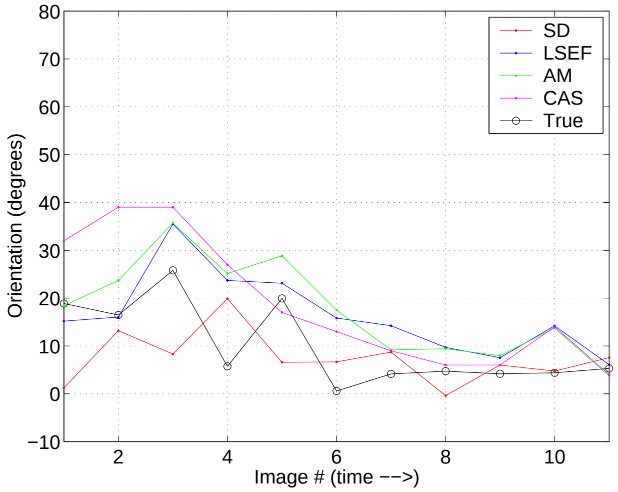

(a) α

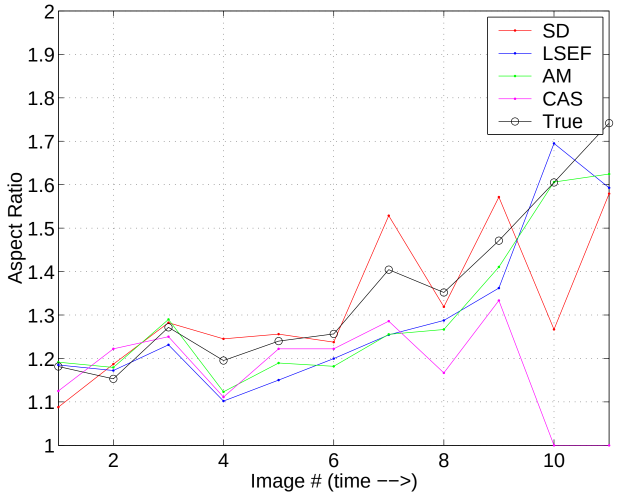

(b) κ

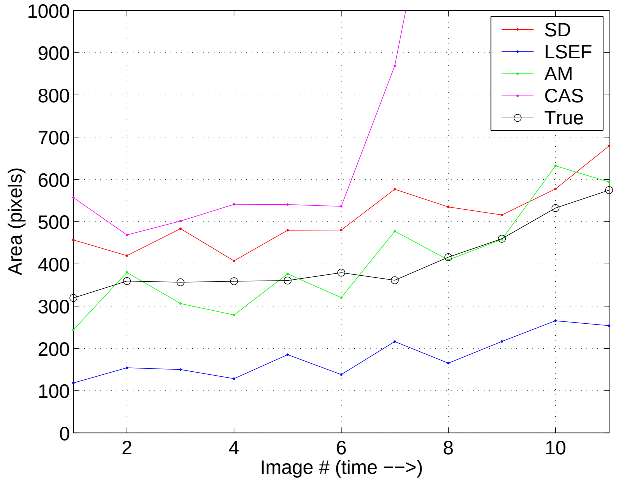

(c) A

Figure 4.7: Composite cell parameter estimations and truths for time-lapse images.

peaks in the $|F|$ image at low frequencies. The LSEF method was stable but underestimated $\hat{A}$ due to its squared distance measure. The SD method is smooth, but overestimated $\hat{A}$ due to the lack of consistent boundary visibility.

These results demonstrate that the LSEF and AM methods give the most stable estimations of the composite cell parameters. The AM method alone, however, showed the highest accuracy.

4.3.8 Method of choice

The AM method achieved the most accurate results in estimating each composite cell parameter. The AM results were also not biased and were stable. The AM method also had a similar complexity and computation time as the other methods. For these reasons, the AM method was chosen as the preferred method to estimate $\hat{\alpha}$, $\hat{\kappa}$ and $\hat{A}$.

4.4 Confidence measures

The user of this algorithm should be alerted if the results for a particular image were expected to be poor. An attempt was made to construct 'confidence measures' to add to the AM method for each of the composite cell parameters.

It was first thought that if the standard deviations of each composite cell parameter ($\sigma_{\hat{\alpha}}, \sigma_{\hat{\kappa}}$ and $\sigma_{\hat{A}}$) could be estimated, then this would allow a measure of confidence to be assigned to each parameter result. The standard deviations, however, could not be easily estimated due to their combined effect on the $|F|$ image. Each parameter variation would just add to the overall blurriness of the $|F|$ image (Fig. 4.8). Adding to this problem was the net blurring effect of variation in cell edge visibility, irregular cell shapes and noise in the spatial image.

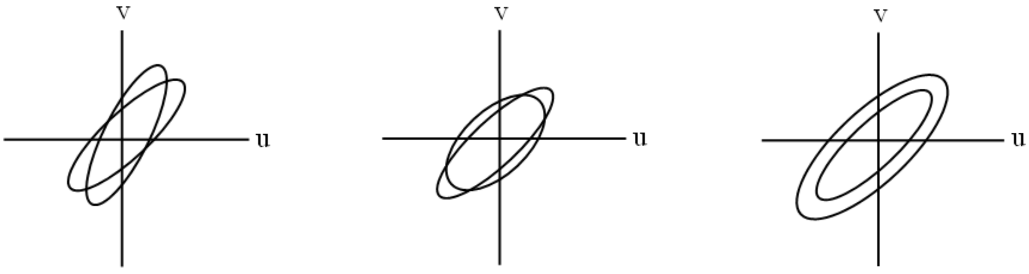

(a) (b) (c)

Figure 4.8: Affect on ellipse in $|F|$ image due to variation in composite cell parameters ((a) α, (b) κ, (c) A). There is not enough information in the $|F|$ image to individually measure these variations.

Two features were then found that had a relationship with the AM parameter estimation errors. The estimated aspect ratio, κ was found to be correlated with the absolute orientation error, $|\zeta_\alpha|$, and the local image contrast was found to be correlated with the absolute aspect ratio error, $|\zeta_\kappa|$. The average c.p.i. was originally correlated with the area error, but this correlation was removed to improve the accuracy of the area estimation, as described in Section 3.2.4.

4.4.1 Orientation confidence measure

The absolute error of the estimated α was found to decrease with increasing κ (Fig. 4.9). This is understandable: as the average aspect ratio of the cells increases, the dominant orientation can more easily be detected. Based on the results presented in this thesis, confidence levels were assigned to α for certain ranges of κ (Tab. 4.6). This confidence measure requires no extra computation since κ has already been estimated.

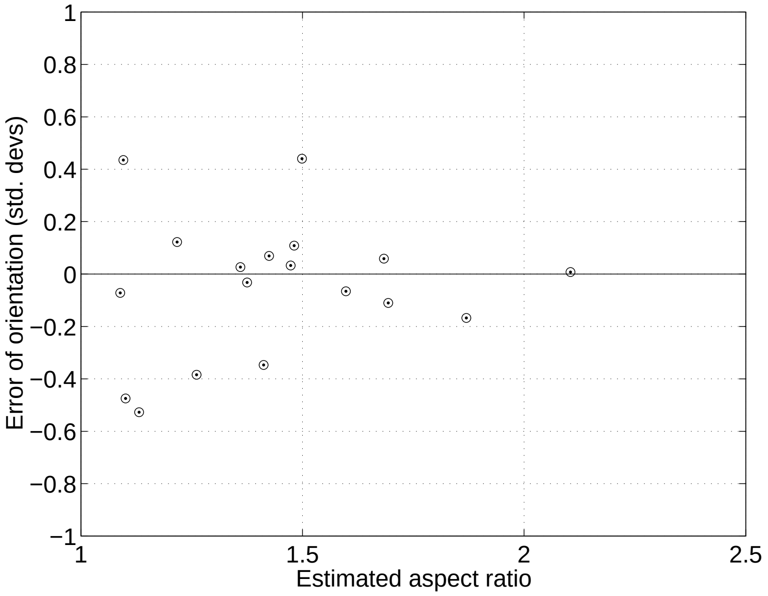

Figure 4.9: Relationship between κ and ζ_α (AM case).

Table 4.6: Confidence levels for α, based on estimated κ.

Feature range	Confidence level
$\kappa > 1.5$	High
$1.3 < \kappa < 1.5$	Medium
$1.1 < \kappa < 1.3$	Low
$\kappa < 1.1$	Very low

4.4.2 Aspect ratio confidence measure

The absolute error of the estimated κ was found to decrease with an increase in the average local image contrast (Fig. 4.10). This relationship can be explained as follows: local contrast gives a measure of the local visibility of cell boundaries; therefore, the *average* local contrast gives a measure of the average visibility of cell boundaries for the entire image. Inconsistent or low-contrast cell boundaries hide the shapes of the cells and make the estimation of κ difficult.

For this use, a new local contrast measure was developed. The Weber contrast measure (see [27]) was used as a foundation for the new contrast measure. The Weber measure defines the contrast, C, of a small object of uniform intensity, L, on a background also of uniform intensity:

$$C = \frac{\Delta L}{L} \tag{4.1}$$

where ΔL is the difference between the intensities of the object and background. The new local contrast measure at an image location (x,y) is:

$$C(x, y) = k \cdot ln \left(1 + \frac{\sigma(x, y)}{\mu(x, y)}\right) \tag{4.2}$$

where $\sigma(x, y)$ and $\mu(x, y)$ are the local standard deviation and mean, respectively, of the pixel intensities in a local window and k is a scalar. The average local contrast, C_{avg}, is simply the average of $C(x, y)$ over all (x,y).

This new contrast measure modifies the Weber measure in three ways. First, while both measures contain a dimensionless ratio of intensities, the new measure removes the implausible assumption that each window contains a small 'object' laying on a 'background'. The numerator in the new measure is a more robust $\sigma(x, y)$ which measures the

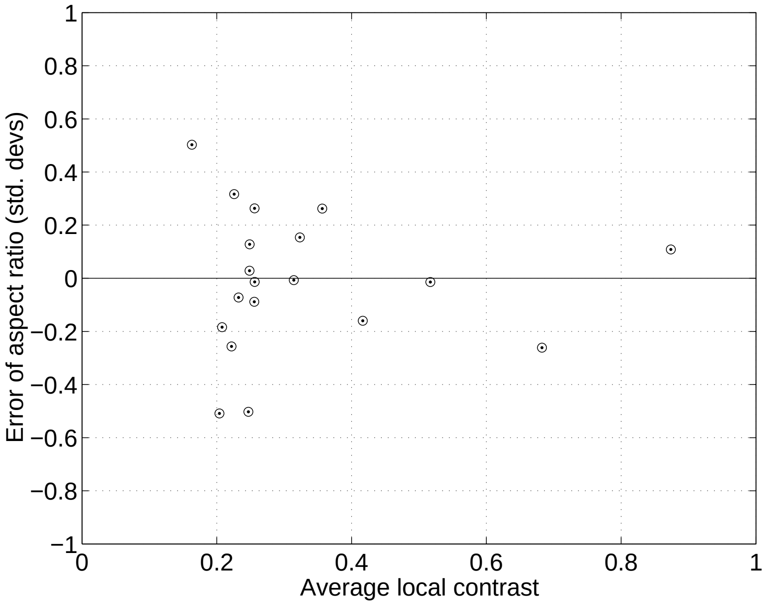

Figure 4.10: Relationship between average local contrast, C_{avg}, and ζ_κ (AM case).

overall variability of pixels in a window, while the denominator remains the same. Second, the natural logarithm operator is used to compress the scale of the measure. Finally, the scalar k makes $C(x, y) = 1$ correspond to a maximum contrast level. A window with a Gaussian intensity histogram (having 98% of its area between 0 and 1) is defined as having maximum contrast.

Implementing this new contrast measure required the choice of a window size and a complexity reduction. The window size was set to $1.5 * L_{major}$ so that it was an appropriate size for the image contents. Processing time was reduced by scaling the image and window dimensions by a factor of 0.25.

The relationship between C_{avg} and the average level of cell boundary visibility is demonstrated for selected images (Fig. 4.11). As before, confidence levels were assigned to κ for certain ranges of C_{avg} (Tab. 4.7). The average computation time to implement this mea-

Table 4.7: Confidence levels for κ, based on average image contrast measure, C_{avg}.

Feature range	Confidence level
$C_{avg} > 0.3$	High
$0.2 < C_{avg} < 0.3$	Medium
$0.1 < C_{avg} < 0.2$	Low
$C_{avg} < 0.1$	Very low

sure (in MATLAB® 6 on a Pentium 4 computer @2.4 GHz) for the 19 test images was 2.0 seconds.

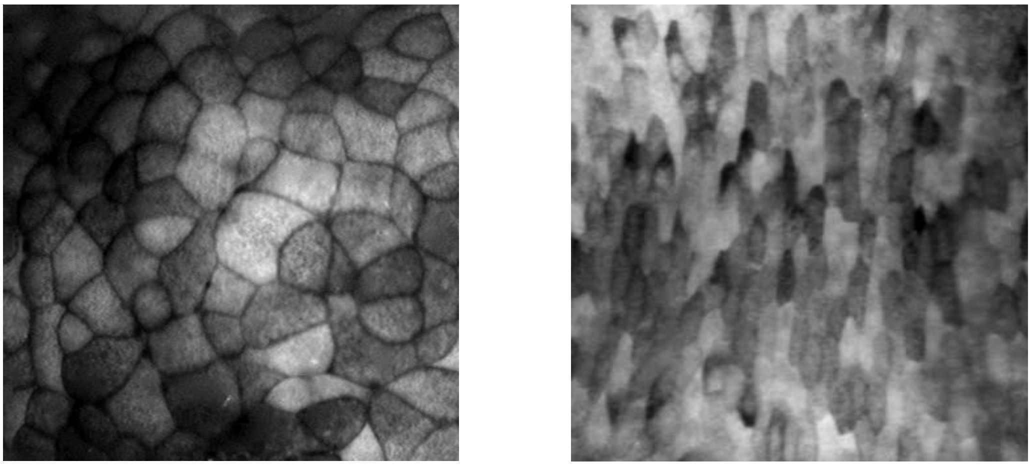

(a) $C_{avg} = 0.87$ (b) $C_{avg} = 0.68$

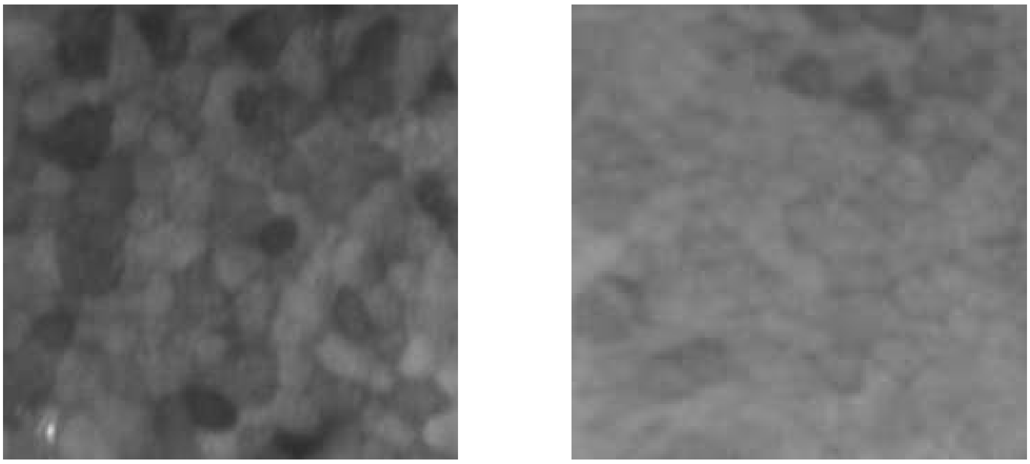

(c) $C_{avg} = 0.36$ (d) $C_{avg} = 0.16$

Figure 4.11: Average local contrast, C_{avg}, of selected images. The relationship between C_{avg} and the average visibility of cell boundaries is demonstrated.

Chapter 5

Conclusions

Five algorithms have been developed to estimate the average orientation, aspect ratio and area of cells in an image of biological tissue. One of these methods was based in the spatial domain, while the four others – least squares ellipse fitting (LSEF), area moments (AM), correlation and axes searching (CAS) and Gabor filters (GF) – estimated the composite cell parameters from the FFT magnitude response image. Four of these methods (all but the GF method) displayed some measure of success when applied to 19 test images and a separate set of 11 time-lapse images. The AM method alone showed high accuracy and stability.

The SD method was accurate in estimating $\hat{\alpha}$, but showed some instability in estimating $\hat{\kappa}$ and a bias in estimating $\hat{A}$. The instability and bias arose because the SD segmentation had trouble segmenting images with variable cell boundary visibility. Another weakness of the SD method was the need to define five parameters.

The LSEF method was accurate in estimating $\hat{\alpha}$, a close second behind the AM method for estimating $\hat{\kappa}$ but consistently underestimated $\hat{A}$. The underestimation occurred because the LSEF method was overly affected by high-frequency energy due to its squared distance measure. All parameter estimations were stable.

The AM method was most accurate in estimating each composite cell parameter, with an average $|\zeta_\alpha| = 0.20$, $|\zeta_\kappa| = 0.20$ and $|\zeta_A| = 0.45$. The area results were found to be correlated with the minor axis frequency, but this problem was contained by enforcing a 10-22 minor axis c.p.i. range for the test images. The AM method achieved the orientation and aspect ratio performance objectives ($|\zeta_\alpha| < 1$ and $|\zeta_\kappa| < 1$) for all the test images and the area objective ($|\zeta_A| < 1$) when the 10-22 c.p.i. range was enforced. When manual selection of the image is done this range is large enough to be easily followed by the user. The AM results were not biased and were stable for every composite cell parameter.

The CAS method was also accurate for estimating $\hat{\alpha}$, but struggled to estimate $\hat{\kappa}$ and $\hat{A}$. This was due to a lack of well-defined maxima in the $|F|$ images. The $\hat{A}$ estimations were also biased high, and the $\hat{\kappa}$ and $\hat{A}$ estimations were not stable.

5.1 Summary

Due to its consistently superior performance, the AM method should be used to estimate all the composite cell parameters: $\hat{\alpha}$, $\hat{\kappa}$ and $\hat{A}$. This method was robust to significant variation in illumination, cell pigmentation, cell shape and cell boundary visibility, and was unbiased and stable. The LSEF and SD methods also performed very well, while the CAS method struggled due to the lack of well-defined maxima in the $|F|$ images. Each method also had low complexity ($O(N^2)$) and a low computation time in Matlab.

Confidence measures were developed for the AM method. A large estimated κ was a good indication of confidence in the α results while a high local contrast measure indicated confidence in the κ results.

5.2 Future work

The AM method is ready to be used in a lab setting. However there are three potential areas for improvement of the algorithm. First, further efforts should be made to remove the $\hat{A}$ estimation range (10-22 c.p.i.). Ideally, the user should be able to choose any window size and get an accurate $\hat{A}$ estimation. Second, a confidence parameter for $\hat{A}$ should be developed. This would complete the functionality of the system by giving every composite cell parameter an associated confidence measure. Finally, an extension to the cell-tracking algorithm that is able to assign an appropriate square test window would aid further time-lapse testing. These additional tests could be used to further study the accuracy and stability of the AM method.

Appendix A

Test images, generated images

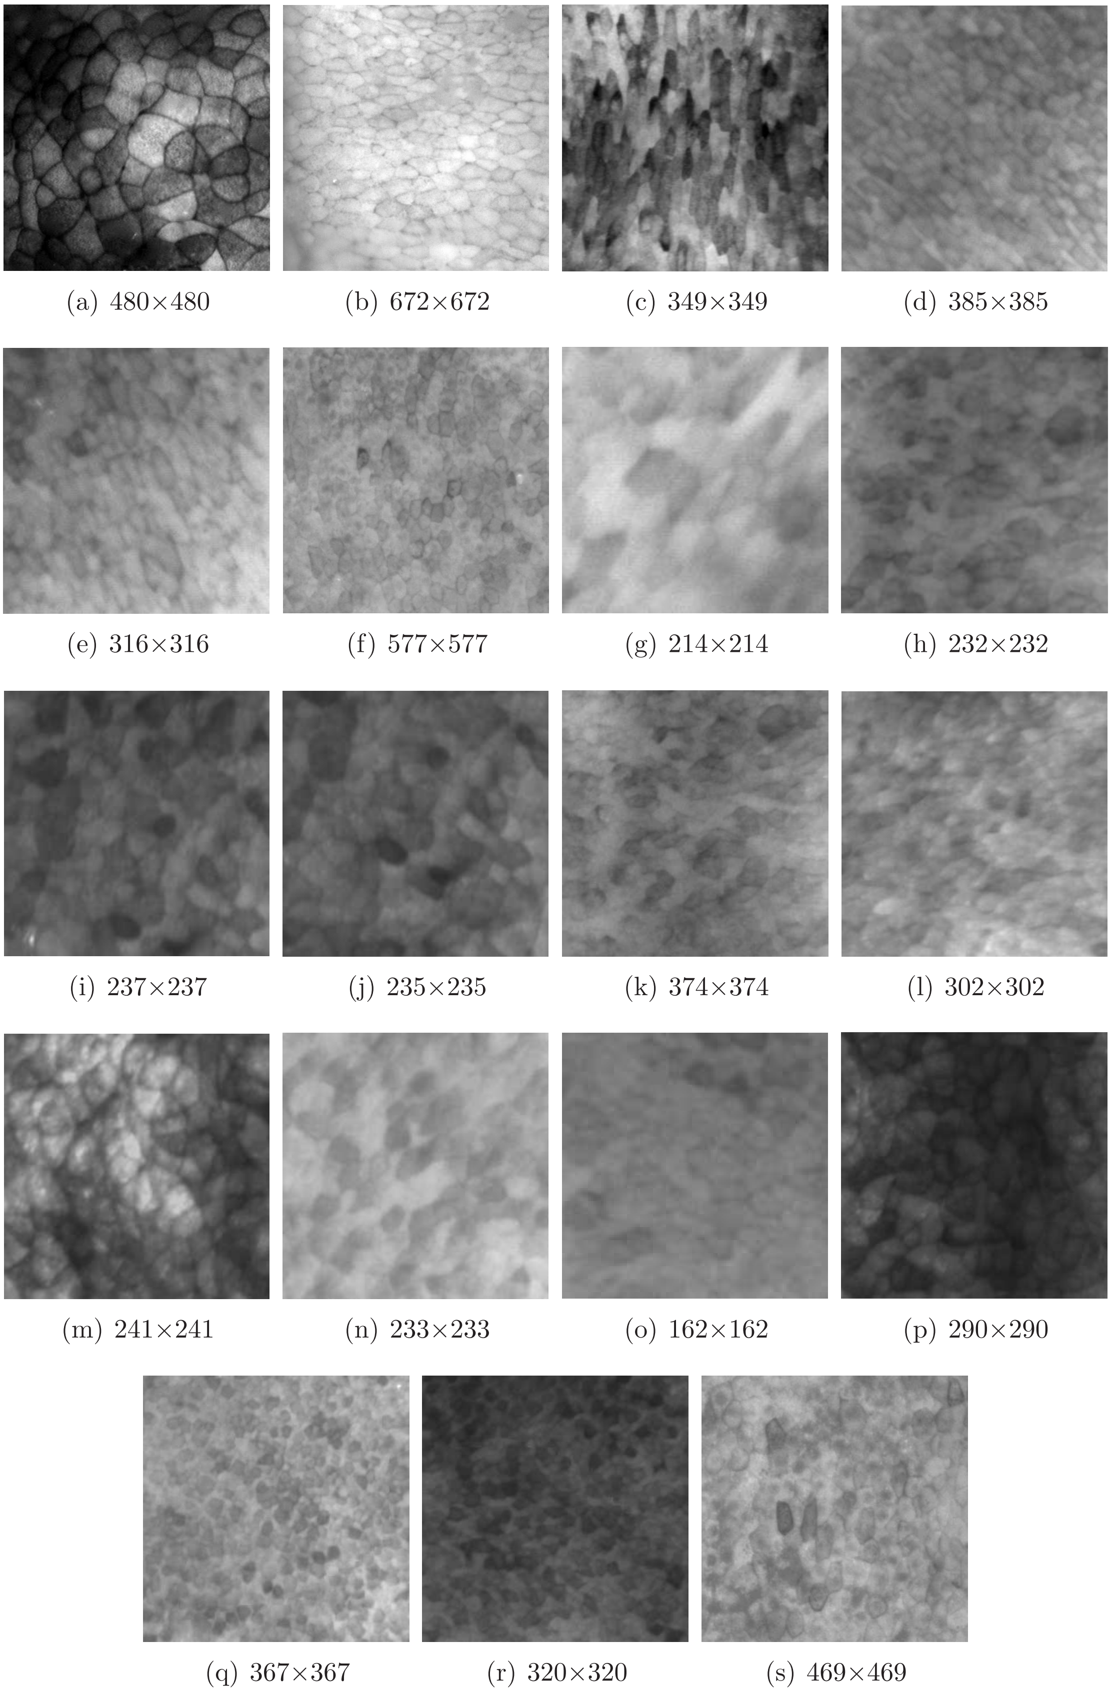

(a) 480×480 (b) 672×672 (c) 349×349 (d) 385×385

(e) 316×316 (f) 577×577 (g) 214×214 (h) 232×232

(i) 237×237 (j) 235×235 (k) 374×374 (l) 302×302

(m) 241×241 (n) 233×233 (o) 162×162 (p) 290×290

(q) 367×367 (r) 320×320 (s) 469×469

Figure A.1: Test images with image dimensions in pixels.

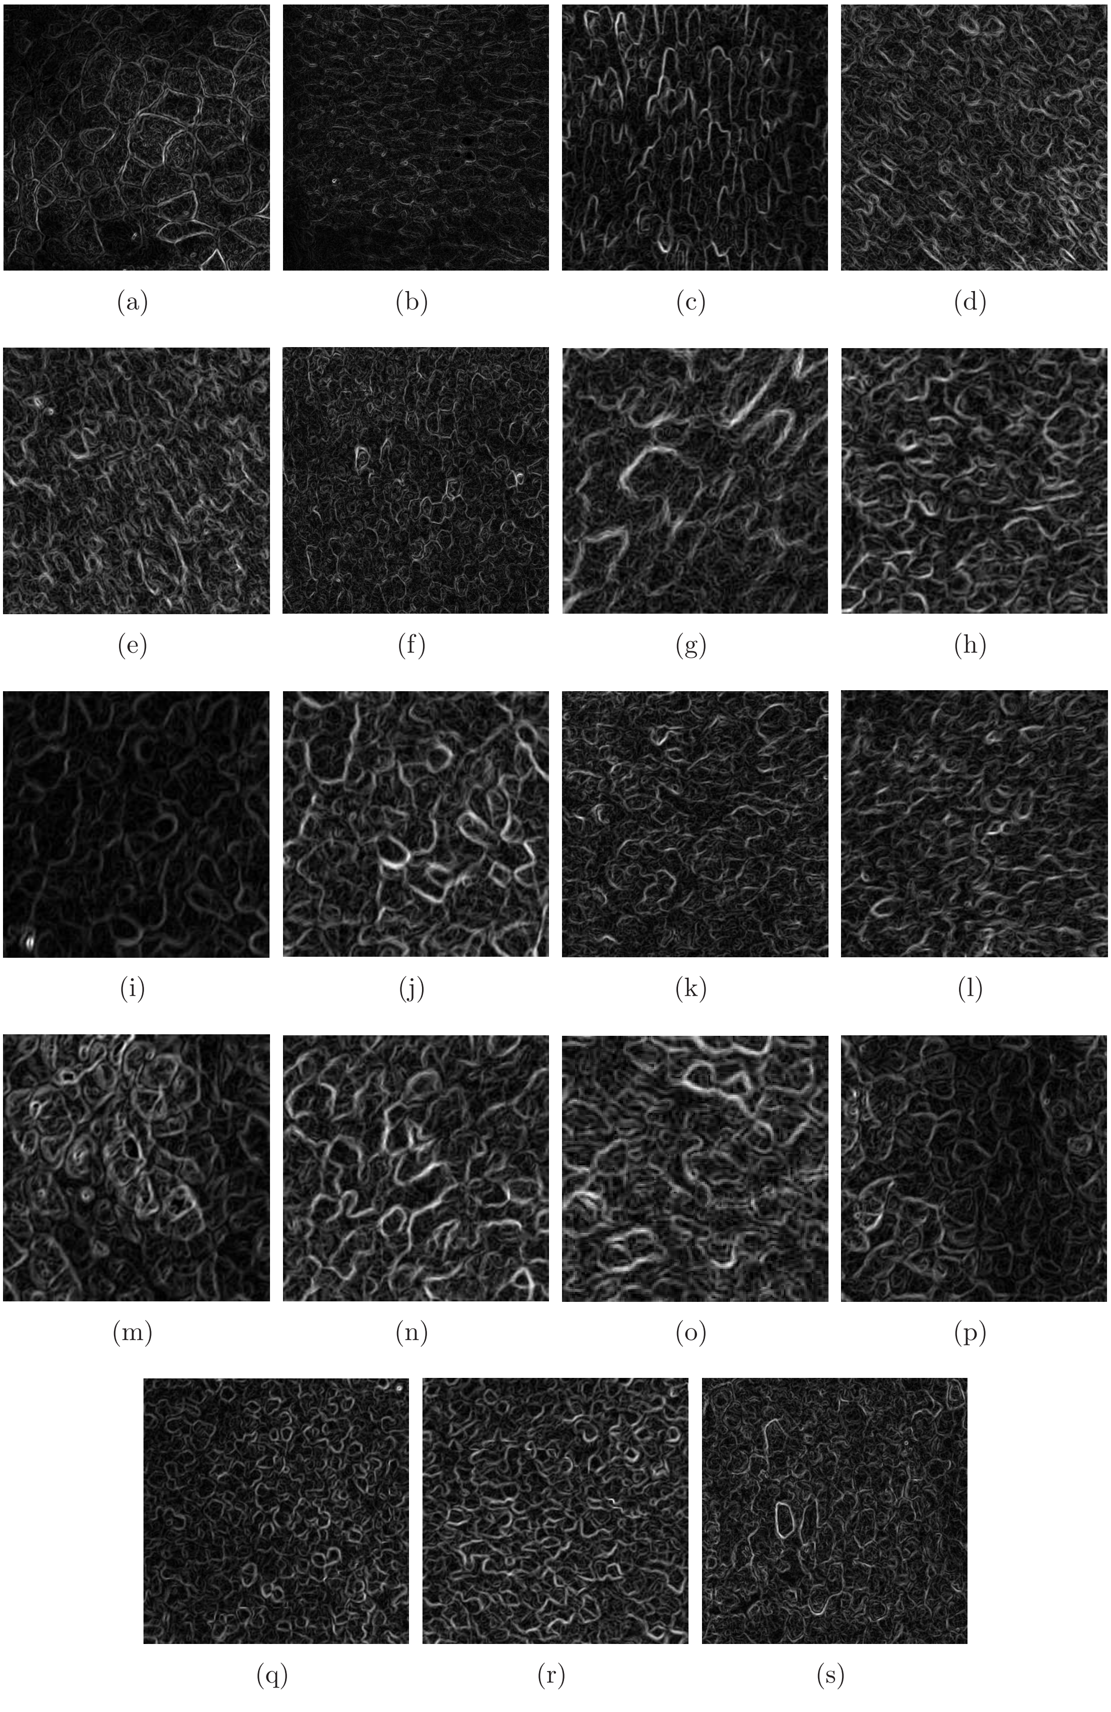

(a) (b) (c) (d)

(e) (f) (g) (h)

(i) (j) (k) (l)

(m) (n) (o) (p)

(q) (r) (s)

Figure A.2: Gradient magnitude images (normalised to [0, 1] to enhance visibility).

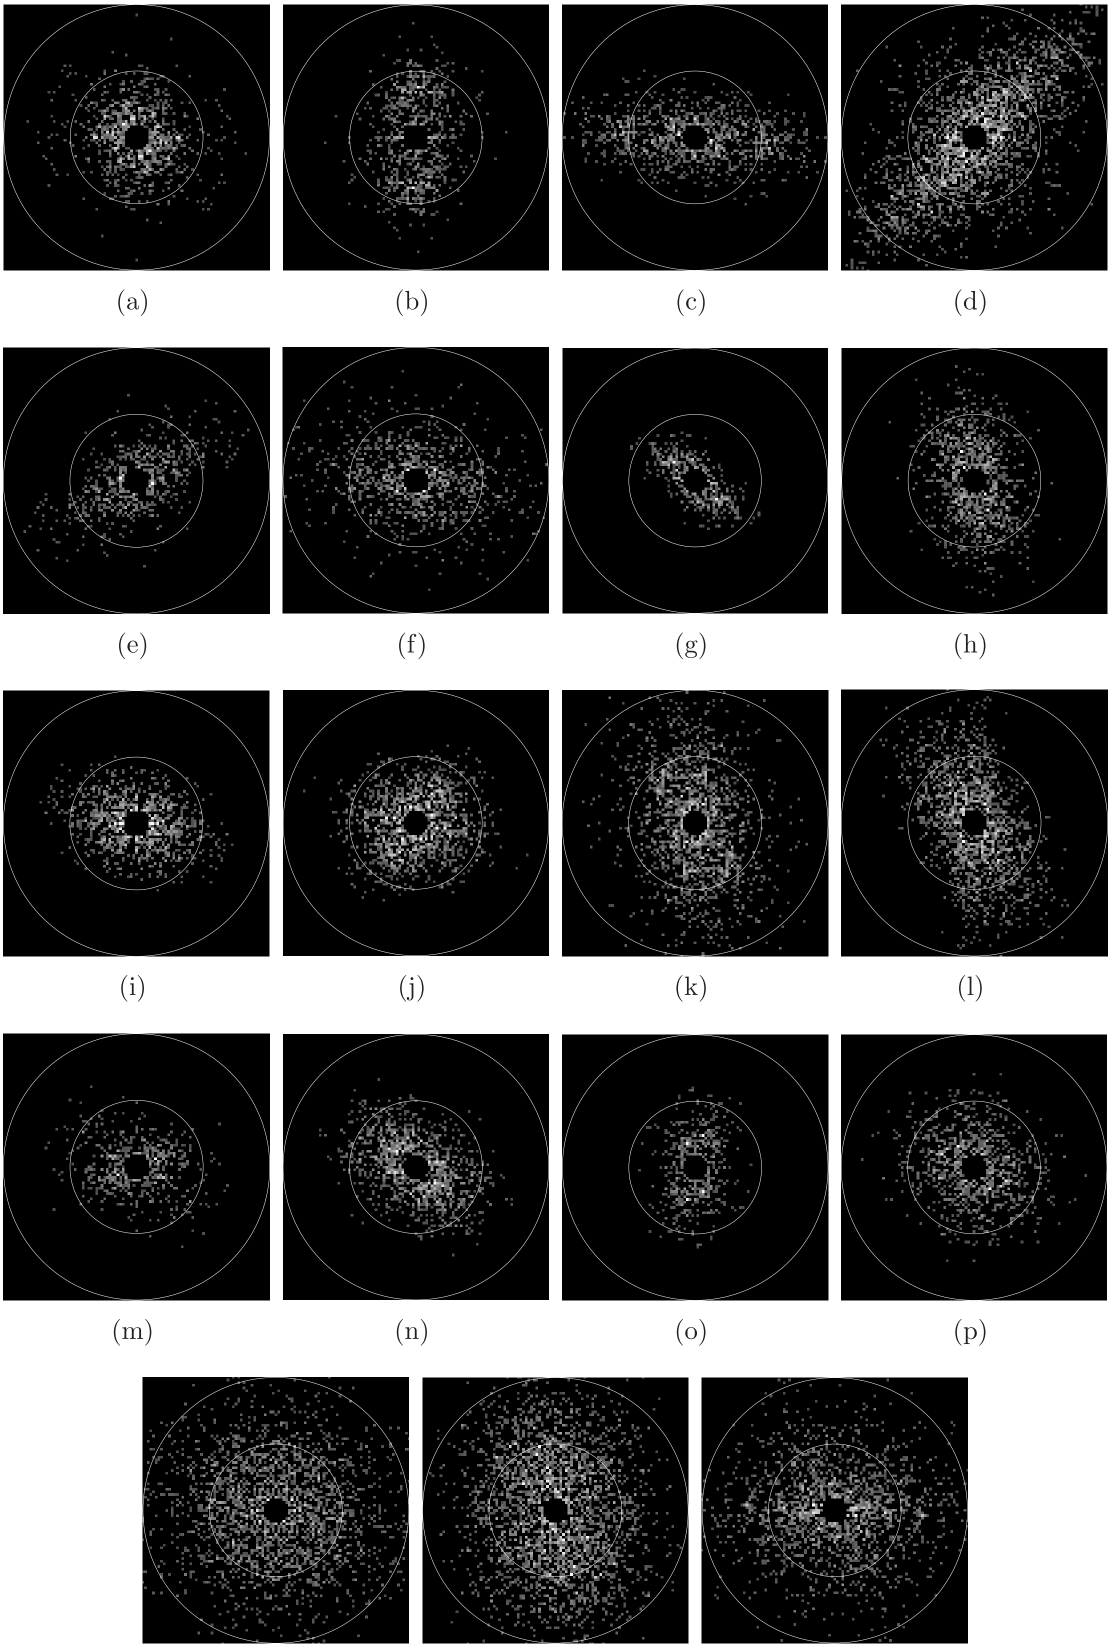

(a) (b) (c) (d)

(e) (f) (g) (h)

(i) (j) (k) (l)

(m) (n) (o) (p)

(q) (r) (s)

Figure A.3: Preprocessed $|F|$ images. Concentric circles show 25 and 50 c.p.i.

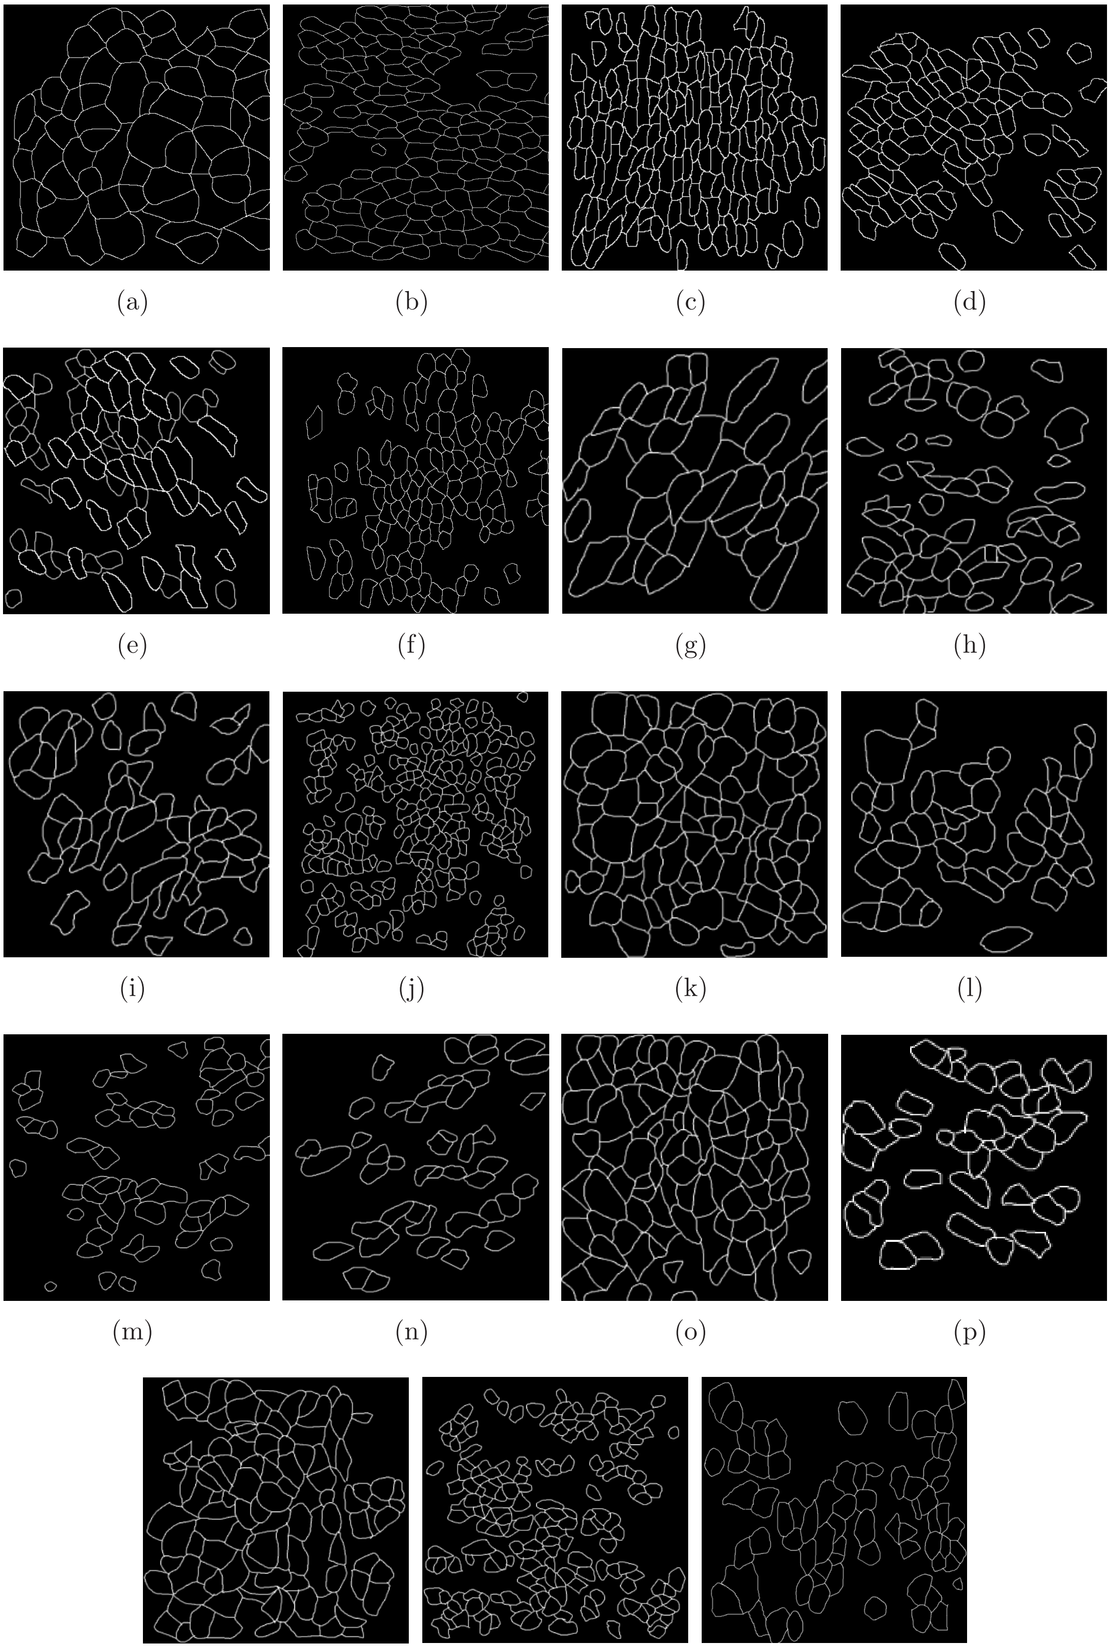

(a) (b) (c) (d)

(e) (f) (g) (h)

(i) (j) (k) (l)

(m) (n) (o) (p)

(q) (r) (s)

Figure A.4: Hand-segmented (truth) images.

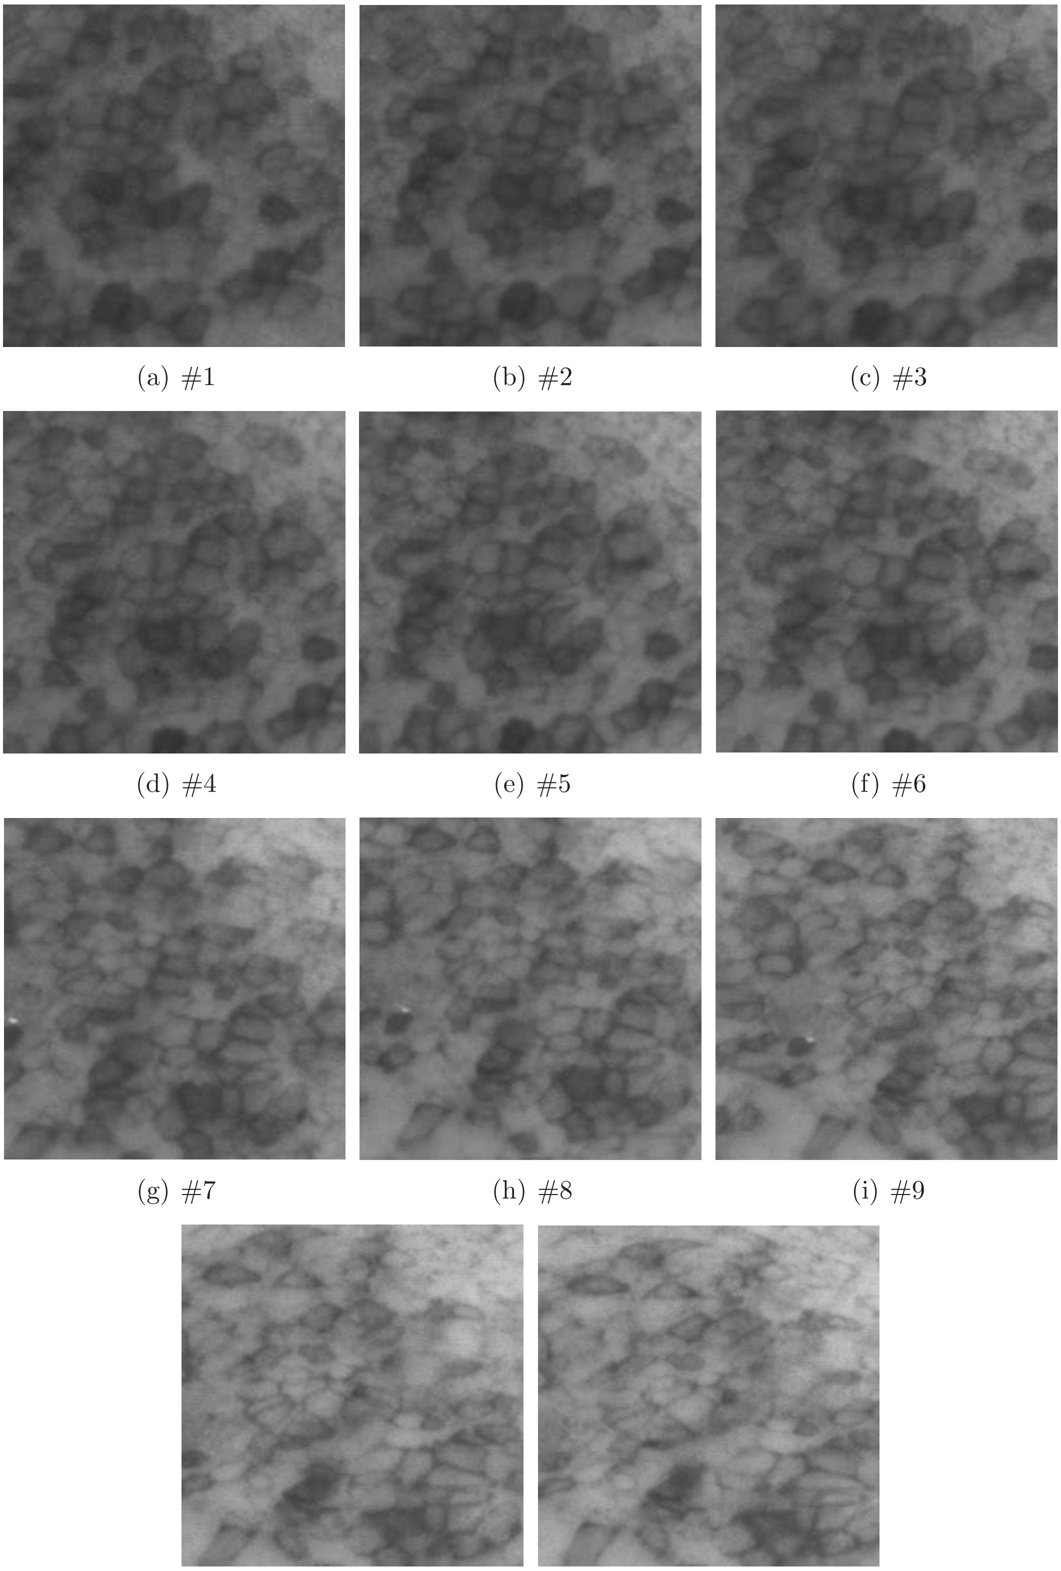

(a) #1 (b) #2 (c) #3

(d) #4 (e) #5 (f) #6

(g) #7 (h) #8 (i) #9

(j) #10 (k) #11

Figure A.5: Time-lapse test images.

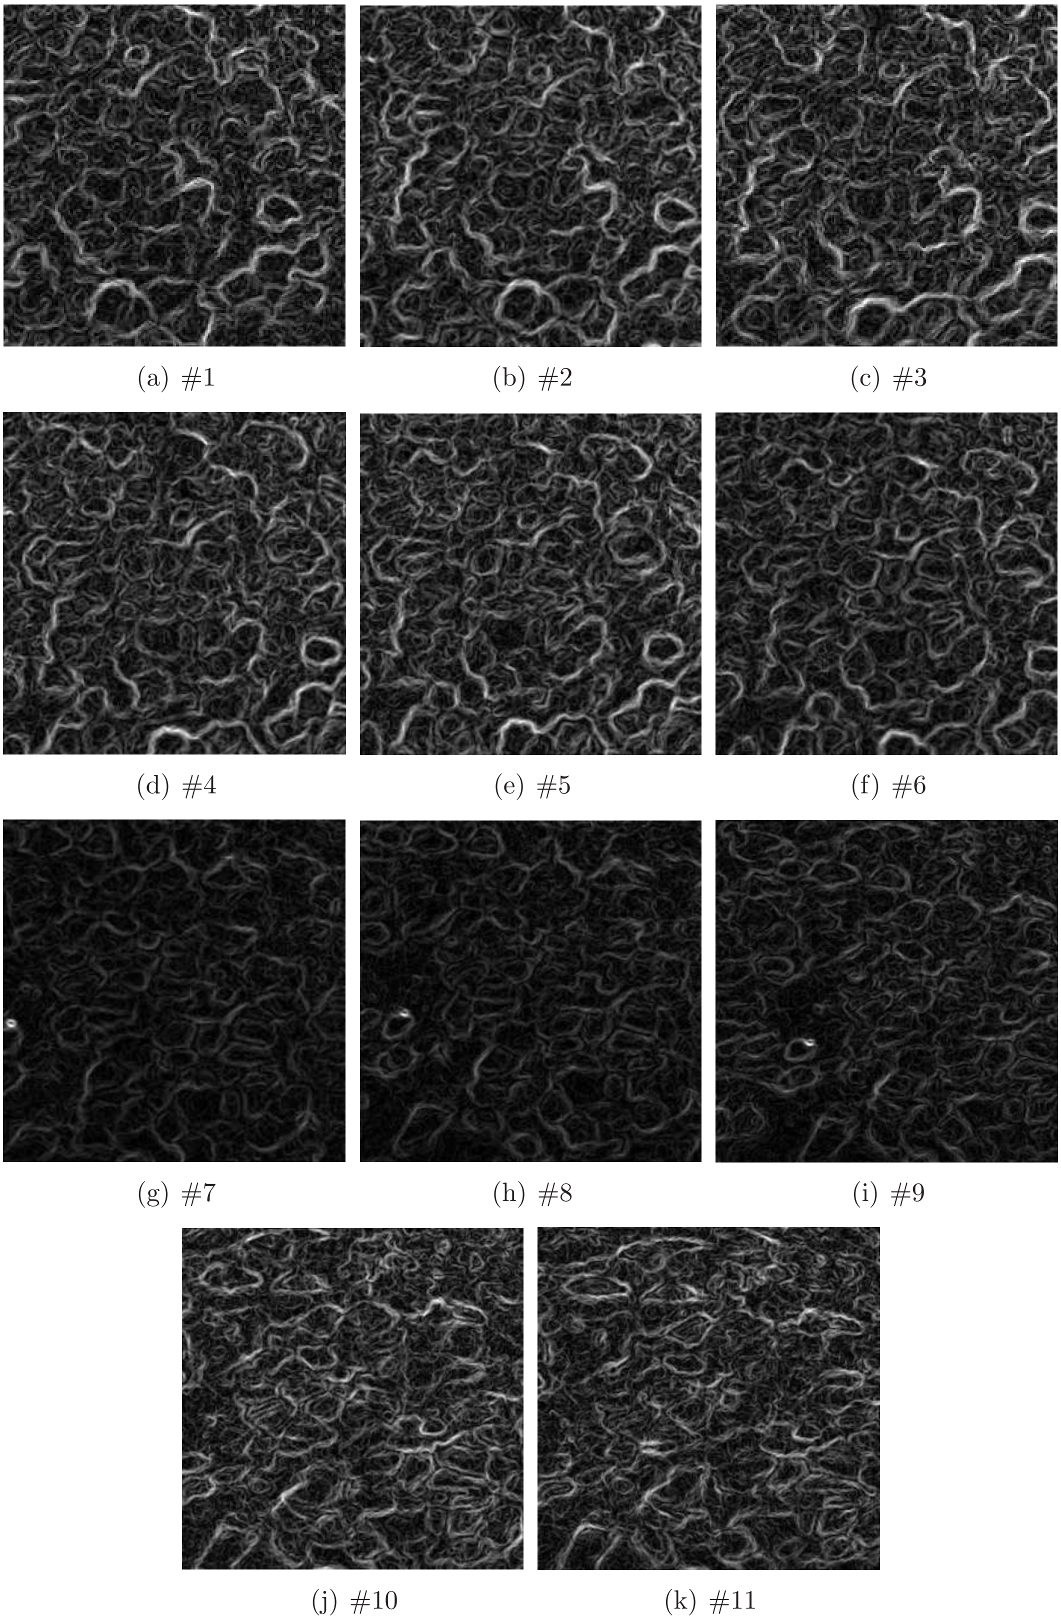

(a) #1 (b) #2 (c) #3

(d) #4 (e) #5 (f) #6

(g) #7 (h) #8 (i) #9

(j) #10 (k) #11

Figure A.6: Time-lapse gradient magnitude images (normalised to [0, 1] to enhance visibility.)

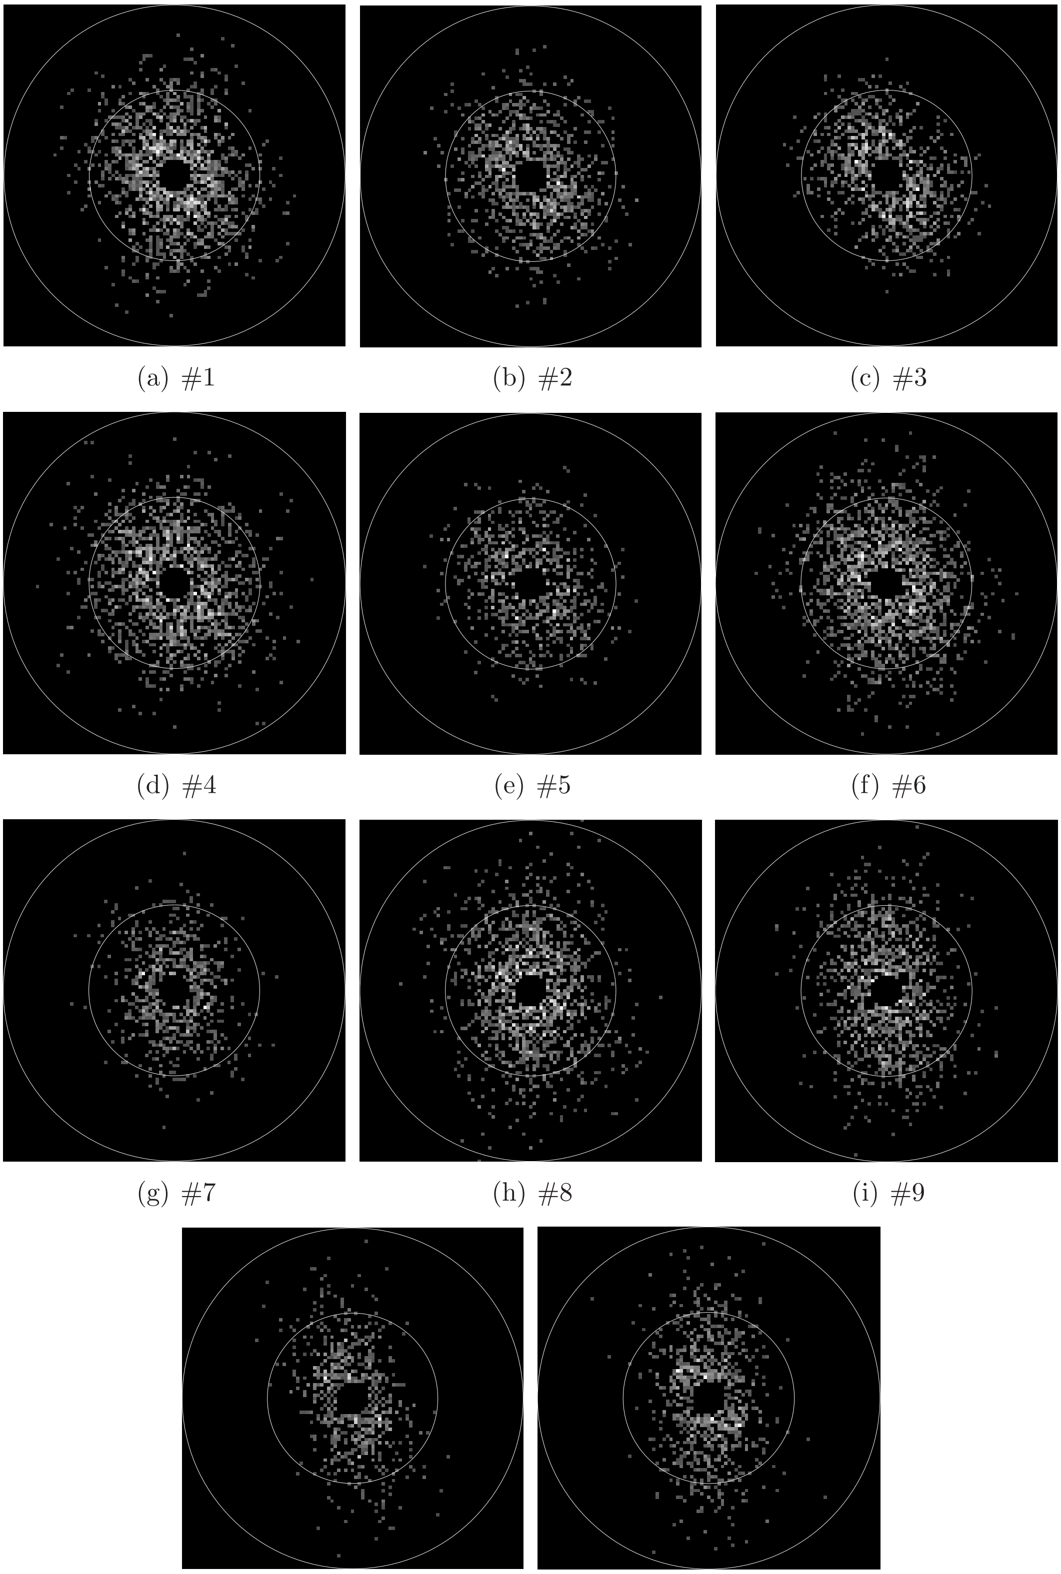

(a) #1 (b) #2 (c) #3

(d) #4 (e) #5 (f) #6

(g) #7 (h) #8 (i) #9

(j) #10 (k) #11

Figure A.7: Time-lapse preprocessed $|F|$ images. Concentric circles show 25 and 50 c.p.i.

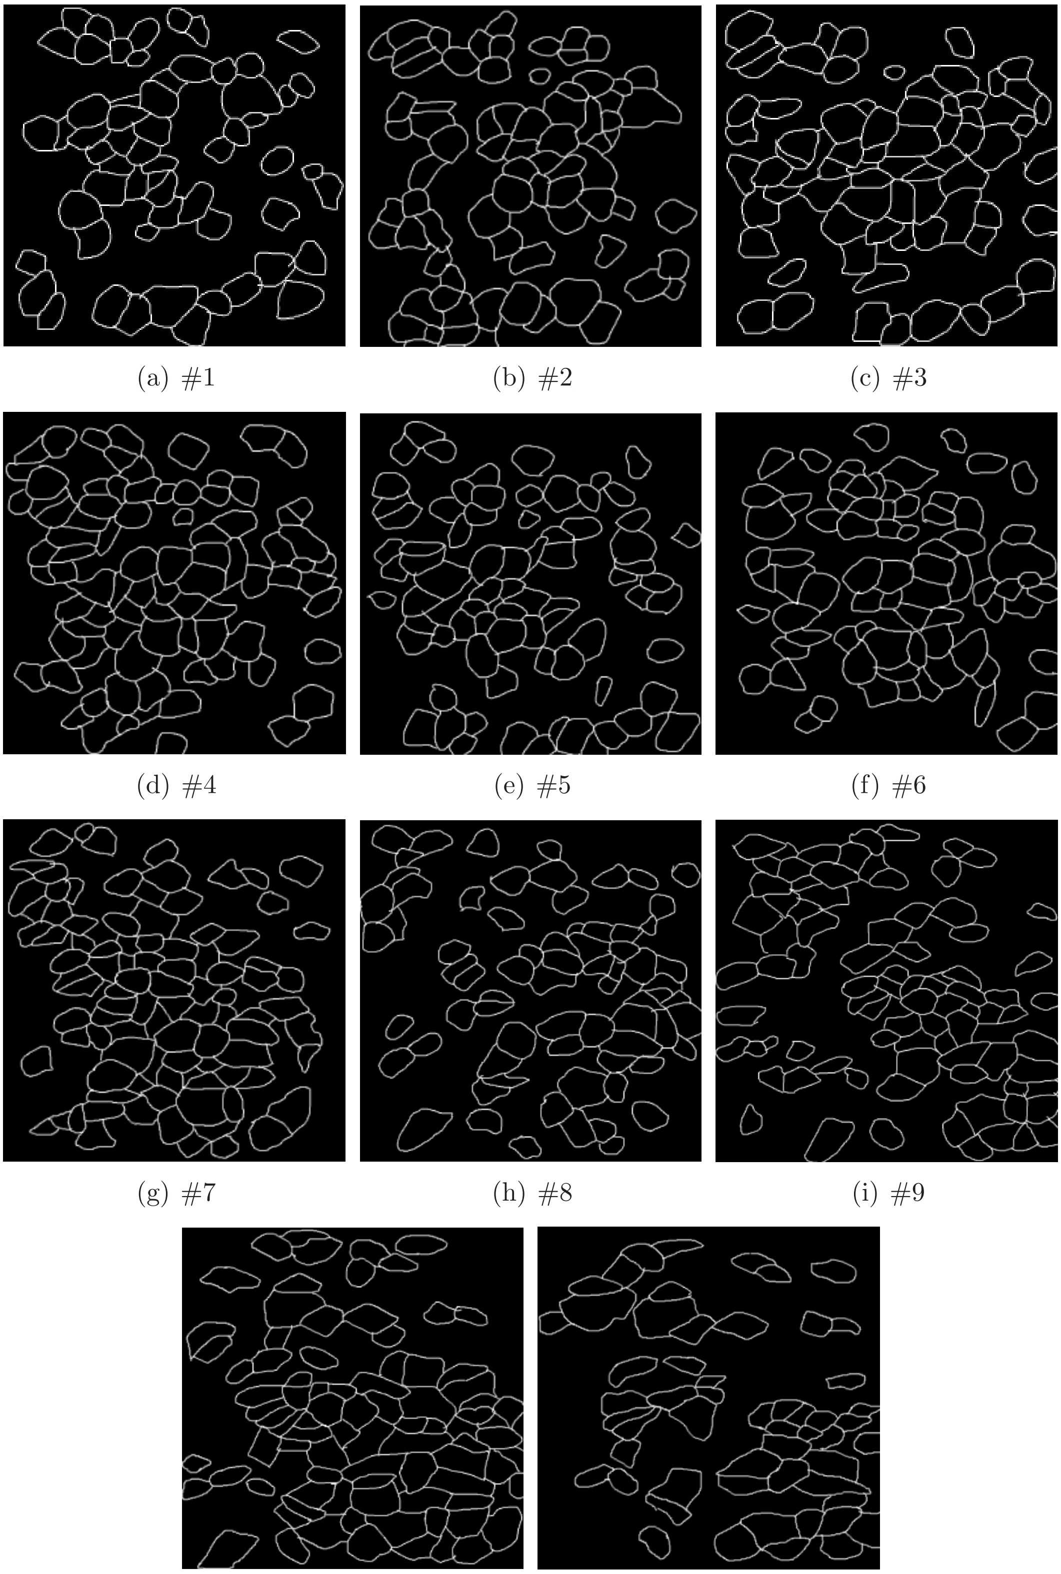

(a) #1 (b) #2 (c) #3

(d) #4 (e) #5 (f) #6

(g) #7 (h) #8 (i) #9

(j) #10 (k) #11

Figure A.8: Time-lapse hand-segmented (truth) images.

Bibliography

[1] L. Bertucco, G. Nunnari, C. Randieri, and A. Sacco V. Rizza. A cellular neural network based system for cell counting in culture of biological cells. In *Proceedings of the 1998 IEEE International Conference on Control Applications*, volume 1, pages 341–345, Sept 1998.

[2] J. Bigün, G. H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, Aug 1991.

[3] A. C. Bovik, M. Clark, and W. S. Geisler. Multichannel texture analysis using localized spatial features. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(1):55–73, 1990.

[4] P. Brodatz. *Textures: A Photographic Album for Artists and Designers.* Dover, 1966.

[5] G. W. Brodland. A computer model for reshaping of cells in epithelia due to in-plane deformation and annealing. *Comput Methods Biomech Biomed Engin*, 6(2):89–98, 2003.

[6] G. W. Brodland. Research homepage (http://www.civil.uwaterloo.ca/brodland/), Accessed: Nov. 2004.

[7] G. W. Brodland and H. H. Chen. Mechanics of cell sorting and envelopment. *Journal of Biomechanics*, 33(7):845–851, 2000.

[8] G. W. Brodland and J. H. Veldhuis. A deformable block-matching algorithm for tracking epithelial cells. *Image and Vision Computing*, 17:905–911, 1999.

[9] G. W. Brodland and J. H. Veldhuis. Computer simulations of mitosis and interdependencies between mitosis orientation, cell shape and epithelia reshaping. *Journal of Biomechanics*, 32:673–681, 2002.

[10] G. W. Brodland, J. H. Veldhuis, and D. I.-L. Chen. How cell shape affects the stresses in a cell sheet. In *ASME IMECE*, pages 79–80, Orlando FL, Nov. 5-10 2000.

[11] G. W. Brodland and C. J. Wiebe. Mechanical effects of cell anisotropy on epithelia. *Computer Methods in Biomechanics and Biomedical Engineering*, 7(2):91–99, April 2004.

[12] K. R. Castleman. *Digital Image Processing*. Prentice Hall, Englewood Cliffs, NJ, 1996.

[13] H. H. Chen and G. W. Brodland. Cell-level finite element studies of viscous cells in planar aggregates. *ASME Journal of Biomechanical Engineering*, 122:394–401, 2000.

[14] Y. Chen, K. Biddell, S. Aiying, P. A. Relue, and J. D. Johnson. An automatic cell counting method for optical images. In *Proceedings of the First Joint BMES/EMBS Conference*, volume 2, page 819, Oct 1999.

[15] D. A. Clausi and M. E. Jernigan. Designing Gabor filters for optimal texture separability. *Pattern Recognition*, 33:1835–1849, 2000.

[16] J. W. Cooley and J. W. Tukey. An algorithm for the machine computation of complex Fourier series. *Mathematics of Computation*, (19):297–301, April 1965.

[17] A. Fitzgibbon, M. Pilu, and R. B. Fisher. Direct least square fitting of ellipses. *Pattern Analysis and Machine Intelligence*, 21(5):476–480, 1999.

[18] R. C. Hibbeler. *Mechanics of Materials.* Prentice-Hall, New Jersey, 6th edition, 2005.

[19] A.K. Jain and F. Farrokhnia. Unsupervised texture segmentation using Gabor filters. *Pattern Recognition*, 24(12):1167–1186, 1991.

[20] M. Kass and A. Witkin. Analyzing oriented patterns. *Computer Vision, Graphics and Image Processing*, 37:362–385, 1987.

[21] H.-B. Kim and R.-H. Park. Extracting spatial arrangement of structural textures using projection information. *Pattern Recognition*, pages 237–245, March 1992.

[22] W. Lin, J. Xiao, and E. Micheli-Tzanakou. A computational intelligence system for cell classification. In *Information Technology Applications in Biomedicine, 1998. ITAB 98. Proceedings. 1998 IEEE International Conference on*, pages 105–109, May 16-17 1998.

[23] The Mathworks Inc. *Image Processing Toolbox: for use with Matlab (Version 3)*, 2001.

[24] T. Matsuyama, S.-I. Miura, and M. Nagao. Structural analysis of natural textures by fourier transformation. *Computer Vision, Graphics and Image Processing*, pages 347–362, Oct. 1983.

[25] D. P. Mukherjee, N. Ray, and S. T. Acton. Level set analysis for leukocyte detection and tracking. *IEEE Transactions on Image Processing*, 13(4):562–572, April 2004.

[26] G. Ongun, U. Halici, K. Leblebicioglu, V. Atalay, S. Beksac, and M. Beksac. Automated contour detection in blood cell images by an efficient snake algorithm. *Nonlinear Analysis, Theory, Methods and Applications*, 47(9):5839–5847, Aug 2001.

[27] E. Peli. Contrast in complex images. *J. Opt. Soc. Amer. A*, pages 2032–2040, 1990.

[28] S. M. Puddister. Estimating bulk geometrical properties of cellular structures. Master's thesis, University of Waterloo, Waterloo, ON, 2003.

[29] N. Ray, S. T. Acton, and K. Ley. Tracking leukocytes in vivo with shape and size constrained active contours. *IEEE Transactions on Medical Imaging*, 21(10):1222–1235, Oct 2002.

[30] P. L. Rosin. A note on the least squares fitting of ellipses. *Pattern Recognition Letters*, (14):799–808, Oct. 1993.

[31] T. Shimada, K. Kato, and A. Kamikouchi. Analysis of the distribution of the brain cells of the fruit fly by an automatic cell counting algorithm. *Physica A: Statistical Mechanics and its Applications*, 350(1):144–149, May 2005.

[32] A. Sourice, G. Plantier, and J.-L. Saumet. Autocorrelation fitting for texture orientation estimation. In *Proceedings of 2003 International Conference on Image Processing*, volume 1, pages 281–284, Sept 14-17 2003.

[33] V. V. Starovoitov, S.-Y. Jeong, and R.-H. Park. Texture periodicity detection: Features, properties and comparisons. (6):839–849, Nov. 1998.

[34] G. Vaidyanathan and P. M. Lynch. Texture direction analysis using edge counts. In *Proceedings of Southeastcon '89*, volume 2, pages 733–738, Apr 9-12 1989.

[35] J. H. Veldhuis, G. W. Brodland, C. J. Wiebe, and G. J. Bootsma. Robotic microscope system determines regional strain rates during amphibian embryo development. *Submitted to Annals of Biomedical Engineering*, Nov 1 2003.

[36] L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(6):583–598, June 1991.

[37] C. Wiebe and G. W. Brodland. Tensile properties of embryonic epithelia measured using a novel instrument. *Journal of Biomechanics. In press.*

[38] F. Ying-Lun, J. C. K. Chan, and R. T. Chin. Automated analysis of nerve-cell images using active contour models. *IEEE Transactions on Medical Imaging*, 15(3):353–368, June 1996.