Chapter 6

MEASUREMENT

The extraction of quantitative information from images is the end-point of much image analysis. The objective may simply be to:

- *count the number* of objects in a scene,
- compute their *areas*,
- or measure *distances* between objects.

Alternatively, it may be necessary to:

- characterize the *shapes* of objects,
- or summarize their *boundaries*,

in order to discriminate between the objects or to summarize information about them.

To illustrate the range of techniques to be covered, let us consider Fig 6.1(a). This shows regions of muscle extracted from the X-ray CT image of a sheep (Fig 1.7(c)). These regions of eye-muscle, or *longissimus dorsi*, the muscles which lie along a sheep's back, were obtained by thresholding the X-ray image and then using morphological openings and closings to reduce the roughness in the boundaries (see Fig 5.9(b)).

What measurements could we take from these regions?

- We could compute the cross-sectional areas of the muscles. These measures of size can be obtained simply by counting the number of pixels in both black regions in Fig 6.1(a). The counts are 1013 and 970 pixels for the left and right muscle respectively. Because 1 pixel is equivalent to an area of 1 mm^2 in this application, these counts convert to areas of 10.1 and 9.7 cm^2 for the two muscles.

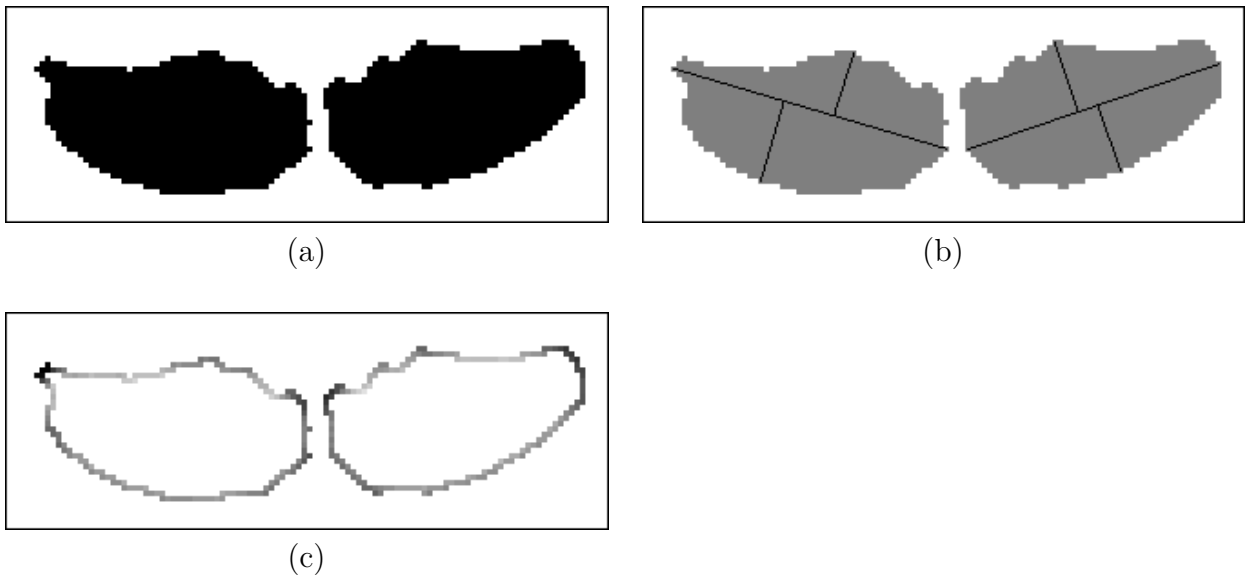

(a) (b)

(c)

Figure 6.1: Eye-muscles in the back of a sheep, from the X-ray CT image, obtained by thresholding the original image followed by morphological opening and closing, and used to illustrate different types of image measurement: **(a)** to estimate cross-sectional areas, **(b)** to measure elongation of the muscle cross-sections, as the length-to-breadth ratios, **(c)** to describe muscle boundaries in terms of local curvature — dark pixels indicate a small radius of curvature and lighter pixels indicate a larger radius.

- Alternatively, we may wish to quantify the shapes of the muscle cross-sections, by using a measure of *elongation*, for example. One definition of this measure is the ratio of the *length* (that is, the maximum distance between two pixels within a muscle) to the *breadth*. The line segments which correspond to the lengths are shown in Fig 6.1(b). There are several ways to define the breadth: the one we have used is the perpendicular distance from one of these lines to the furthest pixel on one side of the muscle, *plus* the distance to the furthest pixel on the other side. The line segments which correspond to the breadths are also shown in Fig 6.1(b). Bear in mind that we consider pixels to be points on a lattice, although we display them as square blocks, and therefore it is appropriate to measure distances between the *centres* of the blocks. For the left and right muscles, the length to breadth ratios are $52.2/27.9 = 1.87$ and $48.7/26.6 = 1.83$ respectively. (Further details of how these distances were obtained are given in §6.1.2.)

- A third possibility is to describe the *boundaries* of the muscles — for example by measuring their *curvature*. Fig 6.1(c) shows the local curvature around the pixellated boundary of both eye-muscles. Darker pixels indicate a smaller radius of curvature and lighter pixels indicate a larger radius (more details are given in §6.3).

Measurements are usually taken from the images output from segmentation algorithms (chapter 4), which have possibly also been processed using morphological operators (chapter 5). Such will be the case for most of the examples considered in this chapter. However, in some applications, measurements can be obtained directly from an image without any preprocessing being required,

and in other cases it may be necessary to combine segmentation and greyscale information. Examples will be given in §6.1.1.

The three forms of image measurement in Fig 6.1 illustrate the coverage in the three sections which follow. In §6.1, ways will be considered in which size measurements can be obtained from an image. Then in §6.2, aspects of shape which are independent of size will be presented. Descriptions of the boundaries of objects will be considered in §6.3. Finally, in §6.4 the main points of the chapter will be summarised.

6.1 Measures of size

It is a straightforward matter to count the number of objects in an image using the connected-components algorithm in §4.2, provided that the segmentation has successfully associated one, and only one, component with each object. If this is not the case, then manual intervention may be necessary to complete the segmentation, as was illustrated for the muscle fibres in Fig 4.10(b). However, short-cuts can sometimes be taken. For example, if the mean size of objects is known, then the number of objects in an image can be estimated even when they are touching, through dividing the total area covered by all the objects by this average size. It is even possible to make allowance for objects overlapping each other provided that this process can be modelled, for instance by assuming that objects are positioned at random over the image and making use of the properties of these so-called **Boolean models** (Cressie, 1991, pp 753-759). For example, Jeulin (1993) estimated the size distribution of a powder.

It is also important to take account of any objects which *overlap the border* of an image. If we wish to estimate average areas or perimeters of objects, then the simplest solution is to ignore such objects, although this may introduce some downward bias because larger objects are more likely to intersect the border. However, if the aim is to estimate the number of objects per unit area, this will be overestimated if we include all objects in the count, but underestimated if we exclude those which overlap the border. An approximately unbiased estimate is obtained if objects only in contact with the bottom or right image borders are ignored.

It is always a sound practice, if possible, to **calibrate** the results produced by image analysis. This involves obtaining a few of the required measurements by some other means which are known to be accurate, and comparing the answers with those from image analysis. This is well worthwhile, even it is time consuming or expensive, because it secures much greater credibility for the results produced by image analysis. If there is a discrepancy between the two types of measurement, then in some cases an adjustment for bias can be applied to the image analysis results.

Fowler *et al* (1990) provide an excellent case study of calibration in MRI (magnetic resonance imaging). Accuracy of volume estimates was assessed by imaging bottles of copper sulphate at a range of positions in the imaging plane. Also, the test object which can be seen in Figs 1.7(a) and(b), and consisted of 80% water, was used as a standard against which to estimate water content of different tissues.

The two most common types of statistics used to describe an object's size are measurements of area and of distance. Measurements of area and related moments will be covered in §6.1.1. Many variants of distance, including perimeters and diameters, will be dealt with in §6.1.2.

6.1.1 Areas and moments

It is not always necessary to segment an image in order to estimate the area of a region. For example, in the Landsat image we may want to measure the area covered by oil-seed rape fields — these are the approximately-rectangular regions which appear as yellow and yellowish-green in Fig 2.10 (colour plate). Fig 6.2 shows a plot of pixel values in band 2 (green) against those in band 1 (blue). Yellow pixels have been identified manually in this scatterplot as those lying in the outlined region with a higher green than blue signal. They total 6301 pixels, which converts to an area of 5.7 km^2, because 1 pixel is equivalent to 900 m^2 on the ground. Of course, not all pixels will have been correctly classified by this criterion, but provided the number of oil-seed rape pixels which have been missed approximately cancels out the number of other pixels incorrectly classified as rape, the overall estimate will be reasonably accurate. This will depend on both misclassification rates and total areas involved, and may be a somewhat optimistic assumption. However, the estimate can be improved and the uncertainty quantified, if the true ground cover can be determined for part of the image (for example, by visiting the region and identifying the crop growing in each field). Then a **ratio estimate** can be constructed, making use of the satellite data as an auxiliary variate (see for example, Cochran, 1977, chapter 6).

Let us now consider a segmented image, and let the object which we are interested in measuring be denoted by all pixel locations (i, j) such that $(i, j) \in A$. We will use the set notation of chapter 5 (see Fig 5.2(a)). The **area** is the number of pixels in A, and can be expressed mathematically as:

$$\text{area} = \sum\sum_{(i,j)\in A} 1,$$

that is, count 'one' for every pixel in A.

In some applications we may be interested in the sum of pixel values within the region specified by A. This is given by

$$\sum\sum_{(i,j)\in A} f_{ij},$$

where f_{ij} denotes the greyscale value of pixel (i, j). For example, the cumulative sum of the pixel values within a spot on an electrophoresis gel (such as in Figs 1.9(a) and (b)), after subtraction of the average background pixel value, can be used to estimate the volume of that protein in the sample. However, to achieve precise results, image digitization should be performed using a scientific laser scanner rather than the cheaper desktop scanner we used to digitize these gels. In another application, Fowler *et al* (1990) used average MRI signals within selected regions, in images such as Figs 1.7(a) and (b), to estimate water content of different tissues.

We shall now consider how the definition of area can be generalized to encompass a class

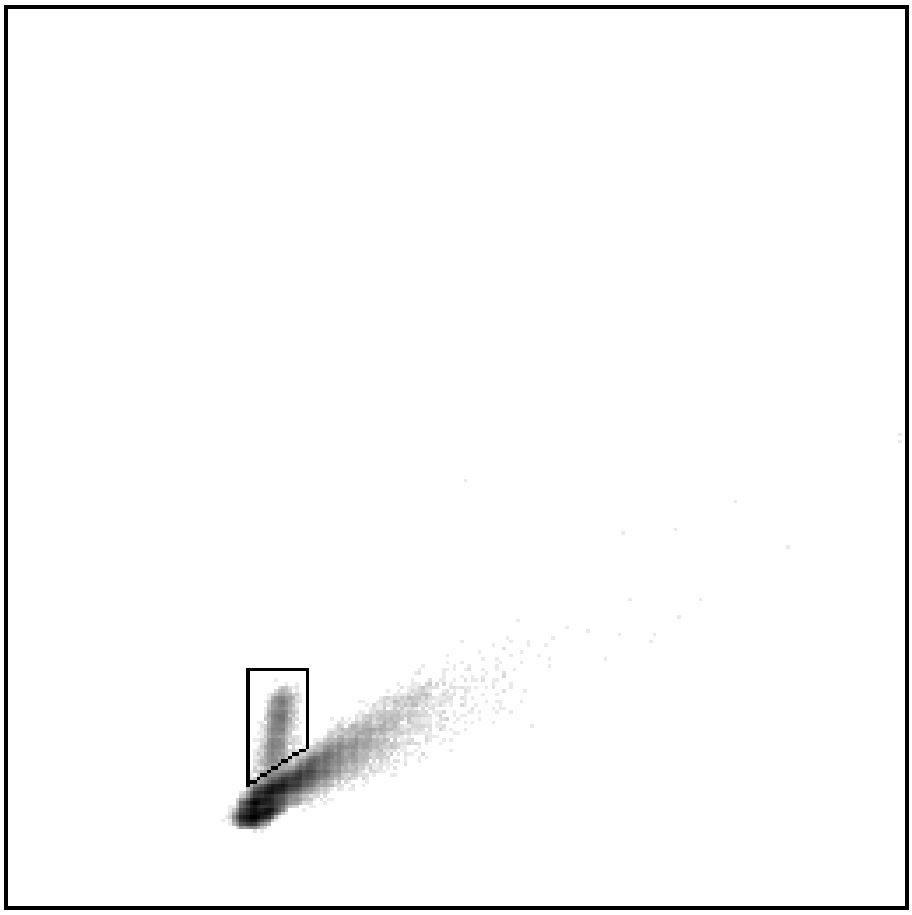

Figure 6.2: A scatter plot of Landsat band 2 (green) against band 1 (blue), with pixel frequency displayed on a log-scale, and the range of values enclosed by the quadrilateral corresponding to pixels in oil-seed rape fields.

of measurements, termed 'moments', which can be used to specify the location and spatial distribution of an object. Instead of counting 'one' for each pixel in the expression above, we count 'i raised to the power k, multiplied by j to the power l' if the pixel is in row i and column j. This is the definition of the **$(k,l)^{\text{th}}$-order moment**:

$$\mu_{kl} \quad = \quad \sum_{(i,j)\in A}\sum i^k j^l \qquad \text{for } k,l = 0,1,2,\ldots \quad .$$

If k and l are both zero, we obtain the zeroth-order moment, μ_{00}, which is the area, because $i^0 = j^0 = 1$.

First-order moments (i.e. $k + l = 1$) can be used to specify the location of an object. The **centre of gravity** or **centroid** of A is a measure of the object's location in the image. It has two components, denoting the row and column positions of the point of balance of the object if it were represented by a piece of cardboard of the same dimensions. The two numbers can be calculated as the mean row number, and the mean column number, of pixels in A. They will

not usually be integers. They are expressible in terms of zeroth- and first-order moments, by

$$\text{centroid} \;=\; \left(\frac{\mu_{10}}{\mu_{00}},\; \frac{\mu_{01}}{\mu_{00}} \right).$$

In applications where the positions of objects are important, the distribution of centroids can be studied using techniques of spatial point pattern analysis (Diggle, 1983).

Moments of higher-order (i.e. $k + l > 1$) are also mainly determined by an object's location. The statistics are more useful if this information (which we already know about from μ_{10} and μ_{01}) is removed. **Central moments** are defined as above, except that the mean row position is subtracted from i and the mean column position is subtracted from j:

$$\mu'_{kl} \;=\; \sum\sum_{(i,j)\in A} \left(i - \frac{\mu_{10}}{\mu_{00}} \right)^k \left(j - \frac{\mu_{01}}{\mu_{00}} \right)^l \qquad \text{for } k + l > 1 \; .$$

Central moments are **location invariant** — that is, two objects which are identical except for having different centroids will have identical values of μ'_{kl}, for all values of k and l. At least this is the case if the centroid changes by an integer number of pixels — for non-integer changes in location the pixel representation of an object may change, with consequent effects on the central moments. The moments can be expressed in terms of the earlier, non-central moments, μ_{kl}. In particular, the second-order central moments are

$$\mu'_{20} = \mu_{20} - \frac{\mu_{10}^2}{\mu_{00}},$$

$$\mu'_{02} = \mu_{02} - \frac{\mu_{01}^2}{\mu_{00}}$$

$$\text{and} \quad \mu'_{11} = \mu_{11} - \frac{\mu_{10}\mu_{01}}{\mu_{00}}.$$

The second-order moments measure how dispersed the pixels in an object are from their centroid: μ'_{20} measures the object's spread over rows, μ'_{02} measures its spread over columns, and μ'_{11} is a cross-product term representing spread in the direction in which both row and column indices increase. They are proportional to the variances and covariance of a bivariate random variable with a distribution which is uniform over the set A.

Central moments are not rotationally invariant — they will change if an object is rotated. If orientation is an important feature of an object, as it will be in some applications, then it is probably desirable for the moments to be sensitive to it. But in other cases orientation is irrelevant and moment statistics are more useful if they are invariant to rotation as well as to location. Again we must invoke the proviso that rotation will affect the pixel representation of an object, and therefore the invariance will hold only approximately. Hu (1962) derived **rotationally-invariant moments**, the first two of which are the following functions of the second-order central moments:

$$\mu'_{20} + \mu'_{02}$$

$$\text{and} \quad (\mu'_{20} - \mu'_{02})^2 + 4{\mu'_{11}}^2.$$

The first of these statistics is the **moment of inertia**, a measure of how dispersed, in any direction, the pixels in an object are from their centroid, whereas the second statistic measures whether this dispersion is isotropic or directional. Note that, Reiss (1991) showed that some of Hu's other results are in error.

There are alternative ways of combining the second-order moments in such a way that the resulting statistics are rotationally invariant. One such method is based on first specifying the direction in which the object has the maximum value for its second-order moment. This direction is:

$$\phi = \frac{1}{2}\tan^{-1}\left(\frac{2\mu'_{11}}{\mu'_{02} - \mu'_{20}}\right) \qquad \text{if} \quad \mu'_{02} > \mu'_{20},$$

and is otherwise this expression plus $\frac{\pi}{2}$, where ϕ is measured clockwise, with the horizontal direction taken as zero, and $\tan^{-1}$ produces output over the range $\frac{-\pi}{2}$ to $\frac{\pi}{2}$. Direction ϕ, which is the **major axis** or **direction of orientation** of the object, has second-order moment:

$$\lambda_1 = \mu'_{20}\sin^2\phi + \mu'_{02}\cos^2\phi + 2\mu'_{11}\sin\phi\cos\phi.$$

The direction perpendicular to ϕ, that is the **minor axis**, has the smallest second-order moment, of

$$\lambda_2 = \mu'_{20}\cos^2\phi + \mu'_{02}\sin^2\phi - 2\mu'_{11}\sin\phi\cos\phi.$$

These second-order moments, λ_1 and λ_2, are rotationally invariant, by definition. For a derivation, see Rosenfeld and Kak (1982, Vol 2, pp 288-290).

Table 6.1 shows several of these moment statistics for the segmented algal cells displayed in Fig 1.12(f). For example, the region representing the cell labelled 1: consists of 358 pixels, is centred at location (18,254), is orientated so that the major axis lies at an angle of 34° to the horizontal in a clockwise direction, and has major and minor second-order moments of 13000 and 8100. Note that values of ϕ for all thirteen cells lie in the range $-11°$ to 65°. This is because segmented cells from a Differential Interference Contrast microscopic image tends to produce elliptical shapes with a common orientation.

If moments of high enough order are recorded from an object, then they uniquely specify the object, which can be recovered precisely (Hu, 1962). However, this is of only theoretical interest because, in practice, such an approach would be numerically highly unstable. Teh and Chin (1988) review other families of moment statistics, such as Zernike and Legendre moments. It is also possible to define moments using boundary pixels alone, by an appropriate change to the definition of set A, and using greyscale pixel values (Rosenfeld and Kak, 1982, Vol 2, p 287); for example, the zeroth-order greyscale moment was introduced earlier in the section. Higher-order moments and their invariant forms can also be derived, though this is less commonly done. Teh and Chin (1986) and Mardia and Hainsworth (1989) concern themselves with the difference between moments evaluated on an integer lattice and those which would be produced if each pixel were regarded as a square of unit size. In particular, for second-order moments

Cell	μ_{00}	$\frac{\mu_{10}}{\mu_{00}}$	$\frac{\mu_{01}}{\mu_{00}}$	μ'_{20}	μ'_{02}	μ'_{11}	ϕ	λ_1	λ_2
1	358	18	254	9600	11500	2300	34	13000	8100
2	547	50	262	22100	26500	4200	31	29000	19600
3	294	98	196	7700	8100	3600	43	11500	4200
4	774	109	174	50600	45900	6500	55	55200	41300
5	566	191	431	24700	27000	4600	38	30600	21100
6	306	200	254	7400	9700	3900	37	12600	4500
7	243	217	266	3900	5800	−400	−11	5900	3900
8	312	243	348	7400	8000	0	0	8000	7400
9	277	261	78	6900	5600	800	65	7300	5200
10	418	285	452	12600	17200	4100	30	19600	10200
11	387	296	100	9500	16000	2800	21	17100	8400
12	261	358	338	5600	5400	800	49	6200	4700
13	440	410	126	13300	19200	4000	27	21200	11300

Table 6.1: Moment statistics for algal cells in Fig 1.12(f)

they suggest leaving μ'_{11} unchanged, but replacing μ'_{20} and μ'_{02} by

$$\mu'_{20} + \frac{\mu'_{00}}{12} \qquad \text{and} \qquad \mu'_{02} + \frac{\mu'_{00}}{12}.$$

However, these corrections do not make much difference, except for objects only a few pixels in size, and so we have not used them in our examples.

6.1.2 Distances, perimeters and diameters

The simplest of all distance measurements is that between two specified pixels (i, j) and (k, l). There are several ways in which distances can be defined on a lattice. The three most common ones are:

1. **Euclidean distance**,
$$\sqrt{(i-k)^2 + (j-l)^2}.$$

2. **Chessboard distance**,
$$\max(|i-k|, |j-l|),$$
the number of moves required by a chess king to travel from (i, j) to (k, l). It is also the length of the shortest chain of 8-connected pixels joining (i, j) and (k, l).

3. **City-block distance**,
$$|i-k| + |j-l|,$$
so-called because it is the distance we would have to travel between two locations in a city such as New York where the streets are laid out as a rectangular grid. It is also the number of pixels in the shortest chain of 4-connected pixels joining (i, j) and (k, l).

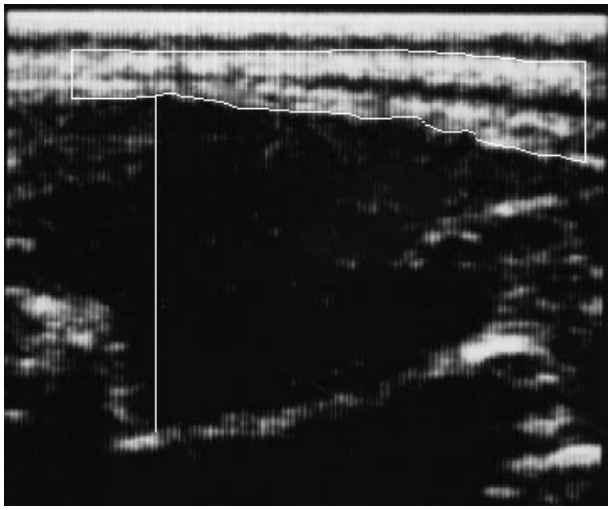

Figure 6.3: Ultrasound image, with hand-drawn lines encircling the fat region and showing the maximum depth of the eye-muscle.

The chessboard and city-block distances can be computed more quickly than the Euclidean distance, because they require only integer arithmetic. However, in the main we will use the Euclidean measure, because it accords with a geometric understanding of the objects to be quantified and is unaffected by the lattice orientation.

To illustrate a simple case of measuring distance, consider Fig 6.3. This shows the ultrasound image (Fig 1.7(d)) after manual interpretation. In particular, a vertical line has been drawn between the top and bottom of the eye-muscle at its maximum depth, using a computer mouse. This line has a length of 201 pixel units, which converts to an eye-muscle depth of 32 mm, because the vertical inter-pixel distance is 0.157 mm.

Another simple way in which distances are used in quantifying objects is that average distances, such as **average breadths**, may be produced. These can be obtained by dividing the area of an object by its length. To illustrate this, the subcutaneous fat layer of the sheep has been manually outlined in Fig 6.3. The left end of the region has been located above a feature in the backbone and the region is of fixed length, 307 pixels. We are interested in the average fat depth. This can be measured by dividing the area of the region by the number of columns it contains, in this case 11987/307, which is 39 pixel units or 6 mm.

Lengths of smooth **curved lines** are sometimes required, such as the total length of fungal hyphae in Fig 1.9(d). Before proceeding with the theory of measuring lengths of curved lines, let us think about the fungal hyphae in more detail.

We could measure the total hyphal length by dividing the area they cover by their average breadth, as we did above for the sheep fat depth. However, we do not know how wide the hyphae are, nor are they easy to measure accurately because they are only about two pixels wide. We can overcome this difficulty in a neat way by skeletonizing (§5.3) the thresholded

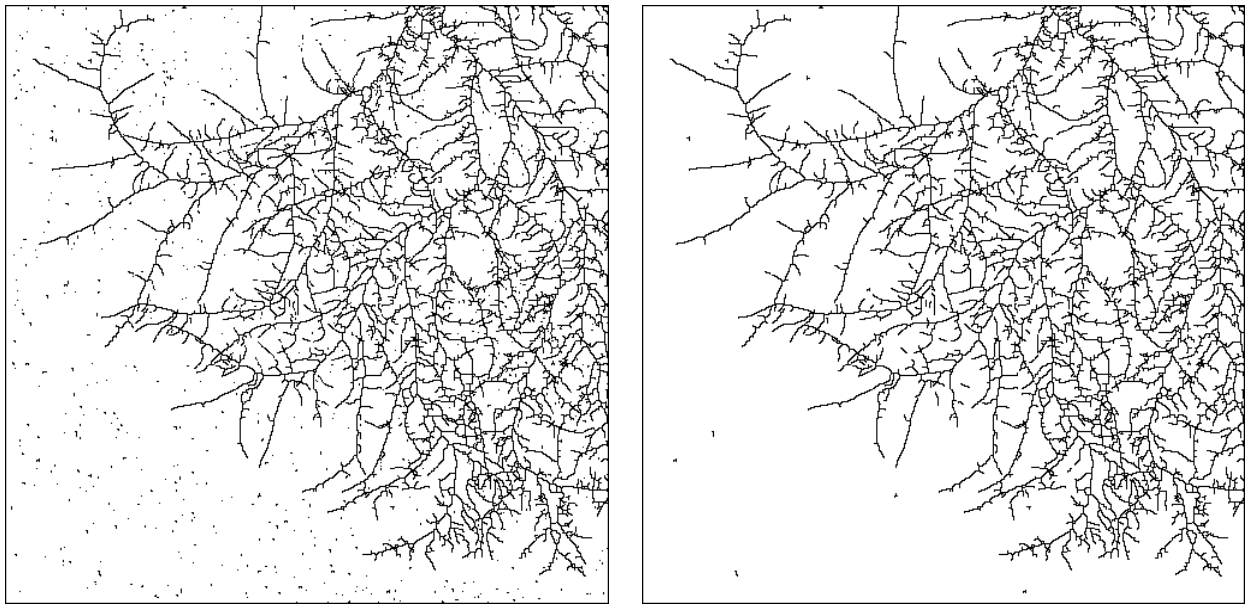

(a) (b)

Figure 6.4: Fungal image: **(a)** thresholded and skeletonized, **(b)** with small connected components removed.

image. Then we know that the breadth will be exactly one pixel unit. Fig 6.4(a) shows the skeleton of the thresholded image. It consists of 1121 distinct objects, with a total area of 25838 pixels: results which were obtained using the connected-components algorithm in §4.2. If all objects less than 5 pixels in area are assumed to be the products of noise rather than hyphae and so are discarded, Fig 6.4(b) is produced. The count is reduced to 231 components and 24013 pixels. We have one further difficulty to contend with, namely that in order to convert the count into an estimate of length it is necessary to consider the relationship between a line in continuous space and on a lattice.

We will consider line lengths in general, before returning to the fungal hyphae application. Figs 6.5(b) and (c) show two possible lattice representations of the straight-line section in Fig 6.5(a), depending on whether the pixels are 4-connected or 8-connected. We made use of connectedness in §4.2: in a 4-connected region all pixels have horizontally- or vertically-adjacent neighbours, whereas in an 8-connected region the neighbours may be diagonally adjacent. The line is 30.4 pixel units in length, although there are 37 pixels between the two end pixels in Fig 6.5(b) and 28 pixels in Fig 6.5(c). Therefore, in order to obtain good estimates of lengths in continuous space, we need to make some adjustment to pixel counts.

Consider any straight line of length d and angle θ. We need only consider θ in the range 0 to $\frac{\pi}{4}$, because for a line at any other orientation, we could if necessary swap or reverse row and column labels until θ did lie in this range. The line covers $d\cos\theta$ columns and $d\sin\theta$ rows. In the 4-connected line shown in Fig 6.5(b), $d\sin\theta$ columns have 2 pixels in them and the rest each have a single pixel, whereas in Fig 6.5(c) every column has one pixel in it. Therefore, the number of pixels are

$$d(\cos\theta + \sin\theta) \qquad \text{and} \qquad d\cos\theta,$$

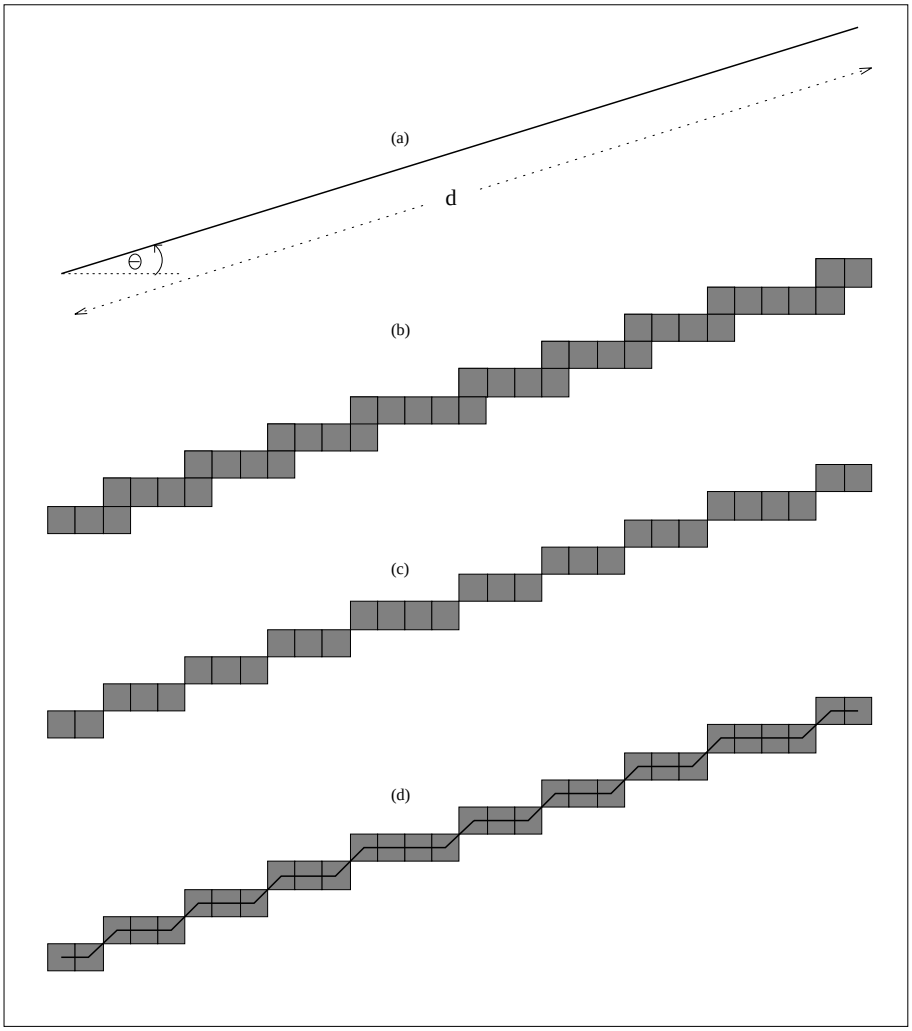

Figure 6.5: Diagram to show 3 ways in which lengths of straight lines can be measured in a pixel image: **(a)** in continuous space the line has length d and makes an angle θ with the horizontal, **(b)** the 4-connected pixels lying on the line, totalling N_4 are shown, **(c)** the 8-connected pixels lying on the line, totalling N_8 are shown, **(d)** a third way of measuring the line length is by measuring the distance between diagonally-adjacent pixels as $\sqrt{2}$, rather than as unity in (c).

for the 4- and 8-connected cases, repectively. If straight line segments at all orientations θ *occur equally often* in the curved line to be measured, then on average the number of 4-connected pixels per unit length can be obtained by integrating over all values of θ between 0 and $\frac{\pi}{4}$, i.e.

$$\frac{4}{\pi}\int_0^{\frac{\pi}{4}}(\cos\theta+\sin\theta)\,d\theta=\frac{4}{\pi}=1.273.$$

Similarly, the number of 8-connected pixels per unit length is given by

$$\frac{4}{\pi}\int_0^{\frac{\pi}{4}}\cos\theta\,d\theta=\frac{4}{\pi\sqrt{2}}=0.900.$$

Therefore, a smooth curved line, which can be approximated by straight-line segments of length 1000, for instance, will be represented on the image lattice by about 1273 pixels in a 4-connected

line, or 900 pixels in an 8-connected line.

Returning to the fungal image, we used a 4-connected skeleton. Also, line segment orientations were shown to be reasonably uniform in §5.4, i.e. there is no strong directional pattern in Fig 6.4(b). Therefore, using the above result, the total length of fungal hyphae can be estimated to be $24013/1.273 = 18860$ pixel units.

Perimeters of objects, that is boundary lengths, can also be measured by making use of the above theory. Let N_8 denote the number of pixels on the boundary of object A, which we count in the following way. Pixel (i, j) is on the boundary if $(i, j) \in A$, but one of its four horizontal or vertical neighbours is outside the object, that is,

$$(i+1, j) \notin A \text{ or } (i-1, j) \notin A \text{ or } (i, j+1) \notin A \text{ or } (i, j-1) \notin A.$$

Similarly, let N_4 denote the number of pixels $(i, j) \in A$, such that either one of the above four neighbouring pixels is not in A, or one of the four diagonal neighbours is outside the object, that is

$$(i+1, j+1) \notin A \text{ or } (i-1, j+1) \notin A \text{ or } (i-1, j+1) \notin A \text{ or } (i-1, j-1) \notin A.$$

The case N_4 gives the number of pixels in the 4-connected boundary of the object, whereas N_8 corresponds to the number of pixels in an 8-connected boundary. Therefore

$$\frac{N_4}{1.273} \quad \text{and} \quad \frac{N_8}{0.900}$$

are two unbiased estimators of the perimeter, provided that the assumption is satisfied that all orientations in the boundary occur equally often.

Note that :

- The above definition of the boundary of A will also include any internal boundaries if the object has *holes*. The **Euler number** of a set is defined to be the number of connected components it consists of, minus the number of holes it contains.
- The use of scaling factors is part of **stereology**, a field which has traditionally been concerned with inference about objects using information from lower-dimensional samples — such as estimating volumes of objects from the areas of intersection with randomly positioned cutting planes (see for example, Stoyan, Kendall and Mecke, 1987, chapter 11). In particular, the scaling factor of 1.273 ($\frac{4}{\pi}$) arises in two of the so-called 'six fundamental formulae' of classical stereology. However, the last ten years have seen a revolution in stereology, with the discovery of the *disector* (sic) and other three-dimensional sampling strategies (Stoyan, 1990).

Fig 6.5(d) shows a third way of estimating the length of a line, by measuring the distance between diagonally-adjacent pixels using the Euclidean measure of $\sqrt{2}$. The number of diagonal links is $(N_4 - N_8)$, and the remaining $\{N_8 - (N_4 - N_8)\}$ links in the 8-connected boundary are of one pixel unit in length. Therefore, the total length is

$$\{N_8 - (N_4 - N_8)\} + \sqrt{2}(N_4 - N_8) \quad = \quad (\sqrt{2} - 1)N_4 + (2 - \sqrt{2})N_8.$$

		Left	Right
Boundary statistics			
	N_4	160	158
	N_8	118	112
Perimeter estimators			
	$N_4/1.273$	125.7	124.1
	$N_8/0.900$	131.1	124.4
	$(0.414N_4 + 0.586N_8)/1.055$	128.3	124.2
	convex hull perimeter	128.6	123.3

Table 6.2: Boundary statistics and perimeter estimates for eye-muscles in Fig 6.1

Repeating the previous calculations for this new expression, the expected number of pixels per unit length is

$$\frac{4}{\pi}\int_0^{\frac{\pi}{4}} \{\cos\theta + (\sqrt{2}-1)\sin\theta\}\, d\theta \;=\; \frac{8}{\pi}(\sqrt{2}-1) \;=\; 1.055,$$

which again can be used as a correction factor to estimate total length.

All three perimeter estimators are unbiased, provided that the condition is satisfied that all line orientations occur equally often. However, they are not all equally precise. The final estimator has the lowest variance and should be used in preference to the other two. These, and more complicated methods for estimating perimeters yet more precisely, are considered by Dorst and Smeulders (1987) and Koplowitz and Bruckstein (1989).

Table 6.2 shows N_4 and N_8 for the sheep eye-muscles shown in Fig 6.1, together with the results obtained using the three perimeter estimators. (The convex hull perimeter will be dealt with later in the section.) Agreement is very close for the right muscle. For the left muscle, differences of about 4% have arisen because the boundary is rougher than the right one, and so the assumption that the line is smooth is not wholely appropriate.

The statistics we have considered are all measures of the **internal perimeters** of objects. If instead the length of the boundary of the background pixels surrounding the object was measured, this would yield the **external perimeter**, which would be greater than the internal one. The difference arises because in one case we are measuring the distances between the centres of pixels on the object side of the boundary, whereas in the other case we are measuring the distances between the centres of pixels on the other side of the boundary. A perimeter intermediate between these two can be obtained by regarding each pixel as a square of unit size, rather than as a point, and measuring the perimeter so that it goes round the boundary sides of each pixel. A short-cut to approximately the same result involves adding π to the internal perimeter. This is the extent to which the perimeter increases if an object is expanded by half the inter-pixel distance in all directions. The effect will be noticeable only for objects

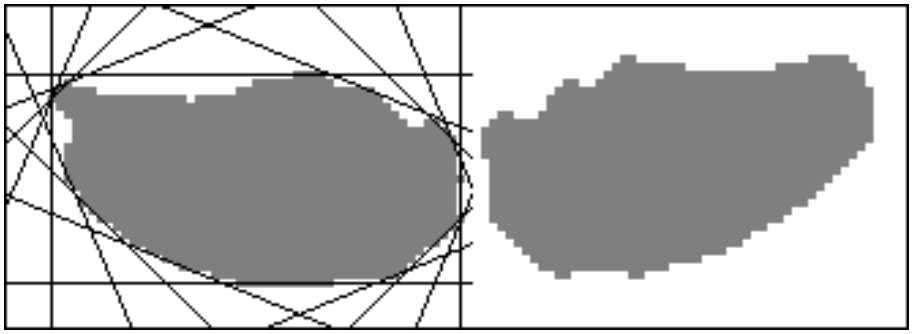

Figure 6.6: Eye-muscles from X-ray CT image, together with eight pairs of parallel tangents around the left eye-muscle at orientations which are multiples of 22.5°. They are used to estimate Feret diameters and approximate the convex hull.

a few pixels in size, and so can usually be safely ignored.

The final distances we shall consider in this section are **Feret diameters**, which are defined to be the distances between parallel tangents touching opposite sides of an object. For example, Fig 6.6 shows eight pairs of tangents to the left eye-muscle at orientations which are multiples of 22.5°. At orientation θ, the Feret diameter is

$$\max_{(i,j)\in A}(i\sin\theta + j\cos\theta) \quad - \quad \min_{(i,j)\in A}(i\sin\theta + j\cos\theta).$$

although it is more efficiently computed using boundary pixels alone. The diameters are also referred to as *caliper diameters* because of the analogy with measuring an object's size using calipers, and as *lengths of projections* because they measure the extent of A when it is projected onto a line at orientation $\theta + \frac{\pi}{2}$.

Certain Feret diameters are of special interest, although the terminology varies between authors.

- The **width** of an object is the range of columns it covers, i.e.

  $$\max_{(i,j)\in A}(j) \quad - \quad \min_{(i,j)\in A}(j).$$

  It is the Feret diameter at angle $\theta = 0°$.

- Similarly, the **height** of an object is the range of rows it covers, i.e.

  $$\max_{(i,j)\in A}(i) \quad - \quad \min_{(i,j)\in A}(i).$$

  It is the Feret diameter at angle $\theta = 90°$. Sometimes, the product of width and height is computed — this is the area of the enclosing rectangle.

- The maximum Feret diameter is one definition of the **length** of an object, which may alternatively be specified as

  $$\max_{(i,j),(k,l)\in A}\sqrt{(i-k)^2 + (j-l)^2}.$$

The corresponding value of θ is an alternative definition of orientation to that considered in §6.1.1.

- The **breadth** of an object does not have a unique definition. Some people define it to be the minimum Feret diameter, whereas others take it to be the Feret diameter perpendicular to the length. To illustrate, Fig 6.1(b), which has already been discussed, shows the lengths for the two eye-muscles, and the perpendicular breadths.

Length can be estimated reasonably precisely by taking the maximum of a few Feret diameters. It will be slightly underestimated, but not by much. If breadth is defined to be the minimum diameter, and is estimated in a similar fashion, as the minimum of a set of diameters, then the overestimation is more severe. Therefore more Feret diameters are needed for precise measurement of the breadth. Houle and Toussaint (1988) give an efficient algorithm.

The **convex hull** of an object is defined to be the smallest convex shape which contains the object — where a convex shape is one in which, if any two points within it are joined by a straight line, then all the points along the line are also within the shape. The convex hull can be obtained as the region enclosed by all the tangent planes to the object, as can be seen for the left eye-muscle in Fig 6.6. If the scale at which measurements are made is varied, by changing the magnification of a microscope, for example, then the perimeter of the convex hull of the object will be found to vary less than the perimeter of the original object. In an extreme case, in which an object has a **fractal** boundary (Mandelbrot, 1982), its estimated perimeter grows infinitely large as the resolution increases. Therefore, the perimeter of the convex hull can provide a more stable measure of an object's boundary length. The perimeter of the convex region defined by N equally-spaced Feret diameters provides an approximation to the perimeter of the convex hull. This perimeter is expressible as

$$2N\bar{D}\tan\left(\frac{\pi}{2N}\right),$$

where $\bar{D}$ is the average of N Feret diameters at orientations which are integer multiples of π/N. Alternatively, the convex perimeter of an object can be estimated using stereological sampling methods. The convex perimeters of the eye-muscles are given in Table 6.2.

Volumes of objects can be estimated, provided that something is known about their three-dimensional shape, using measurements obtained from two-dimensional projections onto the imaging plane. For example:

- If an object is spherical, then its volume (V) can be estimated from the area (A) of its projection, as

$$V = \frac{4A^{\frac{3}{2}}}{3\sqrt{\pi}}.$$

The volume could alternatively be estimated from the perimeter or diameter of the projected circular shape, but this is not recommended because these statistics are more variable than the area. In fact, the perimeters and diameters of circular regions are more precisely estimated as functions of the area than directly.

- If an object is a prolate ellipsoid, i.e. the two minor axes are of equal length, and the projection is along one of the minor axes, then the volume (V) can be estimated from the length (L) and breadth (B) of the projected ellipse, as

$$V = \frac{\pi L B^2}{6}.$$

6.2 Measures of shape

Shape information is what remains once

- location,
- orientation
- and size

features of an object have been dealt with. Therefore two objects which are the same, except that they are in different positions in the image, orientated at different angles, and that one is bigger than the other, will have the same shape. In the computer vision literature, the term **pose** is often used to refer to location, orientation and size. (Of course, invariance will hold only approximately in the case of a lattice representation of an object.)

We will begin our consideration of shape by returning to the very first example presented in this book, that of the turbinate image (§1.1.1). The objective of image analysis in this application was to estimate a *morphometric index*, which was defined as the ratio of air space area in the cross-section of the nasal cavities to air space area plus turbinate bone area. The index is pose invariant and so is a measure of shape. Fig 6.7 shows the turbinate image after the application of a morphological closing to reduce noise levels and improve connectivity, as shown earlier in Fig 5.6(b). The turbinate bone in each nasal cavity has been isolated from the remaining bone areas using a mouse to control a screen cursor. Numbers showing the areas of the largest black and white regions have been included on the figure. From these areas, the morphometric index of the left nasal cavity can be estimated as

$$\frac{(27100 + 1900 + 2100)}{(27100 + 1900 + 2100 + 11000)} = 0.74.$$

Similarly, for the right nasal cavity the index is

$$\frac{(26000 + 2500 + 800)}{(26000 + 2500 + 800 + 10800)} = 0.73.$$

These indices are typical of a pig which is free of the disease atrophic rhinitis. Note, that these statistics are dimensionless, as are all measures of shape because they are all size invariant.

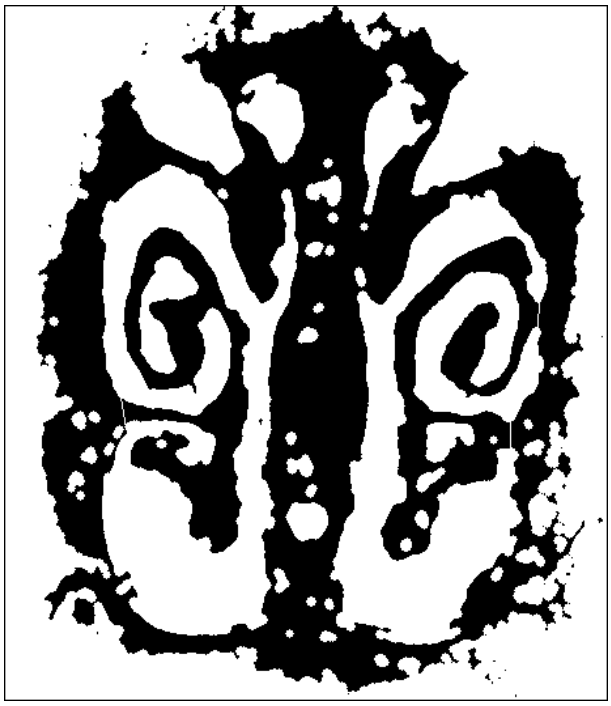

Figure 6.7: Turbinate image after morphological filtering and manual separation of the turbinate bone in each nasal cavity from the remaining bone areas, together with numbers showing the areas of the largest regions.

Returning to generalities, there are a very large number of possible ways of describing shape. However, a few, based on the statistics in §6.1, are widely used. They will be presented here, except for those based on boundaries, consideration of which will be deferred to §6.3.

Probably the most commonly used shape statistic is a measure of **compactness**, which is defined to be the ratio of the area of an object to the area of a circle with the same perimeter. A circle is used as the basic shape with which to make comparisons because it is the object with the most compact shape (although Rosenfeld (1974) shows that for objects which are approximated by pixels, an octagon minimizes the measure). The statistic is:

$$\text{compactness} = 4\pi \frac{\text{area}}{(\text{perimeter})^2} \quad .$$

Note that this is a shape statistic because it is pose invariant: changing an object's location or orientation or doubling its size, say, will leave the measure unchanged; except for the proviso (already commented on) that the estimated perimeter of an object can be sensitive to the resolution at which it is measured. The measure takes its largest value of 1 for a circle. Any departures in the object shape from a circular disc, such as an elliptical shape or a border which is irregular rather than smooth, will decrease the measure.

The responsiveness of the measure of compactness to any departure from a circular disc can sometimes be a disadvantage. In some situations it is useful to have measures which are sensitive only to departures of a certain type from circularity. Such statistics can be devised. For example, a measure of **convexity** can be obtained by forming the ratio of the perimeter of an object's convex hull to the perimeter of the object itself, that is:

$$\text{convexity} = \frac{\text{convex perimeter}}{\text{perimeter}} \quad .$$

This will take the value of 1 for a convex object, and be less than 1 if the object is not convex, such as one having an irregular border. Also, a measure of **roundness**, excluding these local irregularities, can be obtained as the ratio of the area of an object to the area of a circle with the same convex perimeter, that is:

$$\text{roundness} = 4\pi \frac{\text{area}}{(\text{convex perimeter})^2} \quad .$$

Again, as with the measure of compactness, this statistic equals 1 for a circular object and is less than 1 for an object which departs from circularity, except that it is relatively insensitive to irregular borders. The effects on the measure are more complex if an object is non-convex in an overall sense, for example a horseshoe shape.

Another statistic often used to describe shape is a measure of **elongation**. This can be defined in many ways, one of which is obtained by taking the ratio of an object's length to its breadth:

$$\text{elongation} = \frac{\text{length}}{\text{breadth}} \quad .$$

This statistic was derived for the sheep's eye-muscles at the beginning of the chapter. Another measure of elongation is given by the ratio of the second-order moments of the object along its major and minor axes:

$$\frac{\lambda_1}{\lambda_2} \quad .$$

The results are often very similar to the previous measure of elongation.

Another use of moments as measures of shape is that the central moments, μ'_{kl}, can be made scale invariant by dividing them by $\mu_{00}^{\frac{1}{2}(k+l+2)}$. This converts rotationally-invariant moments into measures of shape. Thus, the first **scale- and rotationally-invariant moment** is

$$\frac{\mu'_{20} + \mu'_{02}}{\mu_{00}^2} \quad .$$

This is another measure of compactness. It measures how dispersed the pixels in an object are from their centroid, in comparison with the most compact arrangement of the pixels. It takes its smallest value of $\frac{1}{2\pi} = 0.159$ for a circular object.

Table 6.3 gives some of these shape statistics, and the size statistics from which they were derived, for the thirteen algal cells in Fig 1.12(f). Measures involving the convex perimeter have been omitted because all the cells are nearly convex. For these cells, the measure of elongation

Cell	Area	Perimeter	Length	Breadth	Compact-ness	Elong-ation	$\frac{\mu'_{20}+\mu'_{02}}{\mu^2_{00}}$
1	358	68	23	19	0.96	1.23	0.165
2	547	84	28	24	0.97	1.20	0.163
3	294	65	24	15	0.88	1.62	0.182
4	774	100	34	29	0.97	1.15	0.161
5	566	86	29	25	0.96	1.16	0.161
6	306	66	24	14	0.87	1.65	0.183
7	243	57	19	16	0.94	1.21	0.164
8	312	62	20	19	1.01	1.03	0.159
9	277	60	20	17	0.97	1.17	0.163
10	418	78	26	21	0.87	1.28	0.171
11	387	73	26	19	0.92	1.38	0.170
12	261	59	19	18	0.95	1.07	0.161
13	440	77	27	20	0.93	1.31	0.168

Table 6.3: Size and shape statistics for the algal cells in Fig 1.12(f)

and $(\mu'_{20}+\mu'_{02})/\mu^2_{00}$ are highly correlated. Fig 6.8 shows elongation plotted against compactness, with each point represented by that cell's outline. Although all the cells have very similar shapes, some differences show up in the figure. For example, cell 8 is the one closest to being circular, and it has the lowest measure of elongation and the highest measure of compactness. Cells 3, 6 and 10 are the least compact because their perimeters are proportionally longer than those of the other cells. In the case of cells 3 and 6, this is because they are elongated, and they are therefore located in the top-left corner of the figure. The segmented version of cell 10 has an indentation (although probably the cell itself does not) and this is the reason why it is less compact than the others. Because the cell is no more elongated than any of the others, it is located towards the centre-left in the figure.

The description of shape is an open-ended task, because there are potentially so many aspects to an object even after location, orientation and size effects have been removed. Other approaches include the use of **landmarks** (identifiable points on objects) (Goodall, 1991) and warpings such as **thin-plate splines** and other **morphometric methods** (Bookstein, 1991), which consider image plane distortions needed to move landmarks to designated locations. Further methods are discussed in the reviews of shape analysis by Pavlidis (1978, 1980) and Marshall (1989), such as the use of **skeletons** (§5.3) and of distributions of **chord-length**; chords being the lines joining any pair of points on the boundary of an object.

Allometry is a subject related to shape. It is often found that the sizes of different parts of organisms are linearly related on a logarithmic scale. For example, the lengths (L) and breadths (B) of members of a population may, on average, obey

$$\log L = \alpha + \beta \log B$$

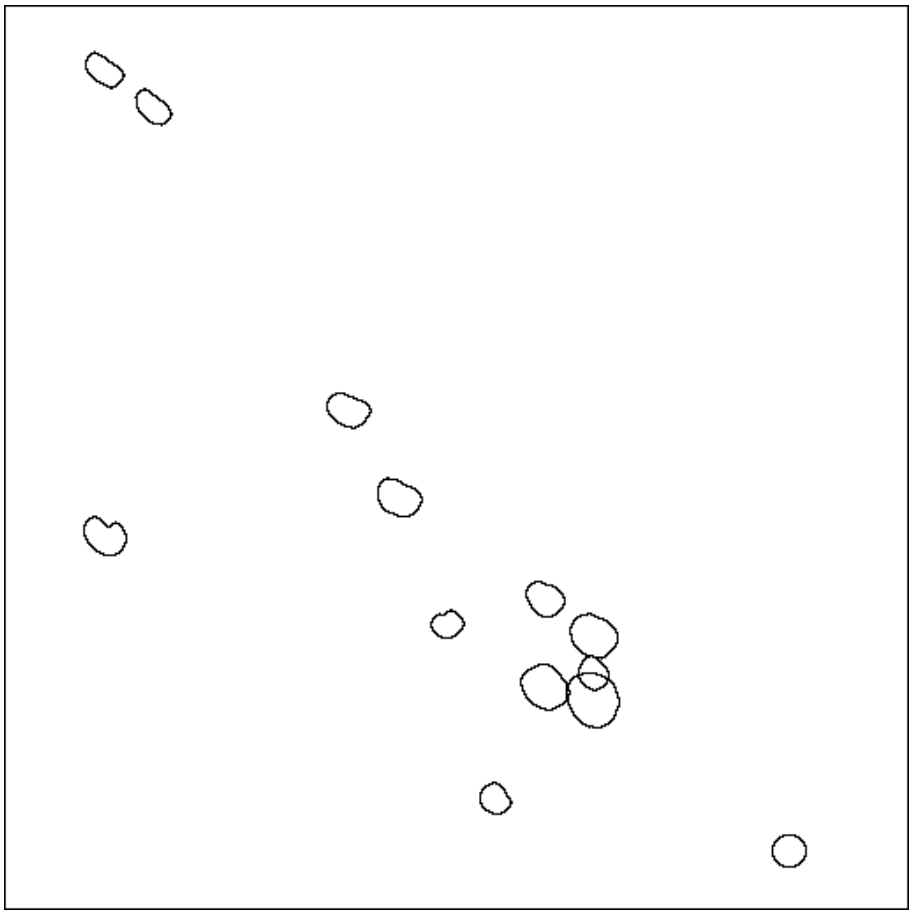

Figure 6.8: A scatter plot of a measure of elongation against a measure of compactness for the segmented algal cells: each point is represented by that cell's outline.

for certain values of the constants α and β (see, for example, Causton and Venus, 1981, chapter 6). This relationship can be re-expressed as

$$L = e^{\alpha} B^{\beta}.$$

Therefore, L/B^{β} is a size-invariant constant for the population, which it may be used to characterize in preference to the shape statistic L/B. To illustrate, Glasbey, McRae and Fleming (1988) found that the log-lengths and log-breadths of potato tubers were normally distributed, with standard deviations of 0.0766 and 0.0646, respectively. The ratio of these standard deviations gives an estimate of β of 1.19. Therefore, the allometric constant ratio is $L/B^{1.19}$, which is consistent with the observation that smaller potatoes tend to be more spherical than larger ones.

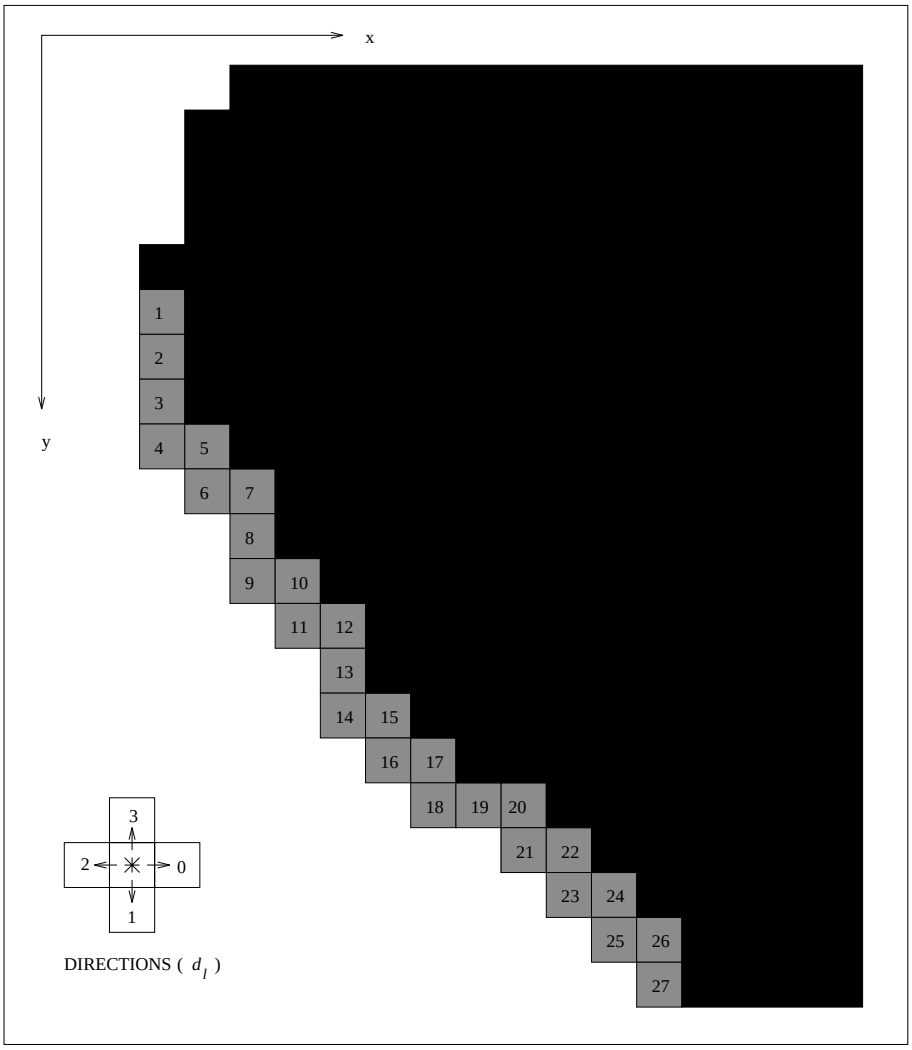

Figure 6.9: The nose of the fish image, after segmentation: a diagram to illustrate the extraction of a boundary chain-code. The inset shows how directions (d_l) used in the chain code are defined.

6.3 Boundary statistics

The first step in studying the boundary of an object is to extract the boundary pixels from the image. We describe an algorithm, first in words and then mathematically, for generating an ordered list of pixels round the 4-connected boundary.

We use d to denote the 4 possible directions we can take moving between pixels. In keeping with orientations used elsewhere in this book, direction 0 is defined to be across a row, so that from pixel location (i, j) we would move to $(i, j+1)$. Proceeding clockwise, directions 1, 2 and 3 will be used to denote changes from (i, j) to $(i+1, j)$, $(i, j-1)$ and $(i-1, j)$ respectively. This code is illustrated in the inset of Fig 6.9. We will also use the notation $\Delta(d)$ to denote the change in pixel label produced by direction d, so that $\Delta(0) = (0, +1)$, etc.

We will track round the boundary from pixel to neighbouring pixel in an *anti-clockwise* direction. This choice of direction is arbitrary, but is the common convention. To start, we need 2 points on the boundary, (y_1, x_1) and (y_2, x_2), where (y_2, x_2) is the next pixel location round the boundary in an anti-clockwise direction after (y_1, x_1). Set d_1 to the direction from (y_1, x_1) to (y_2, x_2), and take the initial direction of search (d_2) for (y_3, x_3) to be $(d_1 + 1)$, except that if $d_1 = 3$ we choose $d_2 = 0$.

The algorithm then proceeds as follows. If $(y_2, x_2) + \Delta(d_2)$ is a pixel location in the object, then (y_3, x_3) is set to this pixel. Otherwise d_2 is decreased by 1, except that when it reaches -1 it is reset to 3, and $(y_2, x_2) + \Delta(d_2)$ is examined to see if it is an object pixel. After trying at most four directions, point (y_3, x_3) will have been found, in a direction we preserve as d_2. The search for (y_4, x_4) then begins in direction $(d_2 + 1)$. Fig 6.9 shows the start of the boundary round the fish (Fig 1.9(e)) starting from the nose, and Table 6.4 shows the stored locations. (The fish image was segmented using a semi-automatic method involving thresholding and the use of a computer mouse to complete the segmentation, rather than by making use of a second, back-illuminated image discussed in §1.2.)

The boundary tracking is complete once we return to the starting point, except for the important proviso that if the object is only one pixel wide at (y_1, x_1), then we will return to this point more than once before having tracked round the whole boundary. To avoid this case, it is simply necessary to check that after reaching (y_1, x_1) we move to (y_2, x_2), and then we know that the algorithm has finished.

The algorithm is formalised as follows:

1. Initialise by setting (y_1, x_1) to a boundary point on the object (set A) and d_1 to an initial direction such that $(y_2, x_2) = (y_1, x_1) + \Delta(d_1)$ is the next pixel location round the boundary in an anti-clockwise direction. Also set $l = 2$.

2. Set the search direction $d_l = d_{l-1} + 1 \pmod 4$, i.e. an integer between 0 and 3, obtained as the remainder after dividing $(d_{l-1} + 1)$ by 4.

3. $(y_{l+1}, x_{l+1}) = (y_l, x_l) + \Delta(d_l)$,
   If $(y_{l+1}, x_{l+1}) \in A$ then go to step (4),
   otherwise $d_l \rightarrow d_l - 1 \pmod 4$, and repeat step (3),

4. If $(y_l, x_l) = (y_1, x_1)$ and $(y_{l+1}, x_{l+1}) = (y_2, x_2)$ then set $N = l - 1$ and the algorithm is completed,
   otherwise $l \rightarrow l + 1$ and go to step (2).

The algorithm can be modified to track an 8-connected boundary, by redefining d and Δ so that 8 directions are considered in anti-clockwise order, including the diagonal neighbours.

The **chain code** (Freeman, 1974), consisting of the starting location and a list of directions $d_1, d_2, \ldots, d_N$, provides a compact representation of all the information in the boundary. It can

l	y_l	x_l	d_l
1	89	30	1
2	90	30	1
3	91	30	1
4	92	30	0
5	92	31	1
6	93	31	0
7	93	32	1
8	94	32	1
9	95	32	0
10	95	33	1
11	96	33	0
12	96	34	1
13	97	34	1
14	98	34	0
15	98	35	1
16	99	35	0
17	99	36	1
18	100	36	0
19	100	37	1
20	100	38	1
21	101	38	0
22	101	39	1
23	102	39	0
24	102	40	1
25	103	40	0
26	103	41	1
27	104	41	1
.	.	.	.
.	.	.	.
.	.	.	.

Table 6.4: First few boundary points and chain code for segmented fish image

also be used directly to derive the summary statistics of §6.1. For example, in simple cases N is equal to the statistic N_4 (§6.1.2) and can be used to estimate the perimeter. The area and moments of an object can also be obtained from its chain code. But, if an object contains *holes*, the two sets of measures will be different. For example, the perimeter defined in §6.1.2 includes the boundary lengths of holes in an object, whereas N is a measure only of the length of a single connected boundary.

The directions (d_l) are estimates of the slope of the boundary, but because they are all multiples of 90° they are very imprecise. Smoothed estimates may be obtained by considering more distant boundary points. The **k-slope** of the boundary at (y_l, x_l) can be estimated from the slope of the line joining $(y_{l-k/2}, x_{l-k/2})$ and $(y_{l+k/2}, x_{l+k/2})$, for some small, even value of k. This

works out as an angle of

$$\tan^{-1}\left(\frac{y_{l+k/2}-y_{l-k/2}}{x_{l+k/2}-x_{l-k/2}}\right),$$

measured in a clockwise direction, with a horizontal slope taken to be the zero. Here, as in §3.2.1, '$\tan^{-1}$' is assumed to produce output over a range of angles of 2π radians.

In a similar fashion, the **k-curvature** of the boundary at point (y_l, x_l) can be estimated from the change in the k-slope:

$$\left\{\tan^{-1}\left(\frac{y_{l+k}-y_l}{x_{l+k}-x_l}\right) \quad - \quad \tan^{-1}\left(\frac{y_l-y_{l-k}}{x_l-x_{l-k}}\right)\right\} \qquad (\text{mod } 2\pi).$$

Fig 6.1(c), which has already been discussed, shows the curvature of the X-ray eye-muscle boundaries, evaluated with $k = 10$.

Use of **Fourier descriptors** is another approach to describing the boundary of an object. Granlund (1972) proposed approximating $y_1, \ldots, y_N$ by a sum of K ($\leq N/2$) sine and cosine terms:

$$\hat{y}_l = \bar{y} + \sum_{k=1}^{K}\left\{a_k \cos\left(\frac{2\pi kl}{N}\right) + b_k \sin\left(\frac{2\pi kl}{N}\right)\right\} \quad \text{for } l = 1, \ldots, N\ ,$$

$$\text{where} \quad a_k = \frac{2}{N}\sum_{l=1}^{N} y_l \cos\left(\frac{2\pi kl}{N}\right)$$

$$\text{and} \quad b_k = \frac{2}{N}\sum_{l=1}^{N} y_l \sin\left(\frac{2\pi kl}{N}\right) \quad \text{for } k = 1, \ldots, K\ .$$

And similarly for x, but with different Fourier coefficients. If K is equal to $N/2$ then the curve passes through all the boundary pixels, whereas if K is less than $N/2$ the curve provides a smooth approximation.

The Fourier coefficients (a's and b's) can be obtained by regressing y on the sine and cosine terms. The sum of squares of the y's about their mean ($\bar{y}$) can be represented as a sum of squares of the Fourier coefficients plus a sum of squares of the differences between the y's and their Fourier approximation:

$$\sum_{l=1}^{N}(y_l-\bar{y})^2 \quad = \quad \frac{N}{2}\sum_{k=1}^{L}(a_k^2+b_k^2) \quad + \quad \sum_{l=1}^{N}(y_l-\hat{y}_l)^2\ .$$

Alternatively, Fast Fourier Transform algorithms can be used to compute $a_1, \ldots, a_{N/2}$ and $b_1, \ldots, b_{N/2}$, provided that N is either a power of 2 or a product of small prime factors. If necessary, extra boundary points can be introduced by duplicating, or interpolating between, pixels to produce such a value for N. (See §3.2 for more details about the Fourier transform.)

For the segmented fish, the 4-connected boundary has length $N = 1416$. The Fourier approximations using one coefficient only, are

$$\hat{y}_l = \ 99.4 - \quad 2.6\cos\left(\frac{2\pi l}{N}\right) + 65.3\sin\left(\frac{2\pi l}{N}\right)$$

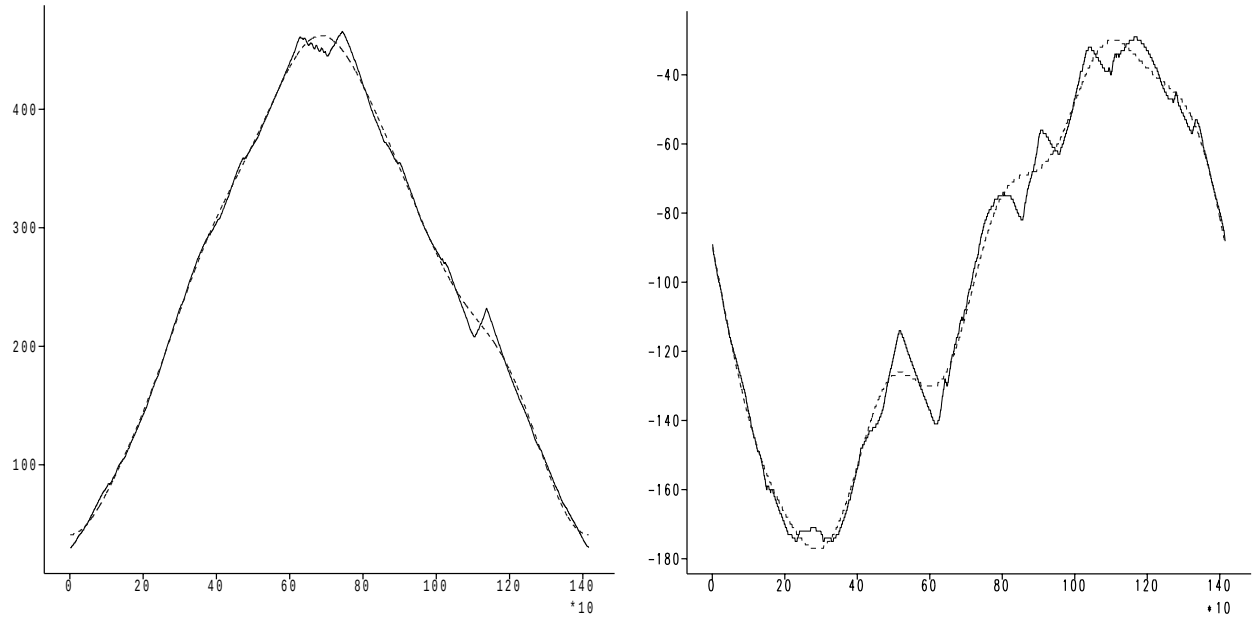

(a) (b)

Figure 6.10: Boundary points of segmented fish image, and their approximation using $K = 5$ Fourier coefficients, plotted against pixel count (l) around the boundary, starting from the nose: **(a)** column coordinate, x_l, **(b)** row coordinate, y_l.

$$\text{and} \quad \hat{x}_l = 256.4 - 188.5 \cos\left(\frac{2\pi l}{N}\right) + 8.5 \sin\left(\frac{2\pi l}{N}\right) \quad \text{for } l = 1, \ldots, N \ .$$

From a total sum of squares of the y's of 3,292,000, the residual sum of squares is reduced to

$$3,292,000 - \frac{1416}{2}(2.6^2 + 65.3^2) = 268,000 \ .$$

Therefore 92% of the variation in the y's has been accounted for. Fig 6.10 shows the fitted curves when $K = 5$, for which 99.2% of the y variability and 99.8% of the x variability has been explained. Fig 6.11 shows the resulting boundaries when $K = 1$, 5 and 20. When $K = 1$ the result is an ellipse. Increasing the number of terms improves the fit, but at the price of using more parameters. For $K = 20$, 99.97% and 99.99% of the variation in y and x, respectively, have been accounted for.

Crimmins (1982) shows how the Fourier coefficients can be transformed to provide measures of shape, i.e. statistics which are invariant to changes in an object's pose. There are other ways of using Fourier descriptors to approximate object boundaries. For instance, Zahn and Roskies (1972) proposed a set of Fourier descriptors which are based on *local boundary slope*, rather than directly on the (y, x) coordinates. Rohlf and Archie (1984) and Mou and Stoermer (1992) compared alternative approaches, and applied Zahn and Roskies' method to describe the outlines respectively of mosquito wings and diatoms. It is also possible to use **polygons** (Pavlidis, 1977), **splines** (Hill and Taylor, 1992) and **conic sections** (Bookstein, 1978) to

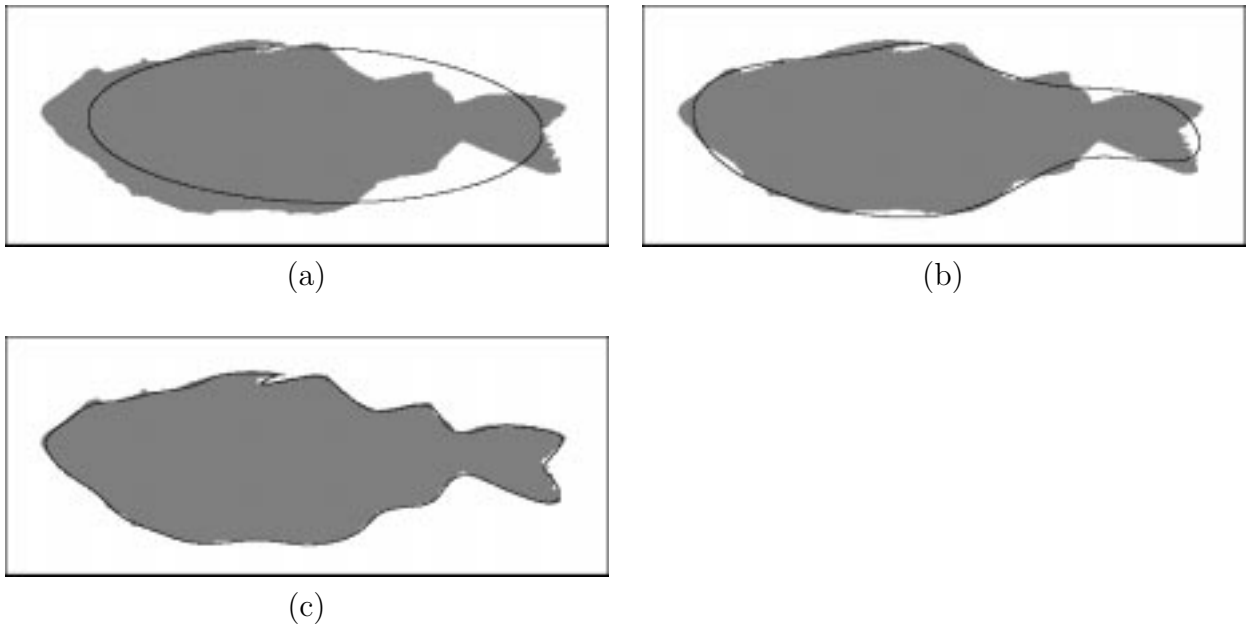

(a) (b)

(c)

Figure 6.11: Segmented fish image, together with approximations to the boundary based on Fourier descriptors using different numbers of coefficients (K): **(a)** $K = 1$, **(b)** $K = 5$, **(c)** $K = 20$.

describe boundaries. Marshall (1989) reviews the use of stochastic descriptions of boundaries based on **circular autoregressive models**.

6.4 Summary

The key points of this chapter are:

- Measurements are usually taken from the images output from segmentation algorithms (chapter 4), which have possibly also been processed using morphological operators (chapter 5). In some applications, measurements can be obtained directly from the original image.
- Three general categories of measurement are:
  - size,
  - shape, independent of size,
  - boundary statistics
- Two types of size statistics were considered:
  - those based on moments, such as:
    - area,

    - centre of gravity or centroid,
    - second-order moments, such as the moment of inertia which is invariant to changes in location and orientation.
  - those based on distances:
    - average breadth,
    - length of a curved line, such as a boundary,
    - Feret diameters, and convex-hull perimeter.

- Shape statistics which were considered included measures of:
  - compactness,
  - convexity,
  - roundness,
  - elongation,

  and scale- and rotationally-invariant moments.

- Boundary descriptors considered were:
  - ordered set of boundary pixels,
  - chain code,
  - local measures of slope and curvature,
  - Fourier descriptors.