Chapter 5

MATHEMATICAL MORPHOLOGY

Mathematical Morphology is a theory which provides a number of useful tools for image analysis. It is seen by some as a self-contained approach to handling images and by others (including the authors) as complementary to the other methods presented in this book. Fig 5.1 shows an example of an operation based on mathematical morphology (hereafter referred to simply as morphology). Fig 5.1(a) shows a subset of the soil image thresholded at a pixel value of 20 (Fig 4.2(d)). Fig 5.1(b) is the result of a morphological operation termed a **closing**, which is described in §5.1. At first it may seem to be just another filter which has smoothed the image, and in fact the closing operation is often used for this purpose. However, it does have different properties from the sort of filters considered in chapter 3. For example, although many black pixels have been changed to white to smooth the image, none have been changed from white to black.

Morphology is an approach to image analysis which is based on the assumption that an image consists of structures which may be handled by set theory. This is unlike most of the rest of

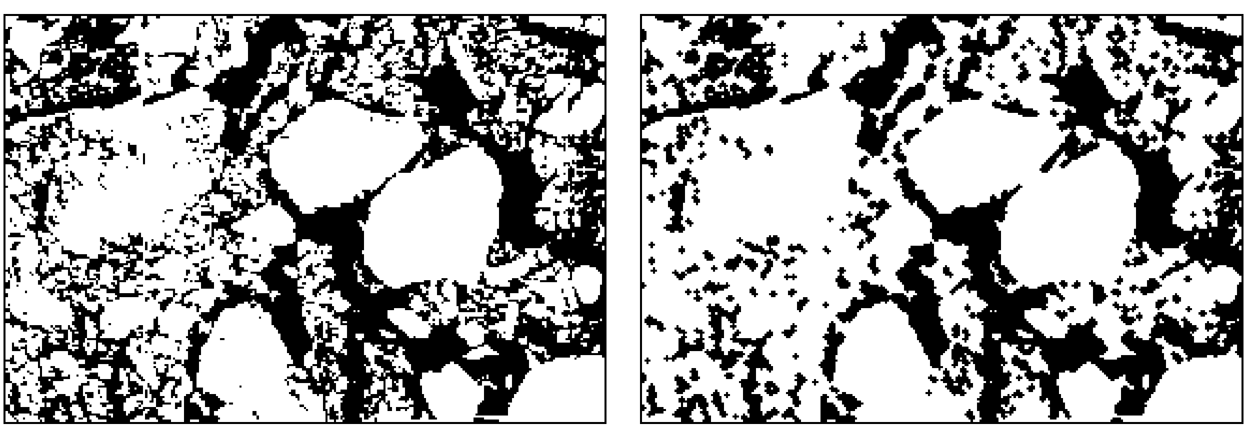

(a) (b)

Figure 5.1: **(a)** Subset of the soil image. **(b)** Morphological closing of (a) using a lattice approximation to a disc of radius 1.0.

the ideas in this book, which are based on arithmetic. It is based on ideas developed by J. Serra and G. Matheron of the Ecole des Mines in Fontainebleau, France. The seminal work is Serra (1982). Other useful introductions may be found in Serra (1986), Haralick, Sternberg and Zhuang (1987) and Haralick and Shapiro (1992, Ch.5). Morphology has become popular in recent years. The basic operations are available in many image analysis software packages. Morphology lends itself to efficient parallel hardware implementations and computers using this are available.

This chapter develops the ideas of morphology mainly for binary images. Many applications use thresholded versions of greyscale images, usually where noise is not a dominant feature. The extension to greyscale images will also be covered. This will be done by regarding a greyscale image as a binary image in three dimensions (§5.5).

As morphological operations are based on sets, set notation will be used in this chapter. A reminder of the basic concepts is given in Fig 5.2, using sets of lattice points. Sets are simply groups of pixels, and the terminology is just a convenient way of describing what pixels lie in particular groups. (The reader interested in acquiring a fuller knowledge of set theory should read Vilenkin (1968).) However, all definitions in set notation will also be explained in words. §5.1 introduces the basic ideas of morphology, §5.2 to §5.4 look at operations which affect particular aspects of images and §5.5 describes how the ideas may be extended to greyscale images. In §5.6, we will summarise the key results.

5.1 Basic ideas

Morphology is based on set theory, and so the fundamental objects are sets. For morphology of binary images, the sets consist of pixels in an image. Readers interested in the full mathematical background are referred to Serra (1988).

Fig 5.3(a) shows an example of an image containing three sets of black pixels. Either the pixels labelled 0 (displayed as black) or those labelled 1 (displayed as white) in a binary image may comprise the sets of interest. Unless stated otherwise, we will take the black pixels to be the sets of interest. In this chapter, when we refer to operations on the *image*, we shall be referring to operations on the set of all black pixels. Usually this will be the union of several separate sets of black pixels — what we think of as individual *objects*, such as A, B and C in Fig 5.3(a). The white pixels are the **complement** (Fig 5.2(e)) of the black pixels, the complement of a set being the set of elements it doesn't contain. Any operation which affects the set of black pixels will also affect the set of white pixels. For example, removing a pixel from the set of black pixels naturally creates a new white pixel.

We can see that the sets in Fig 5.3(a) have different shapes. For example, set B is about the same size (contains about as many pixels) as set A, but is a different shape — it is longer and thinner. One way of describing such differences is in terms of intersections with 'test' sets. If we let D be the simple 2×2 set shown in Fig 5.3(b), then it is possible to *translate* D to positions such that $D \subset A$, whereas no translations can be found to satisfy $D \subset B$. (where $\subset$ denotes 'is a subset of' : see Fig 5.2(c)). This is a consequence only of the shapes of A and B, and so

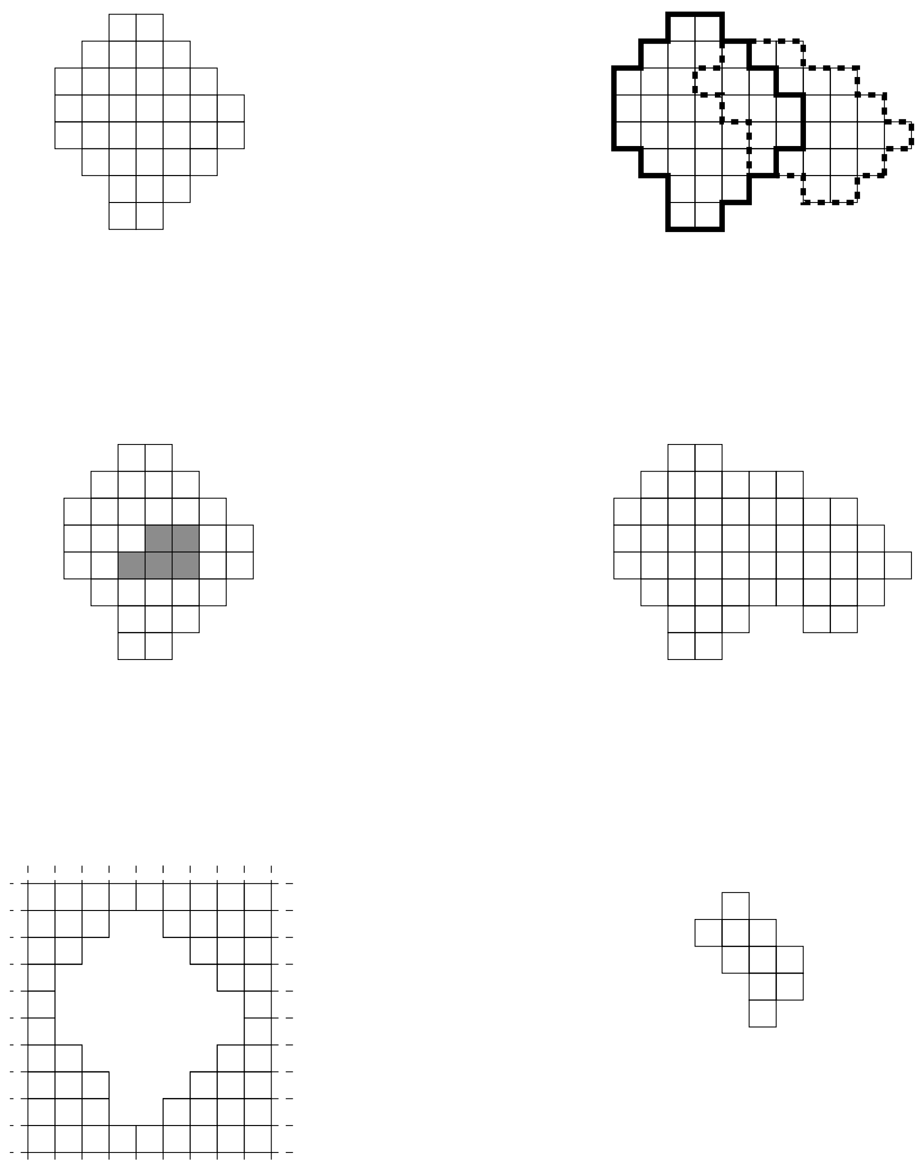

Figure 5.2: Basic ideas of set theory. (a) The set A. The pixel x is an element of A. (written $x \in A$); (b) Two overlapping sets, A and B; (c) Subsets: The shaded pixels (C) are a subset of A. (written $C \subset A$); (d) The union of A and B (written $A \cup B$); (e) The complement of A (written A^c); (f) The intersection of A and B (written $A \cap B$).

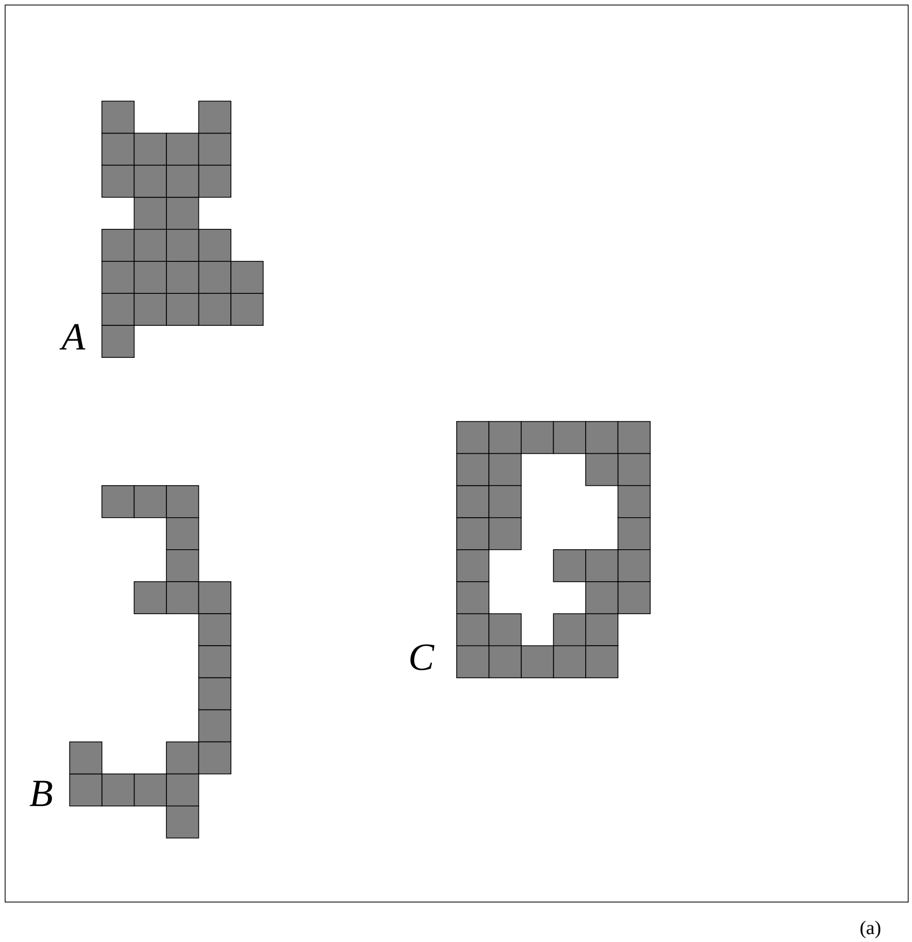

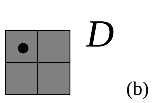

Figure 5.3: **(a)** Example of some sets with different shapes. **(b)** A test set. The • in the top right pixel indicates that we have selected this as the reference pixel.

such tests provide a means of analysing shape information in the image. Test sets such as D will be termed **structuring elements**. It is an important idea that these structuring elements can be placed at any pixel in an image (although rotation of the structuring element is not allowed.) To do this we use some *reference pixel* whose position defines where the structuring element has been placed. The choice of this reference pixel is often arbitrary. We may choose, for example, the centre pixel if the set is symmetric, or one of the corners if it is a polygon. (The only effect of how this choice is made is to translate the position of the result of a morphological operation.)

Morphological operations transform the image. The most basic morphological operation is that of **erosion**. Suppose A is a set (a binary image or part of it) and S is a structuring element. If S is placed with its reference pixel at (i, j) we denote it by $S_{(i,j)}$. Then the erosion of A by S is defined to be the set of all pixel locations for which S placed at that pixel is contained within A. This is denoted $A \ominus S$ and may be written

$$A \ominus S = \{(i, j) : S_{(i,j)} \subset A\}$$

For the example in Fig 5.3, let $S = D$, a block of 4 pixels. If we use the top left corner pixel of D as the reference pixel, then the result of the erosion $I \ominus D$, where I denotes the whole image, i.e. $A \cup B \cup C$, is shown in Fig 5.4. To see how the erosion may be performed on an image we first note that if D is placed with its reference pixel at (i, j), then $D_{(i,j)}$ consists of the 4 pixels $(i, j), (i + 1, j), (i, j + 1)$ and $(i + 1, j + 1)$, i.e.

$$D_{(i,j)} = \{(i, j), (i + 1, j), (i, j + 1), (i + 1, j + 1)\}.$$

If (and only if) all of these are black will the pixel (i, j) in the eroded image be black. If we let g be the eroded image, then an algorithm to perform the erosion will operate as follows, recalling that $f_{ij} = 0$ or 1 according to whether the pixel is black or white:

1. Set $g_{ij} = f_{ij}$. If $f_{ij} = 1$, go to step 5.
2. If $f_{i+1,j} = 1$, set $g_{ij} = 1$ and go to step 5.
3. If $f_{i,j+1} = 1$, set $g_{ij} = 1$ and go to step 5.
4. If $f_{i+1,j+1} = 1$, set $g_{ij} = 1$ and go to step 5.
5. Move to the next pixel in the image and go to step 1.

The above algorithm is efficient and easy to generalise to erosions with other sets. A simpler to program, but less efficient, algorithm for this erosion would be to compute

$$g_{ij} = 1 - (1 - f_{ij}) \times (1 - f_{i+1,j}) \times (1 - f_{i,j+1}) \times (1 - f_{i+1,j+1})$$

since g_{ij} will be zero only if $f_{ij}, f_{i+1,j}, f_{i,j+1}$ and $f_{i+1,j+1}$ are all zero. This is less efficient since it involves arithmetic which is usually slower on a computer than making comparisons.

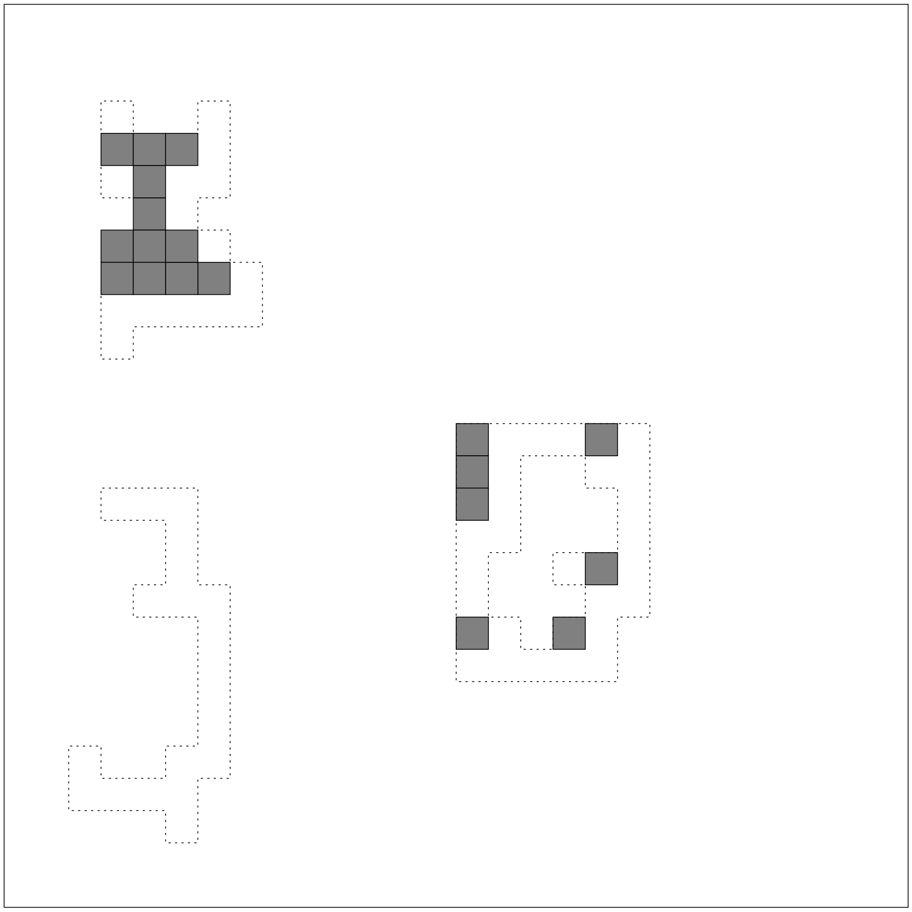

Figure 5.4: The image in Fig 5.3**(a)** after erosion with the structuring element D.

A complementary operation to that of erosion is **dilation**. It is defined simply as the erosion of the complement of a set. If A^c denotes the complement of A, then the dilation of a set A by a set S, denoted $A \oplus S$, is defined by

$$A \oplus S = (A^c \ominus S)^c.$$

This definition easily leads to algorithms for obtaining the dilation.

There is an alternative, equivalent definition of the dilation. Let $S'_{(i,j)}$ denote the reflection of $S_{(i,j)}$ in (i,j), i.e. the rotation of S through 180^o about (i,j). Then $A \oplus S$ contains all the pixels lying in any $S'_{(i,j)}$ for which $(i,j) \in A$. We can think of the dilation as placing a copy of S' at every pixel in A. If the reference pixel is symmetrically placed, then $S' = S$. We may write this definition:

$$A \oplus S = \bigcup_{(i,j) \in A} S'_{(i,j)}$$

where $\cup$ denotes set union (Fig 5.2(d)).

Erosion and dilation are illustrated in Figs 5.5(c) and 5.5(d). The binary image used here is a subset of the the turbinate image (Fig 1.1(a)) thresholded at 128.

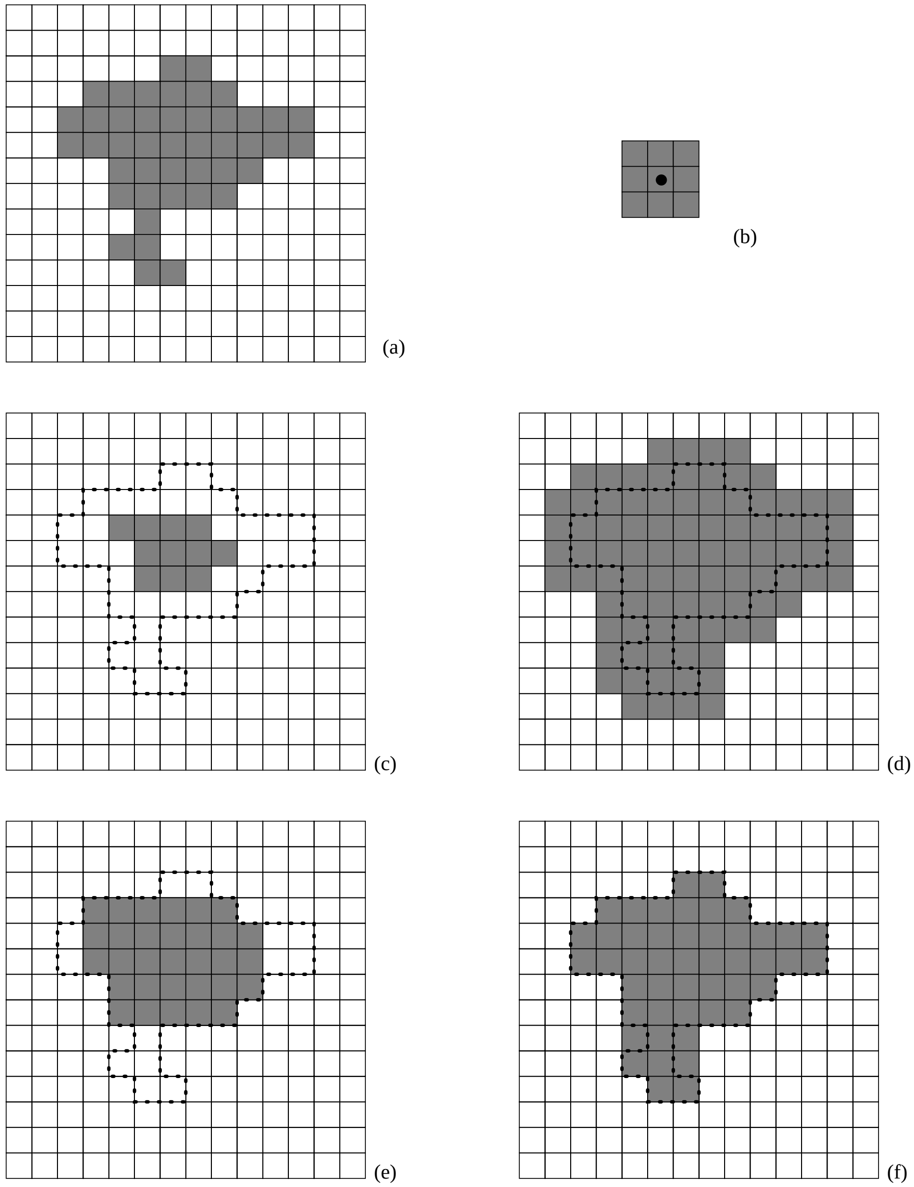

Figure 5.5: Illustration of basic morphological operations: **(a)** Original set, extracted as a subset from the turbinate image, **(b)** structuring element: a square of side 3. The reference pixel is at the centre. **(c)** erosion, **(d)** dilation, **(e)** opening and **(f)** closing.

Two widely used operations are those of **opening** and **closing**, which are often denoted by $\psi_S(A)$ and $\phi_S(A)$ respectively. These are defined as

$$\psi_S(A) = (A \ominus S) \oplus S'$$

$$\phi_S(A) = (A \oplus S) \ominus S'$$

so that an opening is an erosion *followed* by a dilation, and a closing is a dilation followed by an erosion. They are complementary, in that applying one to A is equivalent to applying the other to A^c. Another equivalent definition of the opening is that it is the union of all sets $S_{(i,j)}$ which are contained in A:

$$\psi_S(A) = \bigcup_{S_{(i,j)} \subset A} S_{(i,j)}$$

This differs from the erosion which consists of only the reference pixels, rather than the whole set for which $S_{(i,j)}$ is contained in A. The closing is defined similarly, but applied to A^c. These operations are illustrated in Fig 5.5. For the structuring element used here, both have the effect of smoothing the image. The opening does this by removing pixels from the set, the closing by adding them. Note that both opening and closing are **idempotent** operations, that is applying them more once produces no further effect:

$$\psi_S(\psi_S(A)) = \psi_S(A)$$

and

$$\phi_S(\phi_S(A)) = \phi_S(A).$$

A more general morphological operation than erosion is the **hit-or-miss transform.** The structuring element is a set with two components, $S^1_{(i,j)}$ and $S^2_{(i,j)}$, placed so that both reference pixels are at position (i, j). The hit-or-miss transform of a set A is then defined as:

$$A \circledast S = \{(i, j) : S^1_{(i,j)} \subset A; S^2_{(i,j)} \subset A^c\}$$

that is the set of positions (i, j) for which $S^1_{(i,j)}$ is contained in A and $S^2_{(i,j)}$ lies completely outside A. It follows that for the transform to make sense, $S^1_{(i,j)}$ and $S^2_{(i,j)}$ must not overlap. It can be seen that an erosion is a special case of a hit-or-miss transform, with $S^2_{(i,j)} = \emptyset$ (the empty set).

Operations of erosion, dilation, opening, closing and hit-or-miss can extract a wide range of types of information about a binary image and the components it contains, as well as transforming the image to remove or retain objects satisfying some criteria regarding their structure. What aspects of sets we look at depends on the structuring elements used. Discs or squares of different sizes interact with the size and shape of objects. These are considered in §5.2. Hit-or-miss transforms which interact with the boundary of sets give information on connectivity properties (§5.3), and structuring elements which consist of pairs of separated pixels are affected by the texture of the sets under study (§5.4).

5.2 Operations for size and shape

The most straightforward morphological operations, and the easiest to use, are those relating to the size and local shape properties of objects in images. For this we use structuring elements

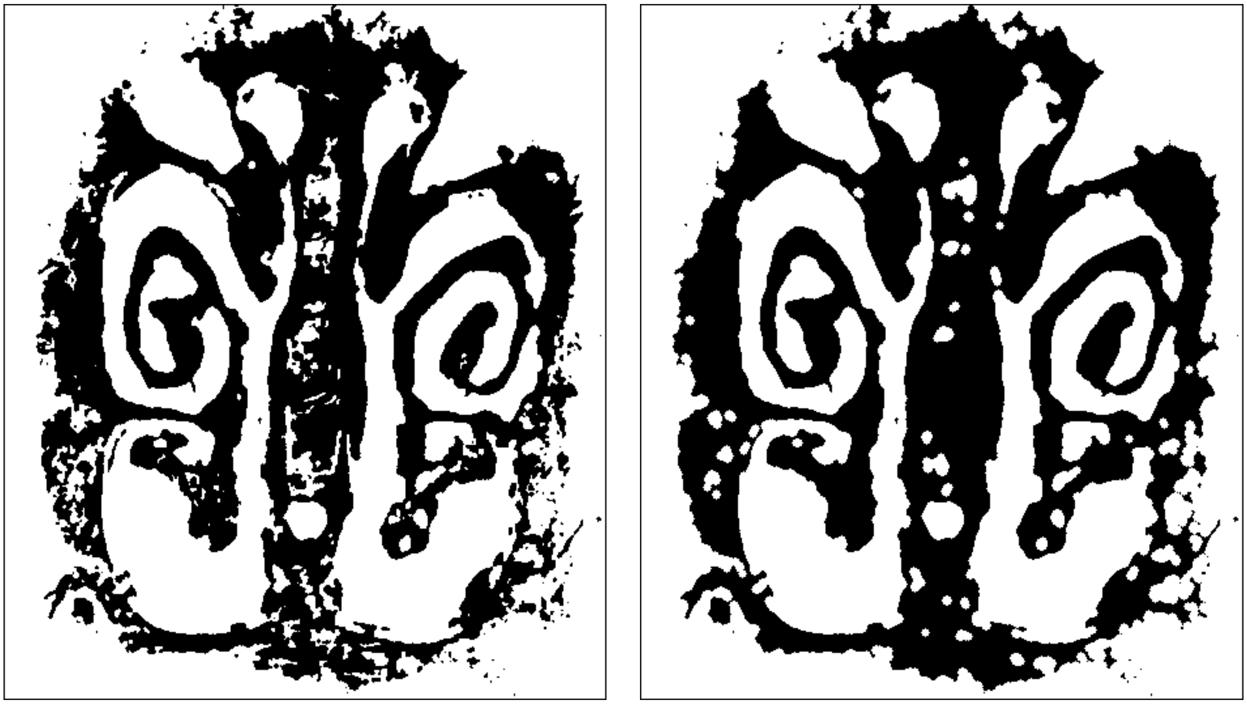

(a) (b)

Figure 5.6: **(a)** The binary turbinate image. **(b)** The turbinate image after closing with a disc of radius 3 pixels.

which are discs or squares. Image components, or their background, will be affected according to whether they can accommodate discs or squares of the size used. Fig 5.6 shows the effect of a closing with a lattice approximation to a disc of radius 3 pixels on the black part of the turbinate image. This could equally be regarded as an opening of the white part of the image. We see that it has had the effect of removing small black speckly features in the white area of the image, when compared with the original, and is akin to the manual cleaning up shown in Fig 1.2(b). Although it has not accomplished all the cleaning desired, it has done much of it, and further operations could go further.

Discs are a natural choice for a structuring element when we are interested in size criteria, since they are **isotropic** (the same in all directions). It is possible to combine operations involving discs. If we denote the disc of radius r as B_r then the following properties hold, where X denotes any image or set:

- Eroding with two discs, one after the other, is equivalent to eroding with one big disc, whose radius is the sum of the radii of the smaller discs:

$$(X \ominus B_r) \ominus B_s = X \ominus B_{r+s},$$

- Opening with two discs is equivalent to opening with the larger disc only:

$$\psi_{B_r}(\psi_{B_s}(X)) = \psi_{B_{\max(r,s)}}(X).$$

There are analogous properties for dilation and closing respectively. Similar properties can be stated for L_r, the square of size $(2r+1) \times (2r+1)$.

The above properties are useful in constructing efficient algorithms for performing operations with large structuring elements. For integer k, the erosion $X \ominus B_k$ is equivalent to $X \ominus B_1 \ominus B_1 \ominus \cdots \ominus B_1$ (k times). However, the exact equivalence of these operations breaks down on the lattice, since no lattice disc corresponds exactly to a disc in continuous space (unlike the square, for which the equivalence *is* exact). Fig 5.7 shows the differences between discs of radius 2, 4, 6 and 8 and what is obtained when they are produced by multiple dilations of the disc of size radius. (Discs of radius r are defined as the set of lattice points whose distance from the centre pixel is $\leq r + \frac{1}{2}$.)

Another useful approach for implementing operations with discs of large radii is the **distance transform**. This replaces the pixel values in one part of a binary image (say the pixels labelled 0) with their distance to the nearest pixel labelled 1. It is then straightforward to perform an erosion with a disc of radius r simply by removing all pixels whose distance label is less than r. A dilation is similarly performed by eroding the background. Forming an exact distance transform is computationally intensive, but efficient algorithms which form good approximations to distance transforms are available (see Danielsson, 1980 and Borgefors, 1986). One easy-to-program and reasonably accurate algorithm is as follows:

The algorithm operates in two passes through the image. We start with the binary image, and assume we want to perform the transform on the black part. We will record for each black pixel where the nearest white pixel is. For pixel position (i, j) denote this by (y_{ij}, x_{ij}). Let its distance from (i, j) be d_{ij}, so that

$$d_{ij} = \sqrt{(i - y_{ij})^2 + (j - x_{ij})^2}$$

This is *Euclidean* distance — see §6.1.2 for other distance measures. On the *first pass* we visit each pixel (i, j) in turn, starting at the top left and moving along each row in turn in a raster scan (§4.2).

- If $f_{ij} = 1$, then $d_{ij} = 0$ and (y_{ij}, x_{ij})=(i, j).
- If $f_{ij} = 0$, then (y_{ij}, x_{ij}) becomes whichever of the locations $(y_{i-1,j-1}, x_{i-1,j-1})$, $(y_{i,j-1}, x_{i,j-1})$, $(y_{i+1,j-1}, x_{i+1,j-1})$ and $(y_{i-1,j}, x_{i-1,j})$ is nearest to (i, j). d_{ij} becomes the distance from (i, j) to (y_{ij}, x_{ij}). Note that $(i-1, j-1), (i, j-1), (i+1, j-1)$ and $(i-1, j)$ will already have been visited. Ties may be resolved arbitrarily.

On the *second pass* we start at the bottom right and move towards the top left going right to left along each row in turn, performing the same steps, but re-setting (y_{ij}, x_{ij}) to whichever of (y_{ij}, x_{ij}), $(y_{i+1,j}, x_{i+1,j})$, $(y_{i-1,j+1}, x_{i-1,j+1})$, $(y_{i,j+1}, x_{i,j+1})$ and $(y_{i+1,j+1}, x_{i+1,j+1})$ is nearest to (i, j), the first of these having been obtained during the first pass. At each step, d_{ij} is re-evaluated.

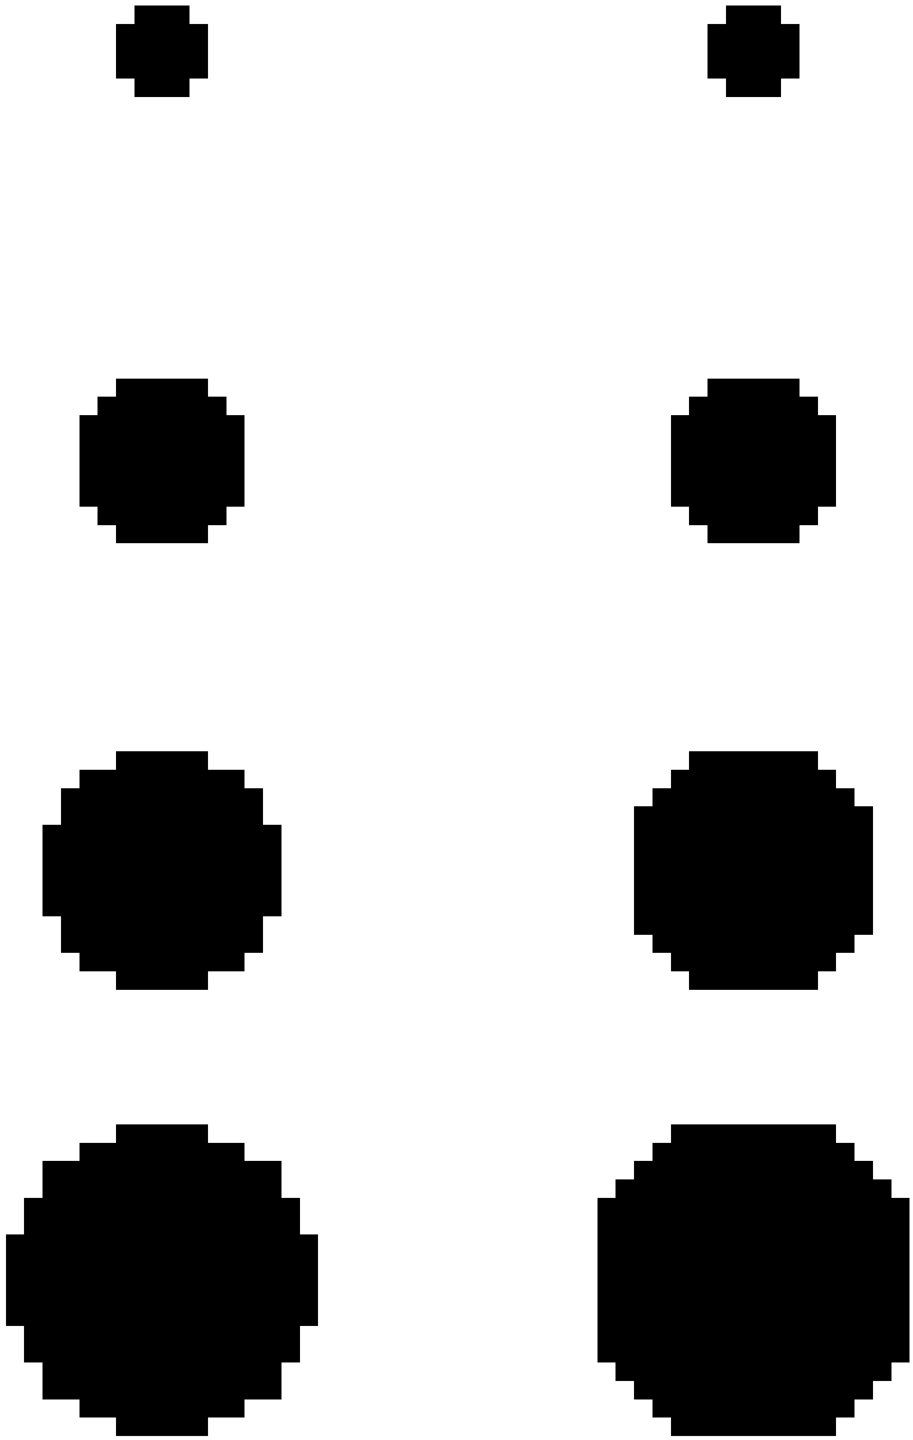

Figure 5.7: The images on the left are dilations of a single pixel with discs of radius (from top to bottom) 2, 4, 6 and 8. The images on the right are obtained by repeatedly dilating with a disc of radius 2.

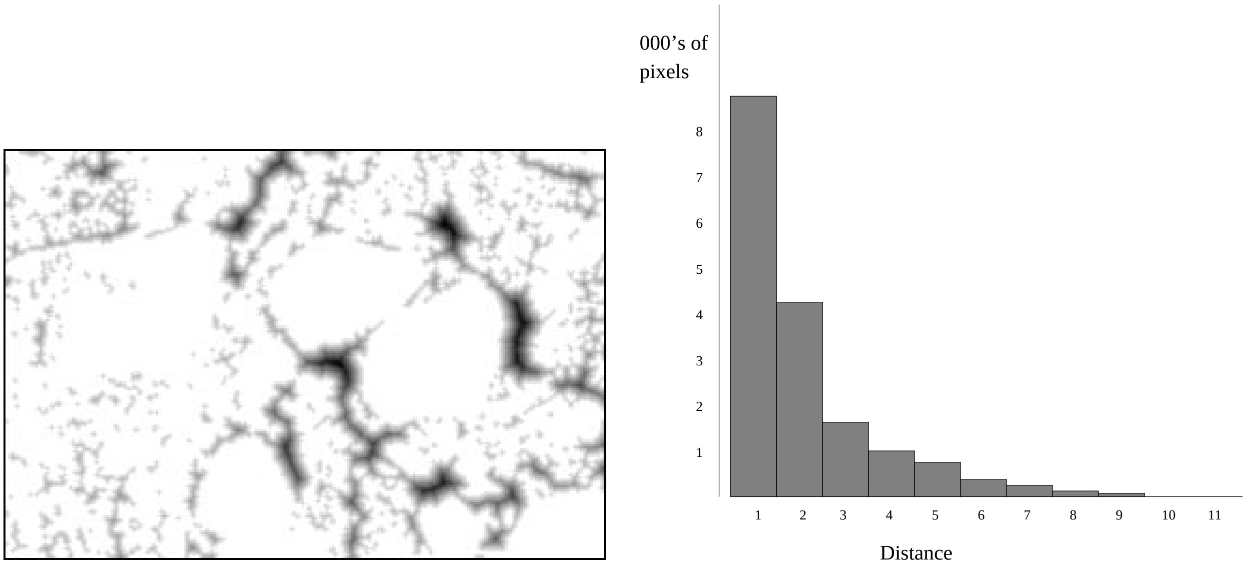

(a) (b)

Figure 5.8: **(a)** Distance transform of the thresholded soil image Fig 5.1(a). The pixel value is proportional to the distance the pixel is from the nearest pixel shown as white in (a). Note that here larger distances/pixel values are shown *darker*. **(b)** A histogram of the distances.

If any black pixels lie on the image boundary, then we will need to define (y_{ij}, x_{ij}) for pixels beyond the boundary. For example, we may set them to be at a great distance (for example $(y_{0j}, x_{0j}) = (-1000, -1000)$ for all j) in which case the distance for boundary black pixels will be to the nearest white pixel in the image. Variations on the above algorithm and a discussion of accuracy, may be found in Borgefors (1986). This includes algorithms based on other definitions of distance on a lattice (see §6.1.2), which are quicker to compute, although possibly less useful in practice.

Fig 5.8(a) shows this approximate distance transform on the thresholded soil image after a closing with a disc of radius 1 (Fig 5.1(b)). The transform has been applied to the dark part of the image, which represents the soil pores. In addition to being of use in morphological erosions and dilations, the distribution of distances can be interpreted in terms of the size distribution of the pores. This is of relevance in studying the movement of gas and microbes through the pore space. For example, we can see what pore space can accommodate microbes of a given size. Fig 5.8(b) shows a histogram of the distances. This distribution can also be used to develop a model of the pore space in order to study its properties in three dimensions (Glasbey, Horgan and Darbyshire, 1991). Distributions such as this, or those based on openings of different sizes, measure properties of the size distribution of objects in an image. This is known as **granulometry**. The morphological study of size is described by Serra (1982, Ch.X). Related to the distance transform is the **erosion propagation** algorithm, which allows us to associate edges with particular objects, even when they overlap. For details and an example, see Banfield and Raftery (1992).

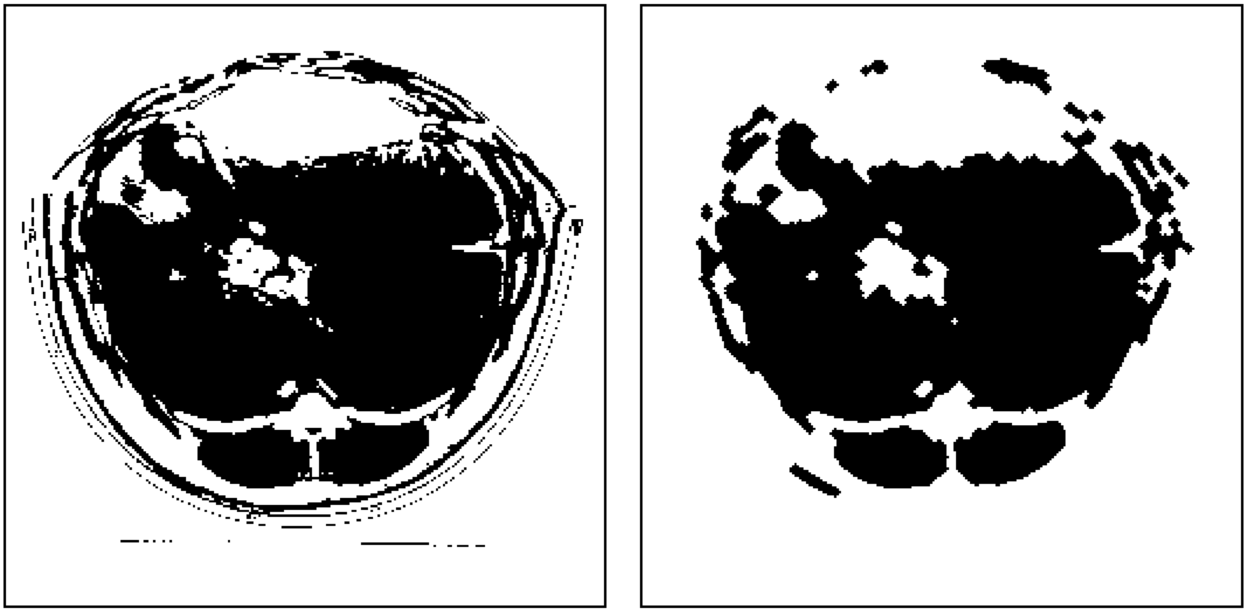

(a) (b)

Figure 5.9: **(a)** The X-ray image, thresholded to select pixel values between -69 and 100. **(b)** The result of opening with a disc of radius 2 followed by closing with a disc of radius 1.

Alternating sequential filters

Openings and closings may be combined when we require an operation which smooths both parts of a binary image. The simplest such combination is an opening followed by a closing or vice versa: $\phi_S(\psi_S(A))$ or $\psi_S(\phi_S(A))$. It can be useful to apply a sequence of such operations, with structuring elements of increasing size. This is because small speckle features, which would interfere with sizing operations with the larger structuring elements, are removed by the smaller structuring elements. If we denote the operation $\psi_k \phi_k \psi_{k-1} \phi_{k-1} \ldots \psi_1 \phi_1 X$ as $G_k(X)$ for integer k, where ψ_k and ϕ_k denote opening and closing with structuring elements of size k, then

$$G_m(G_k(X)) = G_{\max(m,k)}(X),$$

with the special case $m = k$ implying idempotence. Similar properties apply to the operation which begins with an opening. Although these combinations of opening and closing have attractive properties, any combination which seems to suit a purpose can be used. Fig 5.9(a) shows the X-ray image, thresholded to select pixels in the range -69 to 100 Hounsfield units (see §1.1.3), which are displayed as black. Fig 5.9(b) shows the effect of an opening with a disc of radius 2 followed by a closing with a disc of radius 1. This was found to be the most effective for the purpose intended, the measurement of the eye muscles towards the bottom of the image. This measurement is described in §6.1

We conclude this section by noting that openings and closings with elements other than discs or squares can be useful. For example, it is possible to preserve or remove elongated objects in an image by using a structuring element which is a line segment. If the objects have varying orientations, it will be necessary to combine the results of using this structuring element at several orientations.

5.3 Operations for connectivity

Often it is necessary to consider the connectivity of the objects in an image. This is something which is not taken into account in the sizing operations described in the previous section, in which the operators did not interact with the total size of an object, but only its 'local' size, which is an aspect of the shape of the object. If we wish to perform operations pertaining to a connected object as a whole, a different approach is needed. Morphological operations can be used to handle much of this, although in some cases a more direct approach can be more efficient. Note that in dealing with connectivity, we can choose to assume either 4-connectivity or 8-connectivity (§4.2). Unless stated otherwise, we shall assume 8-connectivity.

Counting the number of connected components in an image is a useful operation. Although there are morphological operations to do this efficiently, the simplest approach is to label the connected components individually, as described in §4.2. An immediate by-product of this algorithm is a count of the number of connected components.

A morphological tool which does not change connectivity is that of **conditional dilation**, sometimes called **geodesic dilation**. This involves dilating a set and performing an intersection (Fig 5.2(f), sometimes termed an **"and"** operation) with another set. For example, if we perform an erosion which completely removes objects which do not satisfy some shape criterion, we will also have changed the shape of the objects which have not been eliminated. They can be restored by a sequence of dilations constrained to lie *within* the original image. If I is the original image, and J_0 the result of an operation which has completely removed unwanted objects, but left parts of all other sets, then the iteration using L_1, the 3×3 square with its reference pixel at the centre

$$J_{n+1} = (J_n \oplus L_1) \cap I$$

repeated until $J_{n+1} = J_n$ will produce a result in which the unwanted objects are removed, and the remaining objects are as they were in I. However, it is often easier to work with labelled objects. In the above case, we would simply need to keep track of the labels of the eroded objects, and use this information to filter the image of labels, so that only certain objects remain.

Conditional dilation also allows us to perform other useful image operations. For example, if we wish to remove objects which touch the boundary of the image, then a conditional dilation starting from the set of all black pixels at the boundary will find all pixels in boundary-touching objects. A conditional dilation of the white part of the image, starting from the the white boundary pixels, will fill the image background, but not holes in black objects. This can be used to fill such holes, if this is desired.

Thinning and thickening

The size and shape operations in §5.2 do not preserve connectedness. We may be concerned not simply with whether or not an object is a single connected component, but with properties such as how many branches and holes it contains. Such properties are referred to as the **topology** of the object. We can ask that this topology be preserved. Operations which do this are termed **homeomorphisms**, and they may be described as **homotopic**. They are known as **thinning** and **thickening** operations, depending on whether they add to or remove black pixels from the image.

The general definition of a thinning of a set A by a structuring element S is that we remove from A a part of A specified by the hit-or-miss transform $A \circledast S$. The thinning is denoted by $A \bigcirc S$ and may be written in set notation:

$$A \bigcirc S = A \setminus (A \circledast S)$$

where the set operation $X \setminus Y$ is a subtraction which results in the elements which are in X but not in Y, i.e. $X \cap Y^c$. To use the hit-or-miss transform, $A \circledast S$, the two components of S must be chosen so that connectivity is unaffected.

Thinning is most often used to reduce objects to a thickness of one pixel. Three features of the way this is done are

- A set of rotations of a basic structuring element is applied in sequence
- The thinning is applied repeatedly until no further change occurs.
- Sometimes it is desirable to **prune** the thinned objects to remove small barbs which are not considered to be part of the true structure of the object.

Fig 5.10 shows some standard structuring elements for thinning. We may think of these structuring elements as eating away at the boundaries of a set without breaking local 4-connectedness or introducing holes inside the set. An algorithm to perform this thinning on an image would apply the hit-or-miss transform with the sequence of structuring elements in Fig 5.10 in turn. The sequence would then be repeated until no change was found to occur. To thin the black part of an image with the first structuring element in Fig 5.10 would mean setting

$$g_{ij} = \begin{cases} 1 & \text{if} \quad f_{ij} = 0 \text{ and} \\ & \quad f_{i-1,j-1} = 0 \text{ and } f_{i,j-1} = 0 \text{ and } f_{i+1,j-1} = 0 \text{ and} \\ & \quad f_{i-1,j+1} = 1 \text{ and } f_{i,j+1} = 1 \text{ and } f_{i+1,j+1} = 1 \\ & \\ f_{ij} & \text{otherwise} \end{cases}$$

The second structuring element is then used to thin g etc. Fig 5.11(b) shows the result of applying them in sequence to a thresholded subset of the fungal image until no further change occurs. The position and length of the hyphae have been unaffected, but they are now only one pixel thick. A different set of operations could be used to thin to a minimal 8-connected set of

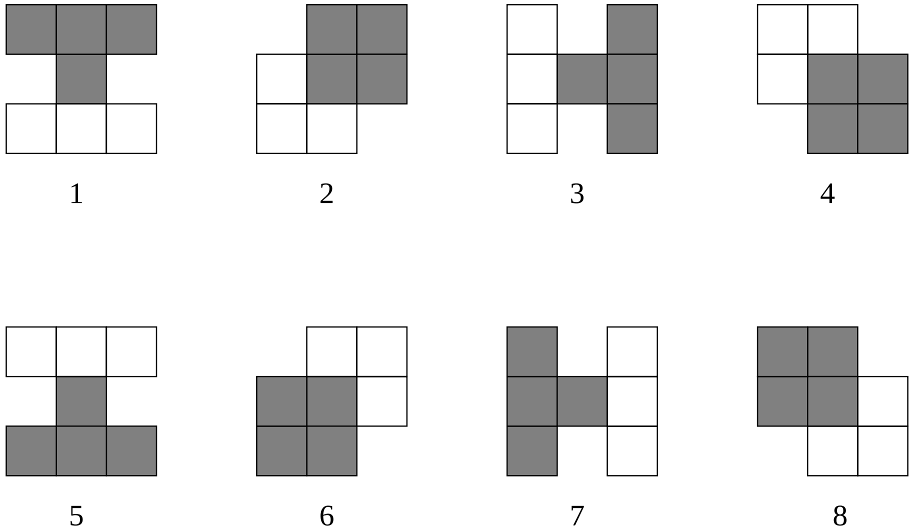

Figure 5.10: Structuring elements for thinning. These are 4 orientations of each of 2 similar structuring elements. The shaded squares are elements of S^1 (i.e. they must be contained in the object) and the empty squares are the elements of S^2, which must lie outside the object. The structuring elements are used in the order $1, 2, \ldots, 8$.

pixels. The above is just one of many algorithms which have been proposed for thinning. For a discussion of some others, see Rosenfeld and Kak (1982; Ch11), Chen and Hsu (1990) and Jang and Chin (1990).

The thickening operation, $A \odot S$, is defined by

$$A \odot S = A \cup (A \circledast S)$$

If S is symmetrical with respect to a set and its complement, then thickening is the equivalent of thinning the complement of the set. One application of doing this is that by thinning the (white) background to a thickness of one pixel, we will have found that part of the background around each black object which is, in some sense, more associated with that object that with other black objects. This is known as the **sk**eleton of **i**nfluence **z**one or **SKIZ**. For more details, see Serra (1982, Ch. XI).

The thinned image above is termed a **skeleton** of the original image — it follows the shape of the image, but is only one pixel thick. This can be a useful transform of the image. For example, with the fungal image, it enables us to estimate the total length of the fungal hyphae by counting the number of black pixels in the skeleton image. (We will need to adjust this for the different lengths of digital distance in the horizontal, vertical and diagonal directions — see §6.1.2.) There are other possible definitions of the skeleton of an object. Our basic requirements would be that

- It preserves the topology of the object.
- It is one pixel thick.

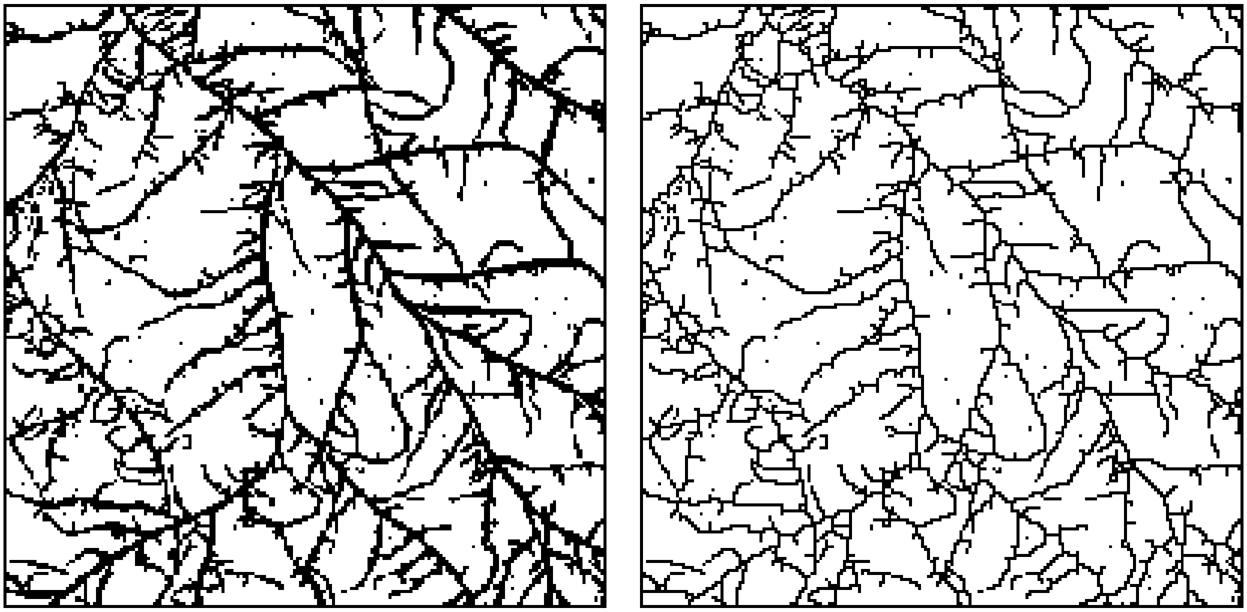

(a) (b)

Figure 5.11: **(a)** Subset of the fungal image thresholded at a pixel value of 150. **(b)** Same image after applying thinning operations of Fig 5.10 until no further change occurs.

- It is in the 'middle' of the object.

The third point could be achieved in various ways. For example, we could ask that any pixel on the skeleton is equidistant from two non-adjacent pixels on the boundary of the object. Such a skeleton is termed the **medial axis**.

The thinning defined above satisfies the first two of these properties, but is only an approximation to the medial axis. Other approaches to obtaining a skeleton are also possible. For example, Kalles and Morris (1993) divide the image into trapezoidal blocks and generate appropriate skeletons for each block, which may then be combined. Other methods are suggested by Arcelli and Sanniti di Baja (1989) and Xia (1989). A discussion of a number of other morphological approaches to defining and constructing the skeleton in a digital image may be found in Serra (1982, chapter XI) and Meyer (1989). A discussion of the relevance of the skeleton and related concepts in studying the shape of biological objects may be found in Blum (1973).

A skeleton can be pruned by looking for end pixels (i.e. pixels in the skeleton with only one neighbouring pixel in the set), and removing those joining another branch of the skeleton within a specified number of pixels along the thinned object. A branch junction can be recognised as one where a pixel has three neighbouring pixels. Pruning can be regarded as a form of thinning, since it involves removing pixels on the basis of some criterion. Pruning is useful if the skeleton is affected by small features which are not of interest or are due to noise. A complete pruning — removing all branches of the skeleton which come to an end point can be used if we wish to preserve only closed loops in a skeleton. (See Serra, 1982, chapter XI).

5.4 Operations for texture

Texture is a concept for which it is difficult to give an exact definition. In one approach (sometimes termed syntactic pattern) it is taken to mean the spatial arrangement of features of an image, and is used when there is some pattern (another difficult word to define!) in this arrangement — when there is something not totally random. These ideas are relevant only where there are features which are repeated in the image. For example, there are several cells in the algal and muscle fibres images, many bands in the DNA image and many hyphae in the fungal image. On the other hand, the features of the X-ray, turbinate and chest images are not repeated, and so the concept of texture as defined above does not arise in these examples.

The most useful and simplest tool for studying the texture of a binary image is erosion by a structuring element consisting of two pixels a specified distance apart, followed by a count of the pixels remaining after doing this. This may be repeated for a number of different distances and orientations. The number of pixels remaining after erosion, as a function of distance, summarises the texture of the binary image. This function may be termed the **auto-crossproduct** function of the image, since it records the number of times pixels a given distance apart have a product of 1 (i.e. are both 1). (Serra (1982, p272) uses the term 'covariance' for this function, although this differs from the usual statistical use of this word.) If the texture is not isotropic, the orientation of the structuring element will be important. The variation in the auto-crossproduct, as a function of the direction of the structuring element, may be used as a description of the **anisotropy** (variation with direction) of the texture.

Fig 5.12 shows the auto-crossproduct function for the thresholded fungal image for four directions: horizontal, vertical, and the two diagonals. There is little evidence of any periodicity, which would be shown by a peak in the function other than at zero. There is also little evidence of any anisotropy. Fig 5.13(a) shows a thresholded section of the DNA image and 5.13(b) shows the auto-crossproduct function for the four directions. Here we see that there is evidence of periodicity in the vertical direction, but not in the horizontal direction. This is a consequence of the regular horizontal banding structure in parts of this image.

5.5 Morphology for greyscale images

The morphology discussed so far was developed for binary images. Morphological techniques are most often used for such images. Morphology is best used with low-noise images, which can often be thresholded with little loss of information. Where the image consists of a few distinct grey levels (or in practice, a few clusters of grey levels), then analysis of the binary images obtained with a few different thresholds may be appropriate. This could be used with the X-ray image, for example.

The most general way of handling a greyscale image is to regard it as a binary set in 3 dimensions. We imagine the image as defining a function in the two-dimensional plane, which produces a map in the third dimension (as in Fig 1.5). The black part of the image is the volume under the function (sometimes termed the **umbra**), and the white part, its comple-

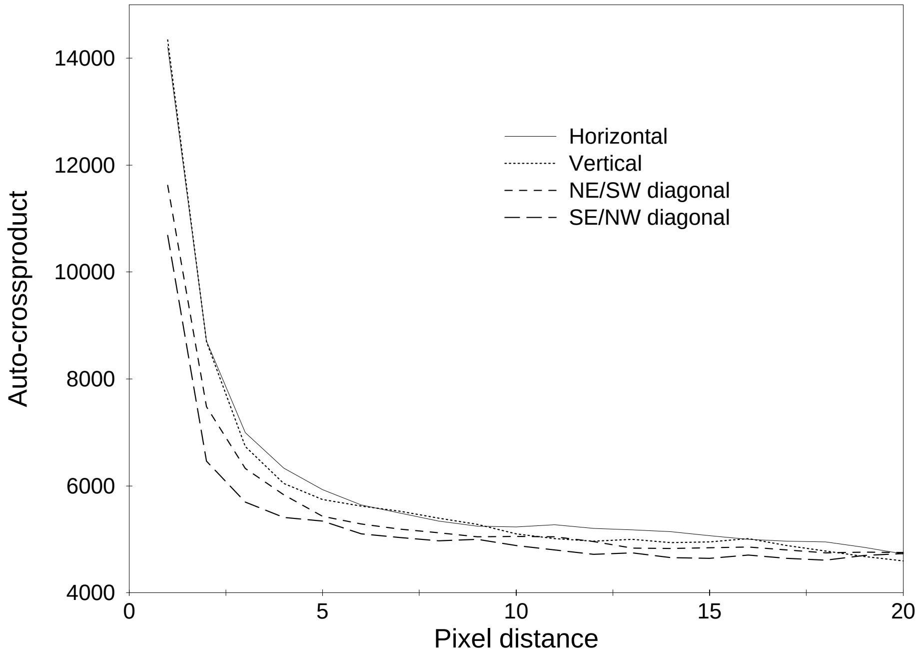

Figure 5.12: Auto-crossproduct function for the thresholded fungal image for four directions: horizontal, vertical, and the two diagonals.

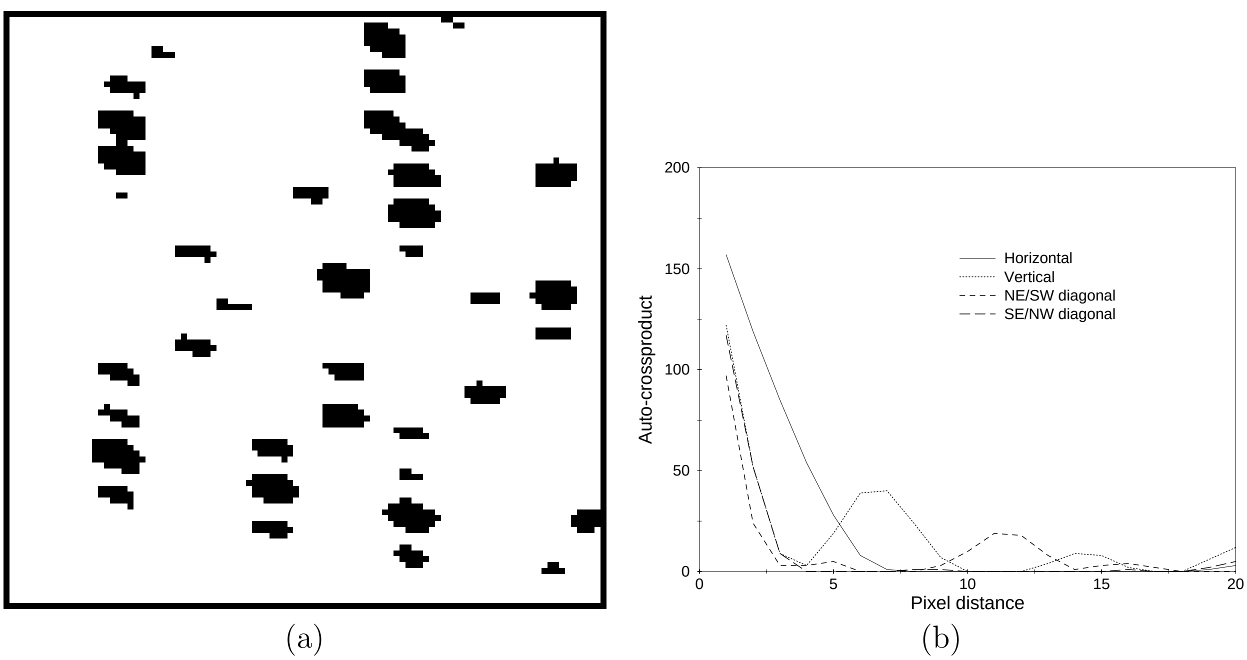

Figure 5.13: **(a)** Section of DNA image, thresholded at 170. **(b)** Auto-crossproduct function for four directions: horizontal, vertical, and the two diagonals.

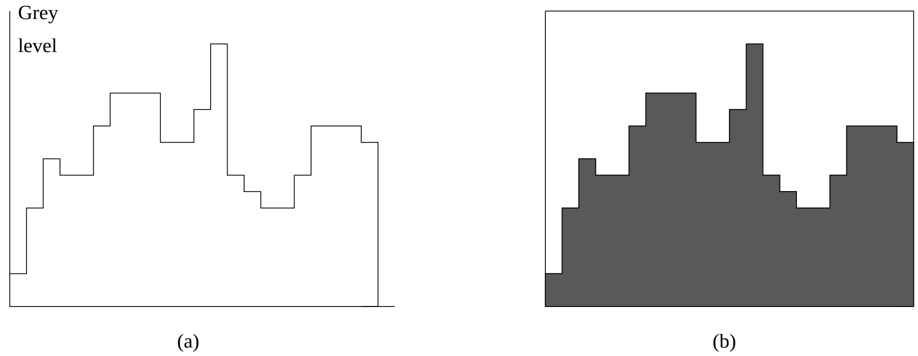

Figure 5.14: **(a)** A one-dimensional function and **(b)** equivalent two-dimensional binary image.

ment, is the volume above the function. This is analogous to using a one-dimensional function to define a set in two dimensions, as illustrated in Fig 5.14. The morphological properties of the two-dimensional set are determined by the features of the one-dimensional function, and so these can be studied by operations on the two-dimensional set. An analogous situation applies to two-dimensional functions and three-dimensional sets, and this is how we approach morphology for greyscale images.

For mathematical morphology in three dimensions, we simply need to consider three-dimensional structuring elements. The only part of the umbra relevant for morphological operations is the top, since the rest is always the same — it continues downward indefinitely. (Similarly, for the one-dimensional function in Fig 5.14, we need only be concerned with the top of the two-dimensional set. The rest provides no information). Similarly, only the top of the structuring element is important, since this is the part which will interact with the top of the umbra. Thus the structuring element can also be considered as a function in 2 dimensions. Commonly used functions are those which are constant over a region, particularly a square or disc (so that the structuring element is a cuboid or cylinder) or are like a hemisphere, so that the structuring element can be thought of as a sphere. A cone is another possibility. Some structuring elements are illustrated in Fig 5.15.

The complement of the binary set we have defined is the region *above* the function. In Fig 5.14(b), this corresponds to the white part of the image. All operations on the volume below the function correspond to a complementary operation on the volume above the function. The area above the function can be thought of as the binary set corresponding to the negative (in the photographic sense) of the image.

Erosion, dilation, opening and closing

As in binary morphology, the operations of erosion, dilation, opening, closing and hit-or-miss play a central role. They are straightforward to obtain, and it can be shown that if the structuring element is flat topped (as in Fig 5.15(f)), the top surface of which is a 2-D set S (which we also use to denote the structuring element itself) then greyscale erosion and dilation are as follows:

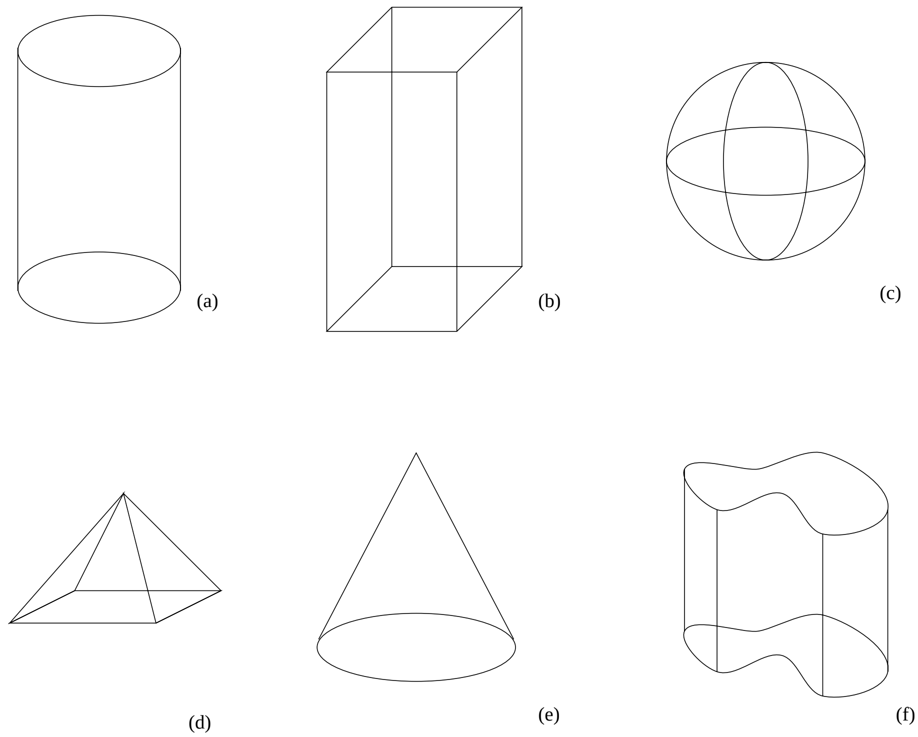

Figure 5.15: Three-dimensional structuring elements: **(a)** cylinder, **(b)** square cuboid, **(c)** sphere, **(d)** square pyramid, **(e)** cone and **(f)** general flat-topped structuring element.

$$(f \ominus S)_{ij} = \min_{(k,l)\in S} f_{i+k,j+l}$$
$$(f \oplus S)_{ij} = \max_{(k,l)\in S} f_{i+k,j+l}$$

where $(k,l) = (0,0)$ is the reference pixel in S. Thus we are transforming f by taking the minimum or maximum of f in a neighbourhood, corresponding to the region S, about each pixel of f. (We also used minimum and maximum filters in §3.4.1.) Fast algorithms for minimum and maximum filters in square or octagonal regions are discussed by van Herk (1992). Openings and closings are defined as before, so that $\psi_S(f) = f \ominus S \oplus S'$ and $\phi_S(f) = f \oplus S \ominus S'$. It can be shown that these operations are equivalent to performing operations with the corresponding two-dimensional structuring elements simultaneously on all possible thresholded versions of the image, and reconstructing a greyscale image from these. When S is a square cuboid or a cylinder, operations with the structuring element S can be decomposed in the same way seen in §5.2. The sequences of openings and closings described in §5.2 can also be used. For an example, see Meyer (1992).

Fig 5.16 shows the effect of applying these operations to the cashmere image using a cylinder of radius 6 as the structuring element. The erosion (Fig 5.16(a)) replaces each pixel with the minimum in a disc of radius 6 centred on that pixel. Since the fibres are darker than their background, or have edges which are darker, this has led to each fibre appearing darker and wider than in the original image. The bright borders on some fibres have disappeared and the fibres which are brighter in the centre no longer have this feature since the pixels in the centre are within 6 pixels of a dark edge. The dilation (Fig 5.16(b)) replaces each pixel with the maximum within a distance of 6 pixels. The bright edges of some fibres are now thicker, and in some cases the whole fibre now appears bright. Some fibres have almost disappeared. The opening (Fig 5.16(c)) is similar to the erosion for this image, but the fibres, while still mostly dark, have now been shrunk back to their original size. The closing (Fig 5.16(d)) has produced a mixture of thick dark fibres and thin bright fibres.

The effect of using a spherical structuring element is demonstrated in Fig 5.17. Comparing its effect with that of a disc is a little like comparing a moving average filter with a Gaussian filter where the weights decay away from the centre. When opening with a sphere, the effect of nearby pixels decreases away from the centre pixel. This can be seen geometrically by thinking of opening with a cylinder as trying to fit a cylinder into the bumps of the function (from below) and similarly with a sphere. Opening or closing with a sphere is sometimes known as a **rolling ball** transform. The differences between Fig 5.16(c) and Fig 5.17 are very subtle. For example, the intersections between the edges of two fibres are less rounded in Fig 5.17. A sphere can be rolled more tightly into the corners of the bright parts of the image between the fibres than a disc can slide. Fast algorithms for erosion and dilation using spheres or discs are described by Adams (1993).

Erosion and dilation with a general structuring element require relatively simple algorithms, but are calculated rather slowly. If the top of the structuring element can be represented by a function g_{kl}, with $g_{00} = 0$ then it follows that

$$(f \ominus g)_{ij} = \min_{(k,l)} \left[f_{i+k,j+l} + g_{kl} \right]$$

$$(f \oplus g)_{ij} = \max_{(k,l)} \left[f_{i+k,j+l} + g_{kl} \right]$$

where (k, l) ranges over the domain of definition of g. A discussion of the use of other structuring elements and of iterated operations may be found in Sternberg (1986). An example of combining operations is described by Skolnick (1986). The use of conditional dilations to extract features satisfying given conditions, and efficient algorithms for doing this, are described by Vincent (1993).

Morphological erosion, openings etc. are often used to smooth an image, and if so they can be thought of as types of non-linear filters. Their effect on the cashmere image is to change the thickness or remove the bright and dark lines associated with the fibres. We may use this to generate an estimate of the fibre diameter distribution, as illustrated in Fig 5.18 (colour plate). Fig 5.18(a) shows the eroded image (Fig 5.16(a)) thresholded at a pixel value of 140. Fig 5.18(b) shows the effect of a **sizing transform** on the black part of the thresholded image. This places discs of maximum radius in the black part of the image — the pixel values are the radius of the largest disc which can be placed in the black part of the image and contain that pixel position. Fig 5.18(b) is shown in pseudocolour. The sizing transform may be obtained in a number of ways. One is to perform openings with discs of increasing integer radii until none of the black part of the image remains. The largest radius at which a pixel remains after the

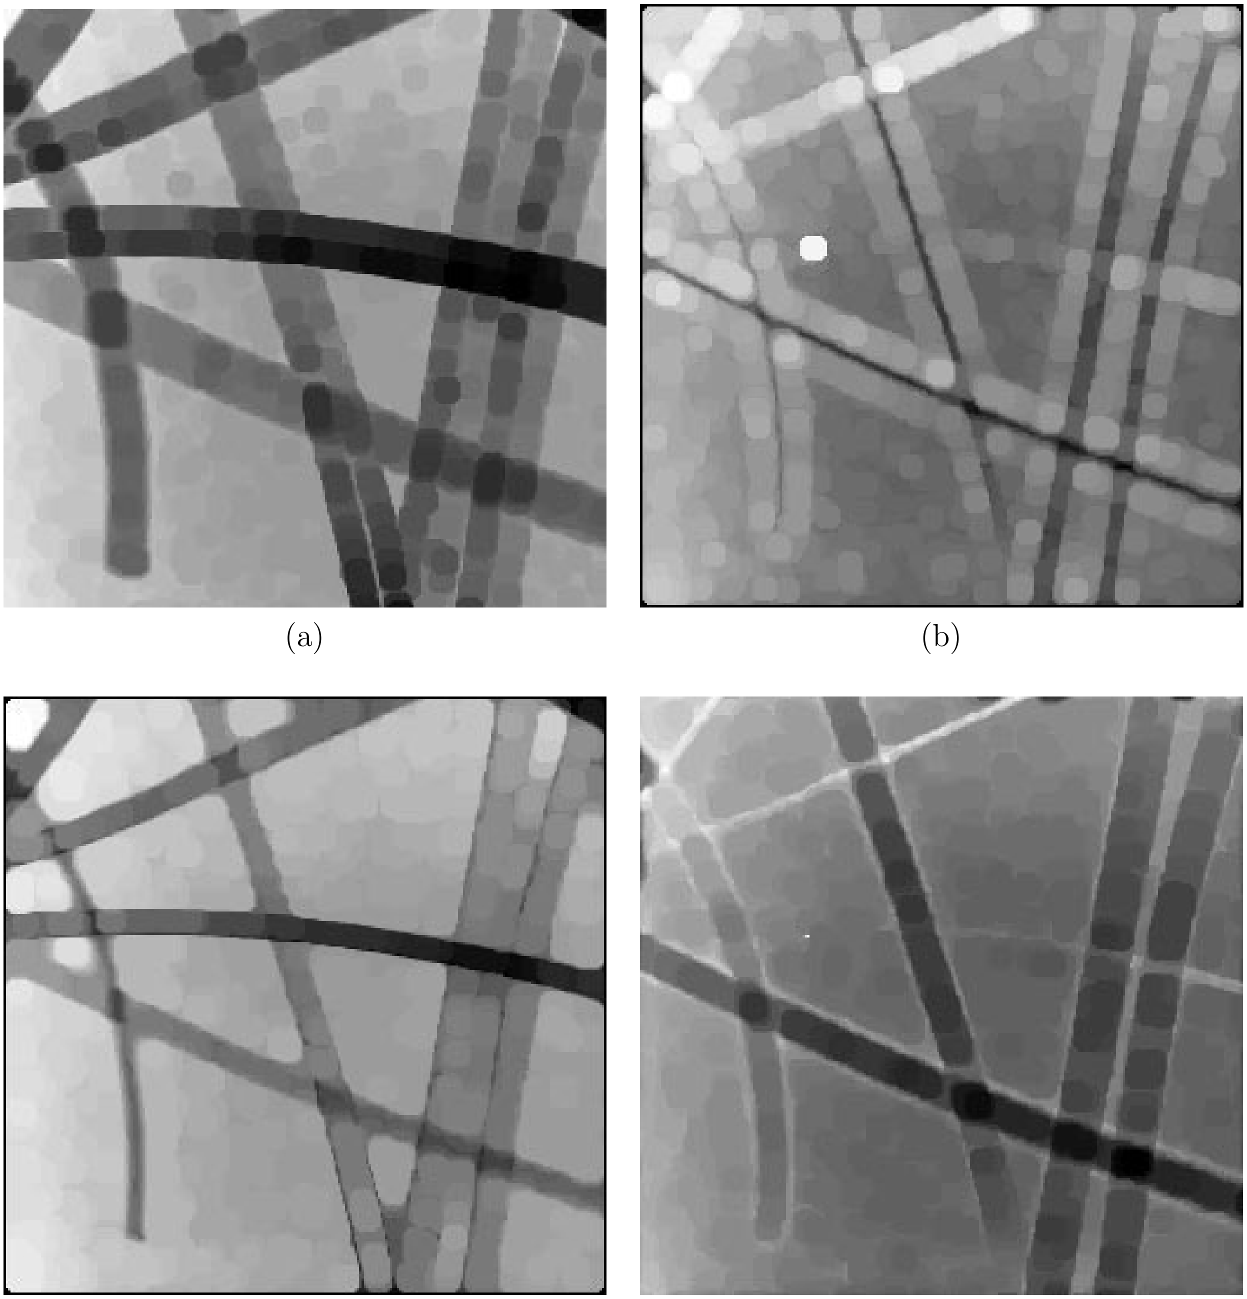

(a) (b) (c) (d)

Figure 5.16: Greyscale morphological operations applied to the cashmere image, using a cylinder of radius 6: **(a)** erosion, **(b)** dilation, **(c)** opening and **(d)** closing.

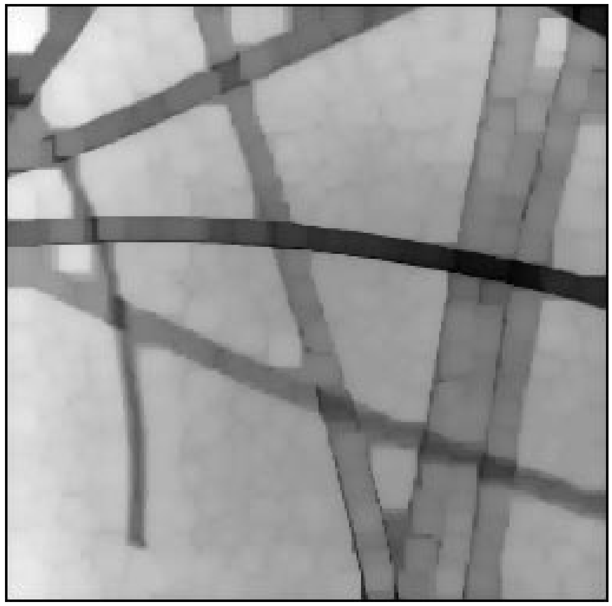

Figure 5.17: Opening of the cashmere image with a sphere of radius 6.

opening is its value in the sizing transform. Fig 5.18(c) shows a histogram of the pixel values in Fig 5.18(b). Each count was divided by the pixel value to produce a length weighted, rather than an area weighted, distribution. In order to use this to estimate average fibre diameters, we will need to remove large values (where fibres coalesce) and small values (due to noise). We will also need to shrink each size value to compensate for the expanding effect of the erosion in Fig 5.16(a). An alternative, which did not work as well for this image, would be to use the opening in Fig 5.16(c). If these difficulties are dealt with, we should be able to use this approach to measure fibre diameters automatically.

The top-hat transform

Openings and closings can be used to derive useful image operations which are not themselves morphological (in the sense of being expressible as hit-or-miss transforms). The range filter, the difference between a greyscale dilation and erosion, is described in §3.4.1. Another one is the **top-hat** transform.

The top-hat transform is used for extracting small or narrow, bright or dark features in an image. It is useful when variations in the background mean that this cannot be achieved by a simple threshold. If C denotes a cylinder (whence the term top hat) with radius r, then subtracting the closing with C from the original image, i.e. taking the transform

$$f - \phi_C(f)$$

will find narrow dark features in the image. This is because the closing will have eliminated them, and they will be apparent when the closing is subtracted from the original image. For finding bright features, the transform $f - \psi_C(f)$, i.e. the subtraction of the opening, is appropriate. Fig 5.19 shows the effect of the top hat transform (subtracting the closing) on the DNA

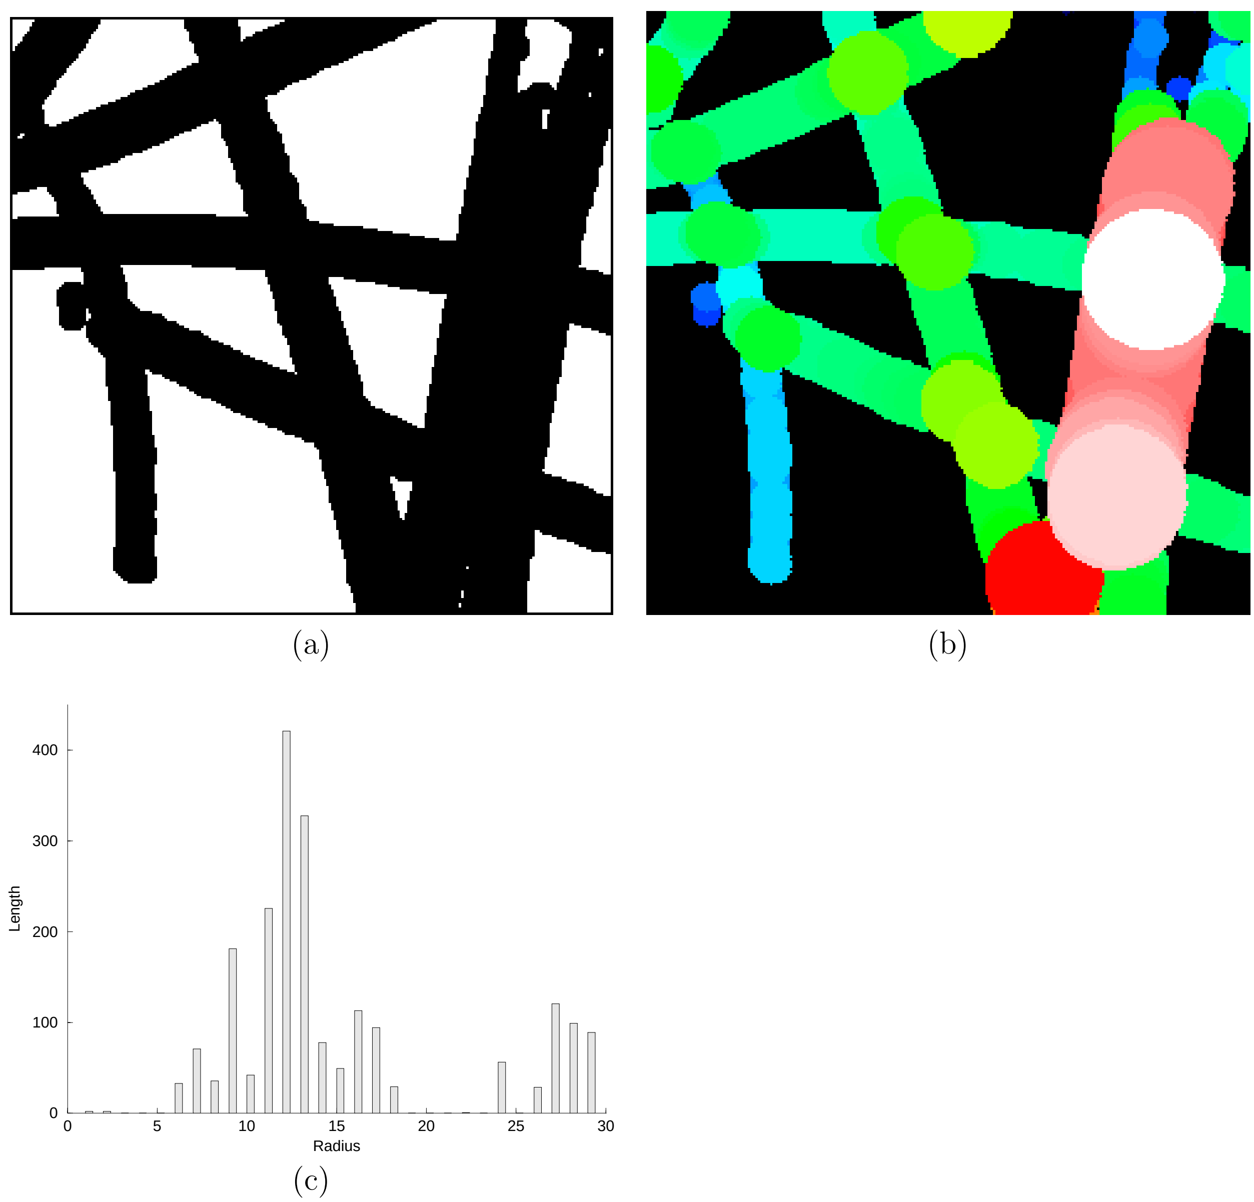

Figure 5.18: Estimating fibre diameters. **(a)** Image of Fig 5.16(a) thresholded at a pixel value of 140. **(b)** Sizing transform of (a), displayed with pseudocolour. **(c)** Histogram of (b).

image. This image shows how dark the bands are relative to the local background. This has removed the trend in the original image for the background to be darker near the top, and in Fig 5.19 (which has been contrast stretched using the minimum and maximum pixel values) the bands stand out more clearly. A disc of radius 4 was used. The radius of the disc should be chosen on the basis of the size of features to be extracted.

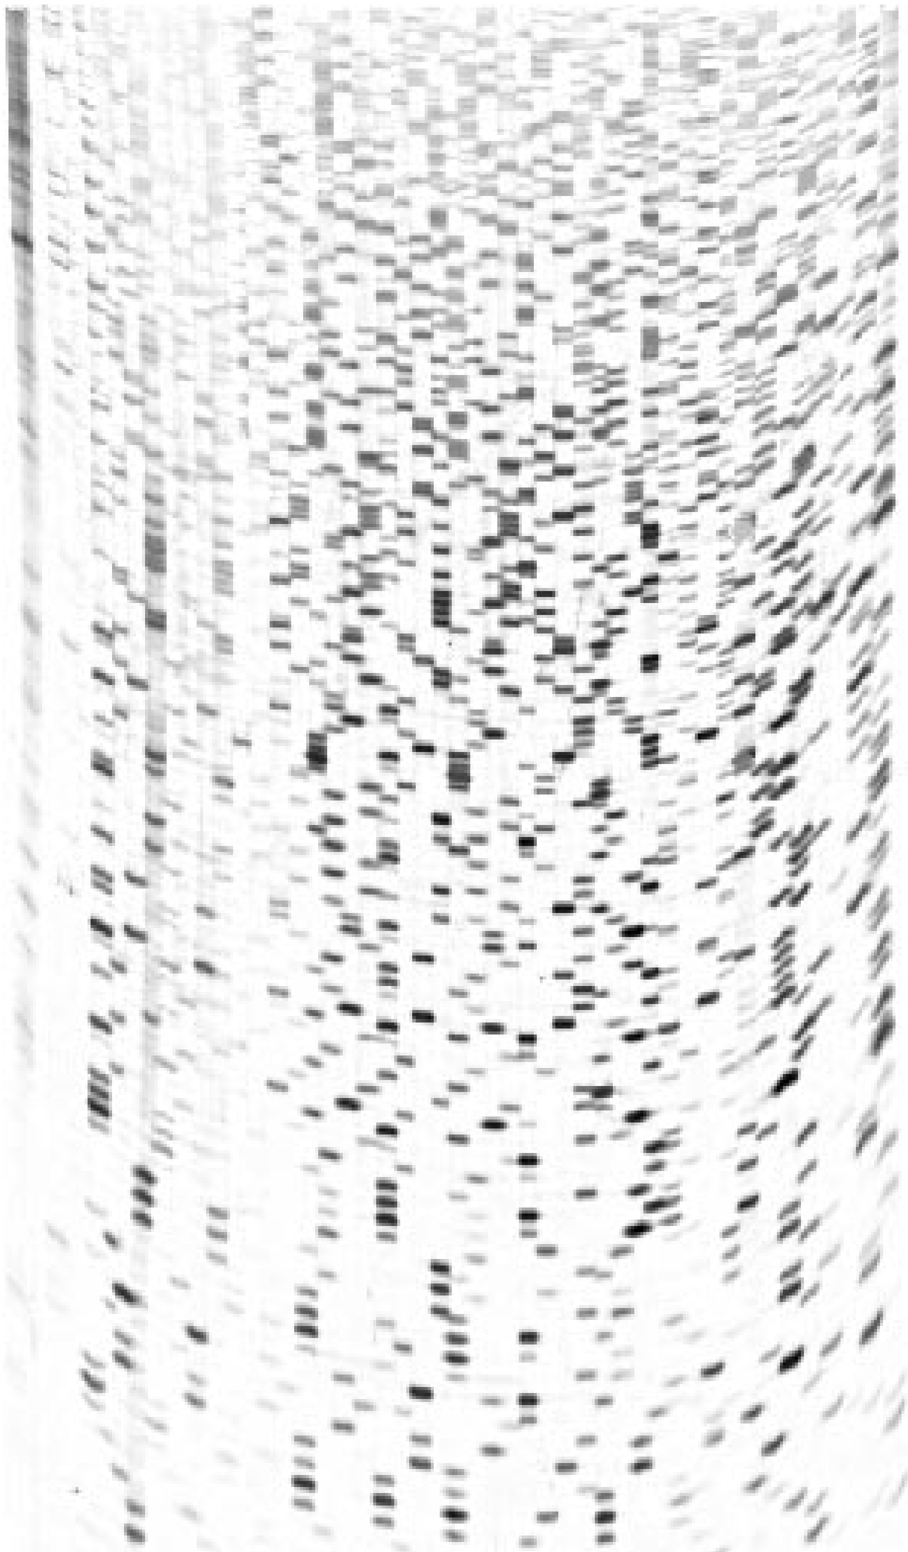

Figure 5.19: Top-hat transform (subtracting the closing) applied to the DNA image. A disc of radius 4 was used.

Greyscale texture

Greyscale texture is more complex to consider than binary texture. Many approaches, both morphological and otherwise have been proposed. A review may be found in Reed and du Buf (1987). Morphological approaches may be based on generalisations of the structuring elements discussed in §5.4. The basic element of two pixels can now have an angle in the third dimension, as well as an orientation. Often, it is simpler to study the texture of a thresholded version of a greyscale image, as was done with the DNA image. If this is an oversimplification, several thresholds may be used to build up a more complete picture.

Greyscale topology

When an image is represented as a function in the 2-dimensional plane (as in Fig 1.5), with the pixel values representing height above the plane, it can be useful to look for features of the resulting ‘landscape’. For example, bright regions in the image can be thought of as hills with the local maxima of brightness being summits or peaks. Similarly, dark regions in the image give rise to channels and pits. We can go further and define concepts such as ridges, valleys, watersheds, divides etc. in the ‘topography’ of an image. These sometimes correspond to features of interest in the image under study, and recognising and extracting them can be a useful technique.

Some ideas based on greyscale topology, such as local minima and watersheds have already been discussed in §4.3. We will not pursue them further here, except to note that morphology often provides an effective method of describing topological features, and finding them in images. Further details may be found in Serra (1982, Ch.XII) and Meyer (1992). Some efficient (although complex) algorithms are discussed in Bleau, Deguise and Leblanc (1992). Useful discussions may also be found in Haralick, Watson and Laffey (1983) and Haralick (1983).

5.6 Summary

The main points of this chapter are:

- Mathematical morphology, which is based on set theory, provides useful tools for image operations.
- Morphology is complementary to the other methods described in this book.
- Morphology is most often used with binary images, but can be adapted to greyscale images.
- A fundamental operation is that of erosion, which looks at where a test set can fit into the image.
- From the erosion we can also define

  - dilation, which is an erosion on the complement of a set.
  - opening, an erosion followed by a dilation.
  - closing, a dilation followed by an erosion.

- The hit-or-miss transform is the most general morphological operation, and is based on a structuring element with two components. The transform involves looking for pixel positions where one component lies within the set of black pixels, and the other lies completely without.

- Operations based on square or circular structuring elements are the most commonly used. They interact with the size of features in the image.

- A distance transform can provide an efficient way to perform morphological operations, and can also be of interest in its own right.

- Morphological thinnings and skeletons preserve the topology of objects while reducing them to single pixel thickness.

- Texture may be studied using structuring elements consisting of pixels separated by a distance.

- Morphology may be extended to greyscale images by considering them as binary images in three dimensions.

- A variety of three-dimensional structuring elements may be used, and erosion, dilations, openings and closings may be defined. They are used to smooth an image, or to help extract particular structures.

- The top-hat transform enables us to extract small or thin, bright or dark objects from a varying background.