

IMAGE ANALYSIS FOR THE BIOLOGICAL SCIENCES

C A GLASBEY and G W HORGAN

Preface

Quantitative image analysis is the extraction of information from data which are in the form of pictures. The aim of this book is to cover the basics of image analysis from a statistical perspective, and with emphasis on methods relevant to the biological sciences. The book is written for both biological scientists and applied statisticians, whom, we assume, are mainly interested in analysing images of static objects in laboratory-type situations. No greater knowledge of mathematics, statistics or computing is necessary than would be gained in a first degree in a biological subject.

Image analysis methods are presented in five self-contained chapters. Detailed algorithms are given for the most important techniques, for the reader who wants to write his/her own computer program. Otherwise, the understanding gained from reading this book can guide the scientist in making intelligent use of the many computer packages on the market.

Fourteen types of image, drawn from a range of applications in microscopy, medical scanning systems, remote sensing, electrophoresis and photography, are used to motivate and illustrate the methods considered. All images, except the Landsat one, are available by anonymous FTP (file transfer) from Internet site peipa.essex.ac.uk. In case of difficulties, contact the publishers.

We are grateful to our collaborators for permission to use their data. They are:

Algal cells	Dr N.J. Martin	Scottish Agricultural College
Cashmere fibres	Dr A.J.F. Russel	Macaulay Land Use Research Institute
DNA sequencing gel	Dr F.G. Wright	Scottish Agricultural Statistics Service
Electrophoretograms	Prof D. Walliker	University of Edinburgh
Fish	Dr N.J.C. Strachan	Torry Research Station
Fungal hyphae	Dr K. Ritz	Scottish Crop Research Institute
Landsat	Dr C.H. Osman	Macaulay Land Use Research Institute
Magnetic resonance images	Dr M.A. Foster	University of Aberdeen
Muscle fibres	Dr C.A. Maltin	Rowett Research Institute
Synthetic aperture radar	Dr E. Attema	European Space Agency
Soil aggregate	Dr J.F. Darbyshire	Macaulay Land Use Research Institute
Turbinate bones	Dr J.F. Robertson	Scottish Agricultural College
Ultrasound image and X-ray CT	Dr G. Simm	Scottish Agricultural College

The Landsat image is ©National Remote Sensing Centre Ltd, Farnborough, Hampshire. We also thank Jim Young, Department of Geography, Edinburgh University, for initially making the data available to us.

We are indebted to our friends and colleagues who commented on earlier versions of the book. We gratefully acknowledge the advice and encouragement of Ian Craw, John Darbyshire, David Elston, Paul Fowler, Sheila Glasbey, Alison Gray, David Hitchcock, John Marchant, Geoff Simm, Norval Strachan, and Mark Berman, along with his colleagues Leanne Bischof, Ed

Breen, Yuchong Jiang, Guy Peden and Changming Sun in the CSIRO Division of Mathematics and Statistics.

We also thank Vic Barnett, the Series editor, for the challenge and guidance in writing this book, and Helen Ramsey and David Ireland of John Wiley and Sons for their technical help. Most of the figures in this book were produced using Zimage, written by the CSIRO Division of Mathematics and Statistics, Sydney.

Finally, we appreciate the backing of Rob Kempton, Director of the Scottish Agricultural Statistics Service, and acknowledge the financial support of the Scottish Office Agriculture and Fisheries Department

Chris Glasbey,
Graham Horgan,
Edinburgh,
February 1994.

Chapter 1

INTRODUCTION

Image analysis is simply *the extraction of information from pictures*. For example, it's what you are doing in reading these words. Sight is (arguably) the most important sense we have. It is said that 'a picture is worth a thousand words' (a view shared by Charlie Brown's little sister, Sally — see Fig 1.1). This book is about images which arise in the biological sciences, and analysing them using a computer. In the three following sections we will address the three questions:

- Why use a computer to analyse images?
- What are the data to be analysed?
- What does image analysis consist of?

Finally, in §1.4 we will summarize the points covered.

1.1 Why use a computer?

If we are so effective at seeing, then why try to use a computer to analyse images? Put simply, the answer is that some image analysis is more easily done by the human eye, but for other

THIS FIGURE CANNOT INCLUDED BECAUSE OF COPYRIGHT
IT SHOWS SALLY SAYING THAT HER HOMEWORK ASSIGNMENT WAS
A 2000 WORD REPORT, BUT A PICTURE IS WORTH 1000 WORDS,
SO SHE HAS DRAWN TWO PICTURES!

Figure 1.1: Peanuts cartoon.

tasks a computer is better. To illustrate the point let us consider a problem which arose in a veterinary investigation.

1.1.1 An example of image analysis

To study atrophic rhinitis, a disease of the upper respiratory tract in pigs, snouts of dead pigs were cut in half (Robertson, Wilson and Smith, 1990). The exposed cross-sections were then inked and printed on paper. Fig 1.2(a) shows a print from a disease-free pig. The disease affects the turbinate bones, the curled structures which can be seen, printed as black, inside each of the nasal cavities. Done, Upcott, Frewin and Hebert (1984) proposed a morphometric index to measure deterioration in these bones. It is calculated as the ratio of air-space-area in the cross-section of the nasal cavities to air-space-area plus turbinate-bone-area. In order to measure these areas, Robertson *et al* enhanced the prints laboriously by hand, using a black pen to fill in areas of bone which had not been printed clearly and typists' correction fluid to whiten places where ink had got by mistake. Fig 1.2(b) shows the result obtained from Fig 1.2(a). The pictures were then converted into digital form by scanning them into a simple computer-based image analysis system. (See the Appendix for more details about computer equipment.) Areas of cavity and turbinate bone were measured, following a further manual operation using a computer mouse to outline the appropriate regions. The whole process had to be repeated on 1200 snouts, which proved to be an exceedingly tedious task and took 300 hours.

This combination of human and computer interpretation is known as **semi-automatic** image analysis. It would be very difficult, maybe even impossible, to program a computer to achieve the enhancement shown in Fig 1.2(b). It would also be very laborious for a scientist to accurately measure the areas. However, it might have been possible to have made more use of the computer in image enhancement, and thereby reduce the human input. (Later in the book, in Chapter 5, we will consider some other methods for cleaning-up the turbinate image.)

Let us return to a general consideration of the relative benefits of human and computer interpretation of images.

1.1.2 Human and computer interpretation of images

The human vision system is superb, particularly at qualitative tasks. For example, look at Fig 1.3. You may need to half-close your eyes, or hold the book at arms length, to recognise the picture. It is remarkable that we can see so much even though the picture is so coarse. (If you don't recognize the picture, we will tell you in two paragraphs time.)

Our eyes can however be fooled. In Fig 1.4(a), both central circles are the same size, although the right-hand one looks bigger. In Fig 1.4(b), the central squares are the same shade of grey, although the right-hand one looks lighter. Therefore, for extracting quantitative information from images, a computer may be better: it may be more consistent from day to day and not




Figure 1.2: Turbinate image: (a) as printed, (b) after manual enhancement.

rely on the opinions of different observers. Furthermore, computers may spare us from much tedious image interpretation such as that already discussed for the turbinate images.

Most of the time we see effortlessly. When scientists first started programming computers to interpret images they expected to make rapid progress. It has proved, however, to be a very difficult task. Perhaps this is because we are not conscious of what processes we go through in looking at, for example, Fig 1.3. To a computer, an image is no more than a matrix of numbers. Look at Table 1.1, which is the numerical version of Fig 1.3, a coarse-resolution approximation to Leonardo da Vinci's painting of The Mona Lisa. Table 1.1 is a 29×18 array of numbers, ranging in value from 0 to 255, while Fig 1.3 is a 29×18 array of square blocks of varying shades of grey. The values in Table 1.1 are matched to the grey levels in Fig 1.3, in that 0 is displayed as black, 255 as white, and in general the higher the value the lighter the shading of the corresponding block. Fig 1.5 gives an alternative graphical representation of these numbers,




Figure 1.3: Example of image with coarse resolution. What is it?




Figure 1.4: Two optical illusions: (a) Which central black disc is the larger? (b) Which central grey square is the darker?

where height is used to represent numerical value. In either case, in looking at Table 1.1 or at Fig 1.5, we are no better off than the computer — we also have great difficulty in recognising The Mona Lisa.

Biological objects tend to be more irregular and variable in shape than man-made ones. Therefore they present an even greater challenge to fully-automatic interpretation. For most practical applications in the biological sciences, the best that can be achieved at present is a semi-automatic system, with the computer reducing the tedious aspects of human image interpretation and carrying out the more quantitative tasks. Fortunately, we are not trying to build a robot or some other autonomous system. If a computer has to recognise an airport runway in order to land a plane, then it had better be right *every single time!* In applications involving medical diagnosis, it is usually better for a computer screening system to err on the cautious side, and flag the likely presence of a tumour (for example) if in doubt, in preference to ever missing a real tumour. In the applications we are going to consider, occasional mistakes can be tolerated. In a semi-automatic system they can either be corrected manually or they can be allowed to remain, in which case they will contribute to the overall uncertainty in the results.

So, what are these applications?





Figure 1.5: Alternative display of Fig 1.3, with grey level replaced by height.

1.1.3 Further examples


In the Scottish Agricultural Statistics Service, which provides statistical and mathematical support to agricultural and environmental research scientists, we encounter many applications of image analysis. The objects which are imaged range in scale from the microscopic to satellite views of the earth. Here we will introduce some of these examples, and the questions that the scientists wish to see answered. Throughout the book we will use these images and questions to motivate the methods we consider, and to illustrate the results.

Fig 1.6 shows four **microscope images**. Such digital images would normally be produced by mounting a digitizing camera (see Appendix) on top of a microscope. Those shown here were produced less directly by taking photographs and using a desktop scanner to digitize them (Glasbey, Horgan and Hitchcock, 1994).


- **Algal cells:** Fig 1.6(a) is a differential interference contrast (DIC) microscope image of some algal cells, collected as part of a research programme to manage algal ponds for waste treatment (Martin and Fallowfield, 1989). The aim is to identify, count and measure cells in such images. Note that the effect of DIC is to make one side of each cell light and




(a)



(b)



(c)



(d)


Figure 1.6: Microscope images: (a) algal cells, (b) muscle fibres, (c) soil aggregate, (d) cashmere fibres.

the other side dark, producing an effect similar to the illumination of 3D objects. DIC microscopy operates by splitting a beam of light, directing one half through the specimen while the other half bypasses it, then combining them back together. Where the two beams of light are in phase, after combination they produce a bright area in the image, whereas when they are out of phase they cancel out and produce a dark area (Holmes and Levy, 1987 and 1988). DIC is particularly effective in viewing specimens that are almost transparent, because such specimens still change the phase of light passing through them.


- **Muscle fibres:** Fig 1.6(b) is a section through one of the muscles in a rat's leg, the soleus muscle. The transverse section has been stained to demonstrate the activity of Ca^{2+} -activated myofibrillar ATPase and allows one to classify three types of fibre: fast-twitch oxidative glycolytic (dark), slow-twitch oxidative (light) and fast-twitch glycolytic (mid-grey) (Maltin, Hay, Delday, Loblely and Reeds, 1989). Information of the numbers and sizes of the fibres are required for research into clenbuterol, a drug which enhances muscle development.
- **Soil aggregate:** Fig 1.6(c) is a montage of backscattered electron scanning micrographs of a soil aggregate embedded in acrylic resin (Darbyshire, Griffiths, Davidson and McHardy, 1989; Glasbey, Horgan and Darbyshire, 1991). The black areas are soil pores and the lighter areas are the inorganic and organic soil matrix. Note that the image is made more difficult to interpret by the variation in brightness between micrographs in the montage and by the visible edges of some prints. The objective was to study porosity and pore-size distribution within a sample of soil aggregates and to relate these characteristics to microbial activity within and outside the aggregates.
- **Cashmere fibres:** Fig 1.6(d) is a back-illuminated image of cashmere goat fibres. The aim was to measure their diameters automatically, in support of a goat breeding programme (Russel, 1991). Measurement is made more difficult by some fibres being out-of-focus, which gives either dark or light edges to the fibres, so-called 'Becke lines'.

Fig 1.7 shows four digital images produced by **medical scanning systems**:


- **Magnetic resonance imaging (MRI):** Figs 1.7(a) and (b) are inversion recovery and proton density images through a supine woman's chest. The woman has a cubic test object between her breasts. Magnetic resonance imaging is a non-invasive technique for viewing sections through living tissue by means of large, oscillating magnetic fields. These images were obtained as part of a study of changes in breast volume during the menstrual cycle (Fowler, Casey, Cameron, Foster and Knight, 1990).
- **X-ray computed tomography (CT):** Fig 1.7(c) is a cross-section through the thorax of a live sheep, obtained in order to estimate the quantity of fat and lean tissue (Simm, 1992). In CT, X-rays are projected through a subject from different directions, and a computer reconstructs an image of the distribution of tissue types from the transmitted X-rays. As is the convention with X-ray plates, light areas in the image denote regions which transmitted less X-rays. The lightest areas are the backbone and the parts of the ribs which intersect the imaging plane. The muscles and internal organs appear slightly



(a)



(b)



(c)



(d)

Figure 1.7: Medical scanning images: (a) MRI inversion recovery image, (b) MRI proton density image, (c) X-ray CT image, (d) ultrasound image.

lighter than the fat tissue because they are slightly more opaque to X-rays. The U-shaped plastic cradle in which the sheep was lying can also be seen. X-ray attenuation is measured in Hounsfield units, which range between -1000 and about 1000 .

- **Ultrasound imaging:** Fig 1.7(d) is an ultrasound image of a cross-section through a sheep's back. The instrument operates by using a transducer to send a pulse of sound waves of very high frequency into a subject. When the ultrasound wave meets a boundary between two tissues, partial reflection occurs. The reflected energy is received by the transducer and converted into electrical signals which are displayed on a video monitor, with time delay interpreted as depth. Unfortunately, in this application it proved impossible to extract the data directly from the ultrasound machine because there was no video output. Therefore the video display was photographed and redigitized, resulting in the vertical display lines which can be seen in Fig 1.7(d). Images were collected for the same reason as the CT image above, namely to estimate sheep body composition (Simm, 1992). The top-most approximately horizontal white line is the transducer-skin boundary, below which are the skin-fat and fat-muscle boundaries. The backbone is on the bottom left, from which a rib can be seen sloping slightly upwards. Ultrasound images are far less clear than X-ray images, but have the advantages that the machines are safer, cheaper and more portable.

Fig 1.8 shows some **remotely-sensed images**:

- **Landsat Thematic Mapper (TM) images:** Figs 1.8(a)-(f) are TM bands 1-5 and 7, for a region between the river Tay and the town of St Andrews on the east coast of Scotland, in May 1987. Spatial resolution is 30m, and the bands correspond to blue ($0.45\text{-}0.52\mu\text{m}$), green ($0.52\text{-}0.60\mu\text{m}$), red ($0.63\text{-}0.69\mu\text{m}$) and three near-infrared ($0.76\text{-}0.90$, $1.55\text{-}1.75$ and $2.08\text{-}2.35\mu\text{m}$) wavelength regions of light reflected from the earth's surface. Many features can be identified. For example, the dark areas at the top and right of Fig 1.8(f) are water, the bright areas to the centre and left of Fig 1.8(b) are fields of oil-seed rape, and the section of land in the top-left corner is part of the city of Dundee.
- **Synthetic aperture radar (SAR) image:** Fig 1.8(g) is a C-band, HH-polarization, SAR image of an area near Thetford forest, England, in August 1989. It was obtained by plane as part of the Maestro-1 campaign (Joint Research Centre, Ispra, report IRSA/MWT/4.90). The resolution is 12m, formed by averaging four adjacent values in the original $3\text{m} \times 12\text{m}$ data. Notice that the intensity values in this image are locally much more variable than in the Landsat images. The apparently random fluctuations are called 'speckle'. Horgan (1994) found good agreement between the observed 'speckle variance' and a theoretically expected value of 0.41, after the data have been log-transformed (see §2.2). Unlike Landsat, SAR is an active sensing system: microwave radiation is beamed down to the earth's surface, a sensor detects the reflected signal, and from this an image is constructed. This sensing system can operate during day or night and is unaffected by clouds unlike Landsat images.

Finally, Fig 1.9 shows five images obtained directly by digitizing views of the objects themselves:







Figure 1.8: Remotely-sensed images: (a) Landsat TM band 1, (b) band 2, (c) band 3, (d) band 4, (e) band 5, (f) band 7, (g) SAR image.




(a)



(b)



(c)



(d)



(e)

Figure 1.9: Further illustrative images: (a) 2-D electrophoretogram of strain of malaria parasite, (b) electrophoretogram of second strain of malaria parasite, (c) DNA sequencing gel autoradiograph, (d) fungal hyphae, (e) a fish.

- **Electrophoretograms:** Figs 1.9(a) and (b) are digitized autoradiographs of SDS-PAGE gel electrophoretograms of two strains of the malaria parasite *Plasmodium falciparum*. Each spot on a gel represents a different protein. Relative locations of spots, identified by making comparisons between gels, are used to identify the malarial strain (Horgan, Creasey and Fenton, 1992). Initially, a mixture of radioactively-labelled proteins was positioned in the top left-hand corner of each gel, from where the individual proteins migrated at different speeds across and down the gel. Then, a photographic plate was placed over the gel. This blackened at protein locations in response to radioactive emissions, thus producing an autoradiograph. Finally, the autoradiograph was digitized using a desktop scanner.
- **DNA sequencing gel autoradiograph:** Fig 1.9(c) is another type of electrophoretogram, produced at one stage in the DNA-sequencing of gene fragments. In this case, about fifty mixtures were positioned as distinct spots along the top of the gel (as it is currently displayed). Each mixture then migrated down the gel, and DNA fragments produced separate, approximately horizontal bands.
- **Fungal hyphae:** Fig 1.9(d) is a digitized photograph of part of a fungal mycelium *Trichoderma viride*, that it a network of hyphae from a single fungal organism, which was grown on cellophane-coated nutrient agar (Ritz and Crawford, 1990). Image analysis was required here to understand the spatial structure of the fungal hyphae in relation to their environment.
- **Fish:** Fig 1.9(e) is a digitized view of a haddock. The data were obtained by Strachan, Nesvadba and Allen (1990a) as part of their work to automate the identification of fish species for MAFF (Ministry of Agriculture, Fisheries and Food) surveys and for the fish industry. Summary statistics on size and shape need to be extracted from the image, in order to make comparisons between different fish species.

The names we have given some images correspond to their subject matter, whereas in other cases we have named them after the instrument used to gather the data. The choice is somewhat arbitrary, but follows convention, and will be maintained throughout the book.

1.2 What are the data?

In this section we will draw some general conclusions about the examples presented in §1.1.3, in order to answer the question: what are the data?

The array of values of *any* variate measured, either directly or indirectly, at regular points on a two-dimensional grid or lattice may be regarded as a computer image, and analysed using the methods to be presented in this book. Let us first illustrate what we mean by a variate, before considering what the lattice is.

The **variate** may be a measure of the intensity of transmitted light, as in the examples of fungal hyphae, muscle fibres and cashmere fibres. It may be a measure of reflected light, as in

the examples of fish and Landsat. It could also depend upon reflected or transmitted radiation in another part of the electromagnetic spectrum (soil aggregate, SAR and X-ray CT), or it could measure emitted radiation after radioactive labelling (electrophoretograms), or reflected ultrasound. It may be measurements of proton density (MRI), or of interference microscopy (algal cells). It could also be measurements of ink uptake and deposition on paper (turbinate bones), or of height, colour, texture, distance, etc.

If the object is *three dimensional*, then an image may be obtained after physically taking a cross-section (as in the muscle fibres, soil aggregate and turbinate bone examples), or by a computer reconstruction based on some physical property of the object (MRI, X-ray CT and ultrasound). Alternatively, an object could be imaged simply by viewing it from a particular direction. The object could have an opaque surface (fish, Landsat, SAR) or be semi-transparent (algal cells, cashmere fibres). In both cases there may be a focal plane outside which parts of an object are blurred. Some sensors, such as confocal microscopes and magnetic resonance imagers, can collect three-dimensional arrays of data. Such datasets are extremely large and require powerful computers to handle them. They are beyond the scope of this book but can be analysed using similar methods.

We will now consider what we mean by a **lattice**. Although most variates are measurable at all points within a defined area, a variate can only be recorded at a finite number of them. It simplifies matters if these points are arranged systematically in some form of lattice. Lattices are usually square, although they can be rectangular (as in the ultrasound image) or hexagonal (Serra, 1982) or the points can even be arranged in concentric rings (Silverman, Jones, Wilson and Nychka, 1990). Each data value, or picture element (**pixel** for short), is the average value of the variate for a small region around a point in the lattice. For example, Fig 1.10(a) shows a region in the bottom centre of the DNA sequencing gel autoradiograph, whereas Fig 1.10(b) shows a display of the digitized image produced by a desktop scanner. (This is the last we will see of any source image — from now on the DNA image and all the other images considered in this book are digital.) Typically, averaging extends further than the distance between pixels, and images appear slightly blurred when viewed at high resolution. This is the case for all the images in Figs 1.6–9. However, for some sensors the extent of averaging is less than the interpixel distance and details of objects can be lost in the gaps between pixels (Purll, 1985). For example, Glasbey, Horgan and Hitchcock (1994) found this to be the case if their desktop scanner was used to sample photographs at too coarse a resolution.

Although we have said that any variate measured on a lattice *can* be analysed using image analysis techniques, not every variate should be. Plots in a field experiment are often arranged spatially in a two-dimensional grid, so the crop yields of the plots could be regarded as an image. However, it is far more sensible to use statistical methods such as analysis of variance to interpret these data. Also, variates which vary smoothly over space, barometric pressure for example, do not lend themselves to image analysis. All the examples in Figs 1.6-1.9 show abrupt changes in pixel value, which are interpreted by our eyes as boundaries or as the edges of objects. These are characteristics which variates will tend to have if they are to be successfully




Figure 1.10: Detail from bottom centre of DNA image: **(a)** original autoradiograph (MISSING), **(b)** digitized version.

analysed using the methods in this book.

Throughout the book we will consider an image to be a two-dimensional array of numbers and will denote pixel location by (i, j) , where i is the **row index**, an integer ranging from 1 to n , and j is the **column index**, similarly between 1 and n . The number of rows may be different from the number of columns, resulting in images which are rectangular rather than square (in fact some of our examples are) but to keep the mathematics as simple as possible we will use

a single dimension, n . The reader should find it relatively simple to generalize to rectangular images. We will call the value of the variate (or variates) at a pixel the **pixel value** and denote it by f_{ij} . By convention, indices begin in the top-left corner of images so that, for example, in the digitized image of Mona Lisa, $f_{11} = 132$, $f_{12} = 128$ and $f_{21} = 139$ (see Table 1.1). We will regard pixels as *discrete points* in the imaging plane, although we will display them as *square blocks*. Fig 1.11 summarises the notation.

Pixel values are **univariate** if only one measurement is made at each pixel, otherwise they are **multivariate**. If univariate, then they may simply take values of 0 and 1, in which case the image is **binary** (the turbinate image, for example), or they may be multilevel, termed **greyscale** (for example, all the individual images in Figs 1.6–1.9). Because images typically contain a large number of pixels, it is convenient to keep the storage space as small as possible for each pixel, so values are often rounded to integers in the range 0 to 255 — which occupy one byte in a computer. (This is the case in Table 1.1.)

The sizes of the images in Figs 1.6–1.9, and details of pixel and display ranges, are summarized in Table 1.2. In most instances, we displayed the smallest pixel value as black and the largest value as white, with shades of grey used to represent intermediate pixel values. However, where this produced a display which was too dark for much detail to be discerned, the largest 1% of pixel values were all displayed as white and a contrast stretch was used to increase the brightness of all other pixels (see §2.2.2).

Multivariate images can be categorized into several types.

Multispectral: the variates measure intensity in different parts of the electromagnetic spectrum (as in the Landsat example and in the use of red, green and blue light in colour images).

Multimodal: the variates measure different physical properties of the same object (MRI inversion recovery and proton density for example).

Multitemporal: a single variate is measured at different times, possibly seconds or years apart.

It is important in any practical application of image analysis to consider whether, by collecting data in a different way, image interpretation could be made simpler. For example, variations in illumination can sometimes be dealt with by collecting two images, one with the specimen present and one of the background alone, then forming a pixel by pixel ratio. In microscopy, samples could be stained or viewed using a different microscope modality or different lighting conditions. Strachan *et al* (1990a) had an elegant solution to the problem of distinguishing between the fish and its background in Fig 1.9(e), a task which looks simple to the human eye but is difficult to program a computer to achieve. They collected two images: one was obtained with illumination from above (Fig 1.9(e)); and the other one with back-lighting only, which showed the fish as a dark object on a white background. By combining the two images they were able both to distinguish between the fish and its background and also to record surface details of the fish.




Figure 1.11: Mathematical notation for digital image.

In this book we assume that the reader is mainly interested in analysing images of static objects in laboratory-type situations where analysis time is not critical, samples can be prepared to some extent and processing can be semi-automatic. Real-time, fully-automatic image processing of biological objects is also of importance, for example in grading of agricultural produce, but is beyond the scope of this book. Topics in computer vision, such as motion, 3D modelling, real-time processing and high-level image interpretation in the domain of artificial intelligence will not be dealt with. Nor will we consider the mathematics of image reconstruction from projections, such as that used to produce the X-ray CT image. The reader is referred to texts such as Rosenfeld and Kak (1982), Gonzalez and Wintz (1987) and Jain (1989) for coverage of these topics.

In the following section we will give an overview of how computers can be used to analyse the types of images we are considering.

1.3 What does image analysis consist of?

One way of describing image analysis is to recognize five distinct stages that follow each other logically, namely: display, filters, segmentation, mathematical morphology and measurement. To illustrate the stages, consider the algal image (Fig 1.6(a)).

Display of an array of pixel values as a picture on a computer screen or on paper, as we have already seen in this chapter, is the first stage in analysing a digital image. Magnification of pixels is also useful for seeing details in an image. Fig 1.12(a) shows two cells from the algal image in close-up.

Filters enhance images by applying transformations based on groups of pixels. They are computationally efficient methods for reducing ‘noise’ levels in images and emphasising edges. (**Noise** is a technical term used by statisticians and engineers, among others, to describe disturbances in data which are either uninterpretable or not of interest.) The moving-median filter smooths flat regions of images, but preserves edges, by replacing each pixel value by the median of the values in a specified local region. Fig 1.12(b) shows the result of applying a 5×5 median filter to the same section of the algal image as Fig 1.12(a). By comparing Fig 1.12(a) with (b) we can see that the noise which caused the speckle has been reduced.

Segmentation divides an image up into regions, which correspond to different objects or parts of objects, by classifying all pixels. In the algal image, the aim is to identify which aggregations of pixels belong to which cells. Fig 1.12(c) shows the distribution of pixel values in the median-filtered image, displayed on a square-root scale. Background pixels (those not belonging to cells) contribute the bulk of this distribution, with values in the range 130 to 160. By thresholding at 120, that is labelling all pixels with values less than 120 as belonging to cells, dark parts of the algal cells are distinguished from the background, as shown in Fig 1.12(d).

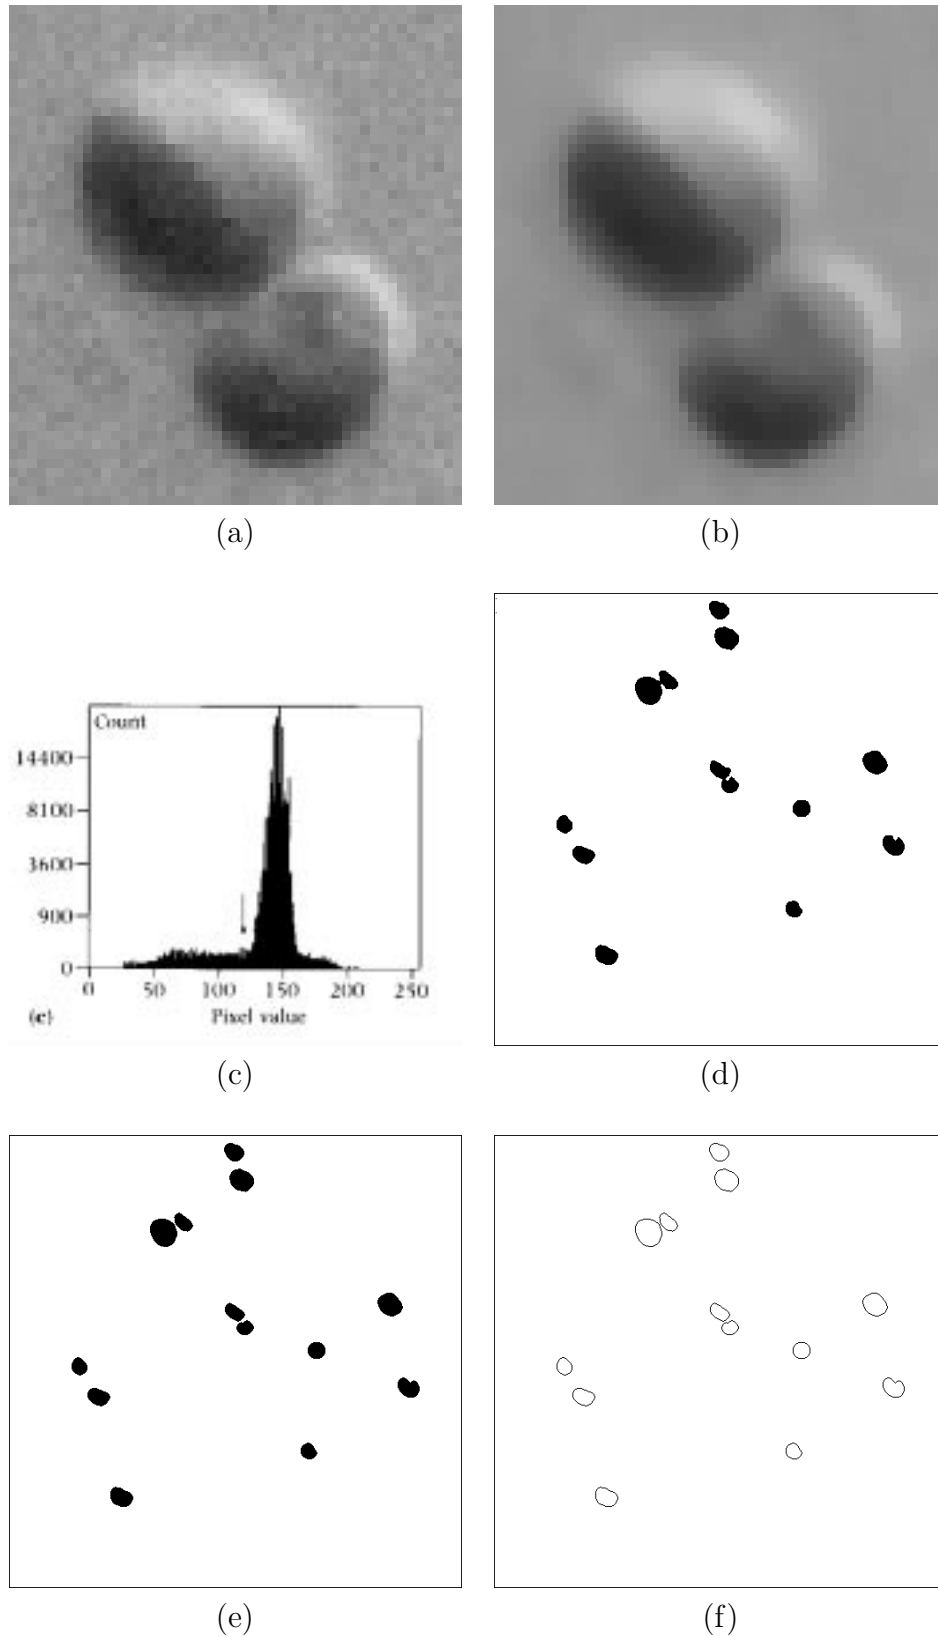


Figure 1.12: Stages in image analysis illustrated using algal image: (a) detail from image; (b) same detail after application of 5×5 moving median filter; (c) histogram, on a square-root scale, of pixel values after filtering, with an arrow to indicate the threshold; (d) result of thresholding image at pixel value 120, to produce a binary image; (e) result of applying morphological opening to (d); (f) separated objects in (e) counted.

Mathematical morphology uses a collection of operations to study the shapes of objects, as expounded mainly by Serra (1982). For example, two pairs of cells are touching in Fig 1.12(d), and boundaries are slightly roughened due to noise, even though the algal cells themselves should be smooth. Fig 1.12(e) shows the result of applying a morphological operation called an opening. Now all the cells are separated and have smooth boundaries.

Measurement, the extraction of quantitative information from images, is the final stage in the analysis. The number of separate objects can be counted, as shown in Fig 1.12(f). Although the algal cells are approximately circular in shape, these objects are more ellipsoidal because only the darker sides of the cells have been identified. By measuring the maximum distances between pixels in each object, the cell diameter can be estimated. Cell 4 appears to be largest, with a diameter of 33.6 pixel units, whilst 12 appears smallest, diameter 19.0.

Analysis of any particular image is likely to require several of these stages, in this order, but sometimes re-using techniques from previous stages. For example, after filtering it is usual to display the transformed version of the image. Of course, the way in which data were collected (§1.2), and the questions to be answered, are of crucial importance in determining how a particular image should be analysed.

There are other ways of approaching image analysis. One alternative is where data are processed in a single step to produce the final result. This is the approach taken by Grenander, Chow and Keenan (1991) in fitting template models to images of hands.

In the five chapters which follow we will deal in greater detail with the five stages identified above. Each chapter is reasonably self-contained, so the reader need not work through the whole book in order. Finally, in chapter 7 we will review the progress we have made in analysing each of the 14 datasets, and give pointers to further work.

1.4 Summary

The key points of this chapter are:

- Image analysis is simply the extraction of information from pictures.
- The human vision system is superb, but qualitative.
- Computers can do better than humans at extracting quantitative information and can reduce the tedious aspects of image interpretation.
- The images to be considered in this book are drawn from microscopy, medical scanning systems, direct photography and remote sensing.
- Any variate measured at regular points on a two-dimensional grid may, after digitizing, be regarded as a computer image.

- The pixel value at location (i, j) will be denoted by f_{ij} . For simplicity, images will be assumed to be square, with both the row index (i) and the column index (j) ranging between 1 and n .
- Images may be univariate or multivariate, and each variate may be binary or greyscale.
- Image analysis can be subdivided into five stages: display, filters, segmentation, mathematical morphology and measurement. These are the topics of the following five chapters.

132	128	126	123	137	129	130	145	158	170	172	161	153	158	162	172	159	152
139	136	127	125	129	134	143	147	150	146	157	157	158	166	171	163	154	144
144	135	125	119	124	135	121	62	29	16	20	47	89	151	162	158	152	137
146	132	125	125	132	89	17	19	11	8	6	9	17	38	134	164	155	143
142	130	124	130	119	15	46	82	54	25	6	6	11	17	33	155	173	156
134	132	138	148	47	92	208	227	181	111	33	9	6	14	16	70	180	178
151	139	158	117	22	162	242	248	225	153	62	19	8	8	11	13	159	152
153	135	157	46	39	174	207	210	205	136	89	52	17	7	6	6	70	108
167	168	128	17	63	169	196	211	168	137	121	88	21	9	7	5	34	57
166	170	93	16	34	63	77	140	28	48	31	25	17	10	9	8	22	36
136	111	83	15	48	69	57	124	55	86	52	112	34	11	9	6	15	30
49	39	46	11	83	174	150	128	103	199	194	108	23	12	12	10	14	34
26	24	18	14	53	175	153	134	98	172	146	59	13	14	13	12	12	46
21	16	11	14	21	110	126	47	62	142	85	33	10	13	13	11	11	15
17	14	10	11	11	69	102	42	39	74	71	28	9	13	12	12	11	18
18	19	11	12	8	43	126	69	49	77	46	17	7	14	12	11	12	19
24	30	17	11	12	6	73	165	79	37	15	12	10	12	13	10	10	16
24	40	18	9	9	2	2	23	16	10	9	10	10	11	9	8	6	10
43	40	25	6	10	2	0	6	20	8	10	16	18	10	4	3	5	7
39	34	23	5	7	3	2	6	77	39	25	31	36	11	2	2	5	2
17	16	9	4	6	5	6	36	85	82	68	75	72	27	5	7	8	0
4	8	5	6	8	15	65	127	135	108	120	131	101	47	6	11	7	4
2	9	6	6	7	74	144	170	175	149	162	153	110	48	11	12	3	5
11	9	3	7	21	127	176	190	169	166	182	158	118	44	10	11	2	5
8	0	5	23	63	162	185	191	186	181	188	156	117	38	11	12	25	33
3	5	6	64	147	182	173	190	221	212	205	181	110	33	19	42	57	50
5	3	7	45	160	190	149	200	253	255	239	210	115	46	30	25	9	5
9	4	10	16	24	63	93	187	223	237	209	124	36	17	4	3	2	1
7	8	13	8	9	12	17	19	26	41	42	24	11	5	0	1	7	4

Table 1.1: Data in Fig 1.3. Each number represents the brightness in the display of a single box in Fig 1.3, ranging from 0 for black to 255 for white.

Image	Number of		Pixel range		Displayed range	
	rows	columns	min	max	black	white
Algal cells	512	512	23	236	23	236
Cashmere fibres	256	256	35	219	35	219
DNA sequencing gel	512	300	29	226	29	226
Electrophoretogram 1	370	400	7	249	7	249
Electrophoretogram 2	370	400	10	213	10	213
Fish	200	500	23	255	23	255
Fungal hyphae	500	500	11	244	11	244
Landsat band1	512	512	61	254	61	100
Landsat band2	512	512	21	134	21	55
Landsat band3	512	512	17	165	17	60
Landsat band4	512	512	8	198	8	140
Landsat band5	512	512	2	254	2	105
Landsat band7	512	512	1	215	1	60
MRI inversion recovery	128	128	0	253	0	130
MRI proton density	128	128	0	253	0	253
Muscle fibres	512	512	7	218	7	218
SAR *	250	250	0.3	70.1	0.3	10
Soil aggregate	512	512	3	164	3	80
Turbinate bones	530	460	0	1	0	1
Ultrasound	300	360	8	239	8	239
X-ray CT *	256	256	-1000	834	-1000	834

Table 1.2: Data sets and displayed ranges of values in Figs 1.6-1.9. * In uses of these images after chapter 2, the SAR data are log-transformed and the X-ray CT data are restricted to the pixel range -250 to 260. All images, except for Landsat, are available by anonymous FTP from Internet site peipa.essex.ac.uk. In case of difficulties, contact the publisher.