

Chapter 1

Introduction to Texture Analysis

E. R. Davies

*Machine Vision Group, Department of Physics
Royal Holloway, University of London
Egham, Surrey, TW20 0EX, UK
e.r.davies@rhul.ac.uk*

Textures are characteristic intensity (or colour) variations that typically originate from roughness of object surfaces. For a well-defined texture, intensity variations will normally exhibit both regularity and randomness, and for this reason texture analysis requires careful design of statistical measures. While there are certain quite commonly used approaches to texture analysis, much depends on the actual intensity variations, and methods are still being developed for ever more accurately modelling, classifying and segmenting textures. This introductory chapter explores and reviews some fundamental techniques.

1.1. Introduction — The Idea of a Texture

Most people understand that the human eye is a remarkable instrument and value highly the gift of sight. However, because the human vision system (HVS) permits scene interpretation ‘at a glance’, the layman has little appreciation of the amount of processing involved in vision. Indeed, it is largely the case that only those working on the brain—or those trying to emulate its capabilities in areas such as machine vision—have any idea of the underlying processes. In particular, the human eye ‘sees’ not scenes but sets of objects in various relations to each other, in spite of the fact that the ambient illumination is likely to vary from one object to another—and over the various surfaces of each object—and in spite of the fact that there will be secondary illumination from one object to another.

In spite of these complications, it is sometimes reasonable to make the assumption that objects, or object surfaces, can be segmented from each

other according to the degree of uniformity of the light reflected from them. Clearly, this is only possible if a surface is homogeneous and has uniform reflectivity, and is subject to uniform illumination. In that case not only will the intensity of the reflected light be constant, but its colour will also be unvarying.

In fact, it will rarely be the case that objects, or their surfaces, can be segmented in this way, as almost all surfaces have a texture that varies the reflectance locally, even if there is a global uniformity to it. One definition of texture is the property of the surface that gives rise to this local variability. In many cases this property arises because of surface roughness, which tends to scatter light randomly, thereby enhancing or reducing local reflectance in the viewing direction. Even white paper has this property to some extent, and eggshell more so. In many other cases, it is not so much surface roughness that causes this effect but surface structure—as for a woven material, which gives rise to a periodic variation in reflectance. There are other substances, such as wood, which may appear rough even if they are smooth, because of the grain of the intrinsic material, and their texture can vary from a fine to a coarse pattern. Ripples on water can appear in the form of a relatively coarse texture, albeit in this case it will have a rapidly varying temporal development. Other sorts of texture or textured appearance arise for sand on the seashore, or a grass lawn, or a hedge. In such cases the textured surface is a composite of grains or leaves, i.e. it is composed of separate objects, but for image interpretation purposes it is usual to regard the surface as having a unique texture. On the other hand, if the scale is altered, so that we see relatively few component objects, as for a pile of large chickpeas—and certainly for a pile of potatoes—the illusion of a texture evaporates. To some extent, then, a texture may be a fiction created by the HVS, and is not unrelated to the limited resolution available in the human eye. All these points are illustrated in Figs. 1.1 and 1.2.

In practice, a surface is taken to be textured if there is an uncountably large number of texture elements (or ‘texels’), and a set of objects if the opposite is true. In general, the components of a texture, the texels, are notional uniform micro-objects which are placed in an appropriate way to form any particular texture. The placing may be random, regular, directional and so on, and also there may be a degree of overlap in some cases—as in the case of a grass lawn. It is also possible to vary the sizes and shapes of the texels, but doing this reduces the essential simplicity of the concept. Actually, what we are seeing here is the possibility of recognition by reconstruction: if a texture can be reconstructed, it will almost certainly have



Fig. 1.1. A variety of textures. (a) Tarmac, (b) brick, (c) carpet, (d) cloth, (e) wood, (f) water. These textures demonstrate the wide variety of familiar textures that are easily recognised from their characteristic intensity patterns.

been interpreted correctly. However, while recognition by reconstruction is generally a sound idea, it is much more difficult with textures because of the random element. Nevertheless, the idea is of value when scene generation has to be performed in a realistic manner, as in flight simulators.

What we have come up with so far is the idea of textured surface ap-



Fig. 1.2. Textured surfaces as composites. In this figure the lentils in (a) and the slice of bread in (c) are shown in (b) and (d) with respective linear magnifications of 4 and 2. Such magnifications often appear to change the texture into a set of composite objects.

pearance, which can be imagined as due to appropriate placement of texels; this texture is a property by which the surface can be recognised; in addition, when different regions of an image have different textures, this can be used to segment objects or their surfaces from each other.

We have also seen that, whatever their source, textures may vary according to randomness, regularity (or periodicity), directionality and orientation. Notice that we are gradually moving away from texture as a property of the surface (the physical origin of the texture) to appearance in the image, as that is what concerns us in image texture analysis. Of course, once image textures have been identified, we can relate them back to the original scene and to the original object surfaces. This distinction is important, because it is sometimes the case (see Fig. 1.3) that 3D object structures can be discerned from information about texture variations.

At this point it is useful to say what is and what is not a texture. If an intensity variation appears to be perfectly periodic, it would normally



Fig. 1.3. Views of a grass lawn. View (a) is taken from directly above, while view (b) is taken at an angle of about 40° to the vertical. Notice how the texture gives useful information on the viewing angle.

be described as a ‘periodic pattern’ and would not be called a texture. Likewise, any completely random pattern would probably be called a ‘noise pattern’ rather than a texture—though this may be a subjective judgement, and might depend on scale or colour. However, if a pattern has both randomness and regularity, then this is probably what most people would call a texture, and is the definition that we adhere to here. In fact, these intellectual niceties will normally be largely irrelevant, as any algorithm designed for texture analysis will almost certainly be able to make some judgements about periodic patterns and about noise patterns. However, the inverse will not be the case: for example, an algorithm designed to discern periodic patterns may give inappropriate answers when presented with textures, because the randomness could partially cancel out the periodicity.

Another feature of a texture is its ‘busyness’: this applies whatever the degree of mix between randomness and regularity, and is not made substantially different if the texture is directional. To a large extent we can characterise textures as having busy microstructures but uniform macrostructures. We can even envisage identifying the busy components and then averaging them in some way to produce uniform measures of macrostructure. As we shall see below, this concept underlines many approaches to texture analysis—at least as a first approximation.

Overall, we have found in this introduction that texture offers another way to segment and recognise surfaces. It will not be useful when surfaces have constant reflectance, in which case the amount of reflected light and its colour will be the sole means by which the surface can be characterised. But when this is not so, texture adds significant further discriminatory

information by which to perform recognition and segmentation tasks. Indeed, certain textures, such as directional ones, provide very considerable amounts of additional information, and it is very beneficial being able to use this for image interpretation; and in the cases where the texture reflects the 3D shapes of objects, even more can be learnt about the scene—albeit not without significant algorithmic and computational effort.

In the remainder of this chapter, we will first examine how effective the busyness idea is when used to perform texture analysis. We will then turn to obvious rigorous approaches such as autocorrelation and Fourier methods. After noting their limitations, we will examine co-occurrence matrices, and then consider the texture energy method which gradually took over, culminating in the eigenfilter approach. After considering potential problems with texture segmentation, an X-ray inspection application will be taken as a practical example of texture analysis in action. At this point the chapter will broaden out to deal with the wider scene—fractals, Markov models, structural techniques, and 3D shape from texture—together with outlines of recent novel approaches. The chapter will close with a summary, and with a forward look to later chapters.

1.2. A Simple Texture Analysis Technique and Its Limitations

In this section we start with the ‘busyness’ idea outlined in the previous section, and see how far it can be taken in a practical situation. In particular, we consider how to discriminate two types of seeds—rape seeds and charlock seeds, the former being used for the production of rape seed oil, and the latter counting as weeds. Rape seeds are characterised by a peaked, almost prickly surface, while charlock seeds have a smooth surface.^a When seen in digital images, the rape seeds exhibit a speckled surface texture, and the charlock seeds have very little texture. Figure 1.4 illustrates a simple procedure for discriminating between the two types of seed. First, non-maximum intensities are suppressed and a non-critical threshold is applied to eliminate low-level peaks due to noise. Then mere counting of intensity peaks in the vicinity of centre locations clearly leads to accurate discrimination between the two types of seed. In this application, intensity and colour alone are unreliable indicators of seed identity, while size is a relatively good indicator, but would almost certainly lead to one error for the image shown in Fig. 1.4.

^aMost rape seeds will have similar numbers of peaks on their surface: hence their appearance exhibits both the randomness and the regularity expected of a texture.



(a)



(b)

Fig. 1.4. Texture processing to discriminate between two types of seed. (a) Original image showing 4 rape seeds (speckled) and 6 charlock seeds (dark and smooth). (b) Processed image showing approximate centre locations (red crosses) and bright peaks of intensity (green dots). Simple counting of centre locations clearly leads to accurate discrimination between the two types of seed.

While the texture analysis technique described above fulfils the immediate demands of the application, it can hardly be described as general or generic. In particular, the seeds are located prior to their classification by texture analysis, and in the majority of applications methods are needed that perform segmentation as an intrinsic part of the analysis. Furthermore, only intensity maxima are considered, and thus the richness of information available in many textures would be disregarded. Nevertheless, the method acts as a powerful existence theorem spurring detailed further study of the many techniques now available to workers in this important area.

1.3. Autocorrelation and Fourier Methods

In Section 1.1 texture emerged as the characteristic variation in intensity of a region of an image which should allow us to recognise and describe it and to outline its boundaries. In view of the statistical nature of textures, this prompts us to characterise texture by the variance in intensity values taken over the whole region of the texture.^b However, such an approach will not give a rich enough description of the texture for most purposes, and will certainly not provide any possibility of reconstruction: it will also be especially unsuitable in cases where the texels are well defined, or where there is a high degree of periodicity in the texture. On the other hand, for highly periodic textures such as those that arise with many textiles, it is natural to consider the use of Fourier analysis. Indeed, in the early days of image analysis, this approach was tested thoroughly, though the results were not always encouraging. (Considering that cloth can easily be stretched locally by several weave periods, this is hardly surprising.)

Bajcsy (1973)¹ used a variety of ring and orientated strip filters in the Fourier domain to isolate texture features—an approach that was found to work successfully on natural textures such as grass, sand and trees. However, there is a general difficulty in using the Fourier power spectrum in that the information is more scattered than might at first be expected. In addition, strong edges and image boundary effects can prevent accurate texture analysis by this method, though Shaming (1974)² and Dyer and Rosenfeld (1976)³ tackled the relevant image aperture problems. Perhaps more important is the fact that the Fourier approach is a global one which

^bWe defer for now the problem of finding the region of a texture so that we can compute its characteristics in order to perform a segmentation function. However, some preliminary training of a classifier may clearly be used to overcome this problem for supervised texture segmentation tasks.




Fig. 1.5. Use of autocorrelation function for texture analysis. This diagram shows the possible 1D profile of the autocorrelation function for a piece of material in which the weave is subject to significant spatial variation: notice that the periodicity of the autocorrelation function is damped down over quite a short distance. © Elsevier 2005.

is difficult to apply successfully to an image that is to be segmented by texture analysis (Weszka *et al.*, 1976).⁴

Autocorrelation is another obvious approach to texture analysis, since it should show up both local intensity variations and also the repeatability of the texture (see Fig. 1.5). In particular, it should be useful for distinguishing between short-range and long-range order in a texture. An early study was carried out by Kaizer (1955).⁵ He examined how many pixels an image has to be shifted before the autocorrelation function drops to $1/e$ of its initial value, and produced a subjective measure of coarseness on this basis. However, Rosenfeld and Troy (1970)^{6,7} later showed that autocorrelation is not a satisfactory measure of coarseness. In addition, autocorrelation is not a very good discriminator of isotropy in natural textures. Hence workers were quick to take up the co-occurrence matrix approach introduced by Haralick *et al.* in 1973:⁸ in fact, this approach not only replaced the use of autocorrelation but during the 1970s became to a large degree the ‘standard’ approach to texture analysis.

1.4. Grey-Level Co-Occurrence Matrices

The grey-level co-occurrence matrix approach^c is based on studies of the statistics of pixel intensity distributions. As hinted above with regard to the variance in pixel intensity values, single pixel statistics do not provide rich enough descriptions of textures for practical applications. Thus it is natural

^cThis is also frequently called the spatial grey-level dependence matrix (SGLDM) approach.

to consider second order statistics obtained by considering pairs of pixels in certain spatial relations to each other. Hence, co-occurrence matrices are used, which express the relative frequencies (or probabilities) $P(i, j|d, \theta)$ with which two pixels having relative polar coordinates (d, θ) appear with intensities i, j . The co-occurrence matrices provide raw numerical data on the texture, though this data must be condensed to relatively few numbers before it can be used to classify the texture. The early paper by Haralick *et al.* (1973)⁸ gave fourteen such measures, and these were used successfully for classification of many types of material (including, for example, wood, corn, grass and water). However, Connors and Harlow (1980)⁹ found that only five of these measures were normally used, viz. ‘energy’, ‘entropy’, ‘correlation’, ‘local homogeneity’ and ‘inertia’ (note that these names do not provide much indication of the modes of operation of the respective operators).

To obtain a more detailed idea of the operation of the technique, consider the co-occurrence matrix shown in Fig. 1.6. This corresponds to a nearly uniform image containing a single region in which the pixel intensities are subject to an approximately Gaussian noise distribution, the




Fig. 1.6. Co-occurrence matrix for a nearly uniform grey-scale image with superimposed Gaussian noise. Here the intensity variation is taken to be almost continuous: normal convention is followed by making the j index increase downwards, as for a table of discrete values (cf. Fig. 1.8). © Elsevier 2005.




Fig. 1.7. Co-occurrence matrix for an image with several distinct regions of nearly constant intensity. Again, the leading diagonal of the diagram is from top left to bottom right (cf. Figs. 1.6 and 1.8). © Elsevier 2005.

attention being on pairs of pixels at a constant vector distance $\mathbf{d} = (d, \theta)$ from each other. Next consider the co-occurrence matrix shown in Fig. 1.7, which corresponds to an almost noiseless image with several nearly uniform image regions. In this case the two pixels in each pair may correspond either to the same image regions or to different ones, though if d is small they will only correspond to adjacent image regions. Thus we have a set of N on-diagonal patches in the co-occurrence matrix, but only a limited number L of the possible number M of off-diagonal patches linking them, where $M = \binom{N}{2}$ and $L \leq M$ (typically L will be of order N rather than N^2). With textured images, if the texture is not too strong, it may be modelled as noise, and the $N + L$ patches in the image will be larger but still not overlapping. However, in more complex cases the possibility of segmentation using the co-occurrence matrices will depend on the extent to which \mathbf{d} can be chosen to prevent the patches from overlapping. Since many textures are directional, careful choice of θ will clearly help with this task, though the optimum value of d will depend on several other characteristics of the texture.

As a further illustration, we consider the small image shown in Fig. 1.8(a). To produce the co-occurrence matrices for a given value of

0	0	0	1
1	1	1	1
2	2	2	3
3	3	4	5

(a)

	0	1	2	3	4	5
0	2	1	0	0	0	0
1	1	3	0	0	0	0
2	0	0	2	1	0	0
3	0	0	1	1	1	0
4	0	0	0	1	0	1
5	0	0	0	0	1	0

(b)

	0	1	2	3	4	5
0	0	3	0	0	0	0
1	3	1	3	1	0	0
2	0	3	0	2	1	0
3	0	1	2	0	0	1
4	0	0	1	0	0	0
5	0	0	0	1	0	0

(c)

Fig. 1.8. Co-occurrence matrices for a small image. (a) shows the original image; (b) shows the resulting co-occurrence matrix for $\mathbf{d} = (1, 0)$, and (c) shows the matrix for $\mathbf{d} = (1, \pi/2)$. Note that even in this simple case the matrices contain more data than the original image. © Elsevier 2005.

\mathbf{d} , we merely need to calculate the numbers of cases for which pixels a distance \mathbf{d} apart have intensity values i and j . Here, we content ourselves with the two cases $\mathbf{d} = (1, 0)$ and $\mathbf{d} = (1, \pi/2)$. We thus obtain the matrices shown in Fig. 1.8(b) and (c).

This simple example demonstrates that the amount of data in the matrices is liable to be many times more than in the original image—a situation which is exacerbated in more complex cases by the number of values of d and θ that are required to accurately represent the texture. In addition, the number of grey-levels will normally be closer to 256 than to 6, and the amount of matrix data varies as the square of this number. Finally,

we should notice that the co-occurrence matrices merely provide a new representation: they do not themselves solve the recognition problem.

These factors mean that the grey-scale has to be compressed into a much smaller set of values, and careful choice of specific sample d , θ values must be made: in most cases it is not at all obvious how such a choice should be made, and it is even more difficult to arrange for it to be made automatically. In addition, various functions of the matrix data must be tested before the texture can be properly characterised and classified.

These problems with the co-occurrence matrix approach have been tackled in many ways: just two are mentioned here. The first is to ignore the distinction between opposite directions in the image, thereby reducing storage by 50%. The second is to work with differences between grey-levels; this amounts to performing a summation in the co-occurrence matrices along axes parallel to the main diagonal of the matrix. The result is a set of first order difference statistics. While these modifications have given some additional impetus to the approach, the 1980s saw a highly significant diversification of methods for the analysis of textures. Of these, Laws' approach (1979, 1980)¹⁰⁻¹² is important in that it has led to other developments which provide a systematic, adaptive means of tackling texture analysis. This approach is covered in the following section.

1.5. The Texture Energy Approach

In 1979 and 1980 Laws presented his novel texture energy approach to texture analysis (1979, 1980).¹⁰⁻¹² This involved the application of simple filters to digital images. The basic filters he used were common Gaussian, edge detector and Laplacian-type filters, and were designed to highlight points of high 'texture energy' in the image. By identifying these high energy points, smoothing the various filtered images, and pooling the information from them he was able to characterise textures highly efficiently and in a manner compatible with pipelined hardware implementations. As remarked earlier, Laws' approach has strongly influenced much subsequent work and it is therefore worth considering it here in some detail.

The Laws' masks are constructed by convolving together just three basic 1×3 masks:

$$L3 = [1 \quad 2 \quad 1] \quad (1.1)$$

$$E3 = [-1 \quad 0 \quad 1] \quad (1.2)$$

$$S3 = [-1 \quad 2 \quad -1] \quad (1.3)$$

The initial letters of these masks indicate *Local* averaging, *Edge* detection and *Spot* detection. In fact, these basic masks span the entire 1×3 subspace and form a complete set. Similarly, the 1×5 masks obtained by convolving pairs of these 1×3 masks together form a complete set:^d

$$L5 = [1 \ 4 \ 6 \ 4 \ 1] \quad (1.4)$$

$$E5 = [-1 \ -2 \ 0 \ 2 \ 1] \quad (1.5)$$

$$S5 = [-1 \ 0 \ 2 \ 0 \ -1] \quad (1.6)$$

$$R5 = [1 \ -4 \ 6 \ -4 \ 1] \quad (1.7)$$

$$W5 = [-1 \ 2 \ 0 \ -2 \ 1] \quad (1.8)$$

(Here the initial letters are as before, with the addition of *Ripple* detection and *Wave* detection.) We can also use matrix multiplication to combine the 1×3 and a similar set of 3×1 masks to obtain nine 3×3 masks—for example:

$$\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} [-1 \ 2 \ -1] = \begin{bmatrix} -1 & 2 & -1 \\ -2 & 4 & -2 \\ -1 & 2 & -1 \end{bmatrix} \quad (1.9)$$

The resulting set of masks also forms a complete set (Table 1.1): note that two of these masks are identical to the Sobel operator masks. The corresponding 5×5 masks are entirely similar but are not considered in detail here as all relevant principles are illustrated by the 3×3 masks.

All such sets of masks include one whose components do not average to zero. Thus it is less useful for texture analysis since it will give results dependent more on image intensity than on texture. The remainder are sensitive to edge points, spots, lines and combinations of these.

Having produced images that indicate local edginess, etc., the next stage is to deduce the local magnitudes of these quantities. These magnitudes are then smoothed over a fair-sized region rather greater than the basic filter mask size (e.g. Laws used a 15×15 smoothing window after applying his 3×3 masks): the effect of this is to smooth over the gaps between the texture edges and other micro-features. At this point the image has been transformed into a vector image, each component of which represents energy of a different type. While Laws (1980)¹² used both squared magnitudes and absolute magnitudes to estimate texture energy, the former corresponding

^dIn principle nine masks can be formed in this way, but only five of them are distinct.

to true energy and giving a better response, the latter are useful in requiring less computation:

$$E(l, m) = \sum_{i=l-p}^{l+p} \sum_{j=m-p}^{m+p} |F(i, j)| \quad (1.10)$$

$F(i, j)$ being the local magnitude of a typical microfeature which is smoothed at a general scan position (l, m) in a $(2p + 1) \times (2p + 1)$ window.

Table 1.1. The nine 3×3 Laws masks.
© Elsevier 2005.

L3[⊤]L3	L3[⊤]E3	L3[⊤]S3
1 2 1	-1 0 1	-1 2 -1
2 4 2	-2 0 2	-2 4 -2
1 2 1	-1 0 1	-1 2 -1
E3[⊤]L3	E3[⊤]E3	E3[⊤]S3
-1 -2 -1	1 0 -1	1 -2 1
0 0 0	0 0 0	0 0 0
1 2 1	-1 0 1	-1 2 -1
S3[⊤]L3	S3[⊤]E3	S3[⊤]S3
-1 -2 -1	1 0 -1	1 -2 1
2 4 2	-2 0 2	-2 4 -2
-1 -2 -1	1 0 -1	1 -2 1

A further stage is required to combine the various energies in a number of different ways, providing several outputs which can be fed into a classifier to decide upon the particular type of texture at each pixel location (Fig. 1.9): if necessary, principal components analysis is used at this point to help select a suitable set of intermediate outputs.

To understand the process more clearly, consider the use of masks $L3^{\top}E3$ and $E3^{\top}L3$. If their responses are squared and added, we have a very similar situation to a Sobel operator. An alternate result can be obtained for directional textures by using the same mask responses and applying the arctan function—which can be regarded as enhancing the classifier (Fig. 1.9) in a particular way.

Laws' method resulted in excellent classification accuracy quoted at (for example) 87% compared with 72% for the co-occurrence matrix method,




Fig. 1.9. Basic form for a Laws texture classifier. Here I is the incoming image, M represents the microfeature calculation, E the energy calculation, S the smoothing, and C the final classification. © Elsevier 2005.

when applied to a composite texture image of grass, raffia, sand, wool, pigskin, leather, water and wood (Laws, 1980).¹² He also found that the histogram equalisation normally applied to images to eliminate first-order differences in texture field grey-scale distributions gave little improvement in this case.

Research was undertaken by Pietikäinen *et al.* (1983)¹³ to determine whether the precise coefficients used in the Laws' masks are responsible for the performance of his method. They found that so long as the general forms of the masks were retained, performance did not deteriorate, and could in some instances be improved. They were able to confirm that Laws' texture energy measures are more powerful than measures based on pairs of pixels (i.e. co-occurrence matrices).

1.6. The Eigenfilter Approach

In 1983 Ade¹⁴ investigated the theory underlying the Laws' approach, and developed a revised rationale in terms of eigenfilters. He took all possible pairs of pixels within a 3×3 window, and characterised the image intensity data by a 9×9 covariance matrix. He then determined the eigenvectors required to diagonalise this matrix. These correspond to filter masks similar to the Laws' masks, i.e. use of these 'eigenfilter' masks produces images

which are principal component images for the given texture. Furthermore, each eigenvalue gives that part of the variance of the original image that can be extracted by the corresponding filter. Essentially, the variances give an exhaustive description of a given texture in terms of the texture of the images from which the covariance matrix was originally derived. Clearly, the filters that give rise to low variances can be taken to be relatively unimportant for texture recognition.

It will be useful to illustrate the technique for a 3×3 window. Here we follow Ade (1983)¹⁴ in numbering the pixels within a 3×3 window in scan order:

1	2	3
4	5	6
7	8	9

This leads to a 9×9 covariance matrix for describing relationships between pixel intensities within a 3×3 window, as stated above. At this point we recall that we are describing a texture, and assuming that its properties are not synchronous with the pixel tessellation, we would expect various coefficients of the covariance matrix C to be equal: for example, C_{24} should equal C_{57} ; in addition, C_{57} must equal C_{75} . It is worth pursuing this matter, as a reduced number of parameters will lead to increased accuracy in determining the remaining ones. In fact, there are $\binom{9}{2} = 36$ ways of selecting pairs of pixels, but there are only 12 distinct spatial relationships between pixels if we disregard translations of whole pairs—or 13 if we include the null vector in the set (see Table 1.2). Thus the covariance matrix takes the form:

$$C = \begin{bmatrix} a & b & f & c & d & k & g & m & h \\ b & a & b & e & c & d & l & g & m \\ f & b & a & j & e & c & i & l & g \\ c & e & j & a & b & f & c & d & k \\ d & c & e & b & a & b & e & c & d \\ k & d & c & f & b & a & j & e & c \\ g & l & i & c & e & j & a & b & f \\ m & g & l & d & c & e & b & a & b \\ h & m & g & k & d & c & f & b & a \end{bmatrix} \quad (1.11)$$

C is symmetric, and the eigenvalues of a real symmetric covariance matrix are real and positive, and the eigenvectors are mutually orthogonal. In addition, the eigenfilters thus produced reflect the proper structure of

the texture being studied, and are ideally suited to characterising it. For example, for a texture with a prominent highly directional pattern, there will be one or more high energy eigenvalues with eigenfilters having strong directionality in the corresponding direction.

Table 1.2. Spatial relationships between pixels in a 3×3 window.

<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>	<i>h</i>	<i>i</i>	<i>j</i>	<i>k</i>	<i>l</i>	<i>m</i>
9	6	6	4	4	3	3	1	1	2	2	2	2

This table shows the number of occurrences of the spatial relationships between pixels in a 3×3 window. Note that *a* is the diagonal element of the covariance matrix *C*, and that all others appear twice as many times in *C* as indicated in the table. © Elsevier 2005.

1.7. Appraisal of the Texture Energy and Eigenfilter Approaches

At this point, it will be worthwhile to compare the Laws and Ade approaches more carefully. In the Laws approach standard filters are used, texture energy images are produced, and then principal component analysis may be applied to lead to recognition; whereas in the Ade approach, special filters (the eigenfilters) are applied, incorporating the results of principal component analysis, following which texture energy measures are calculated and a suitable number of these are applied for recognition.

The Ade approach is superior to the extent that it permits low-value energy components to be eliminated early on, thereby saving computation. For example, in Ade's application, the first five of the nine components contain 99.1% of the total texture energy, so the remainder can definitely be ignored; in addition, it would appear that another two of the components containing respectively 1.9% and 0.7% of the energy could also be ignored, with little loss of recognition accuracy. However, in some applications textures could vary continually, and it may well not be advantageous to fine-tune a method to the particular data pertaining at any one time.^e In addition, to do so may prevent an implementation from having wide generality or (in the case of hardware implementations) being so cost-effective.

^eFor example, these remarks apply (1) to textiles, for which the degree of stretch will vary continuously during manufacture, (2) to raw food products such as beans, whose sizes will vary with the source of supply, and (3) to processed food products such as cakes, for which the crumbliness will vary with cooking temperature and water vapour content.

There is therefore still a case for employing the simplest possible complete set of masks, and using the Laws approach.

In 1986, Unser¹⁵ developed a more general version of the Ade technique. In this approach not only is performance optimised for texture classification but also it is optimised for discrimination between two textures by simultaneous diagonalisation of two covariance matrices. The method has been developed further by Unser and Eden (1989, 1990):^{16,17} this work makes a careful analysis of the use of non-linear detectors. As a result, two levels of non-linearity are employed, one immediately after the linear filters and designed (by employing a specific Gaussian texture model) to feed the smoothing stage with genuine variance or other suitable measures, and the other after the spatial smoothing stage to counteract the effect of the earlier filter, and aiming to provide a feature value that is in the same units as the input signal. In practical terms this means having the capability for providing an r.m.s. texture signal from each of the linear filter channels.

Overall, the originally intuitive Laws approach emerged during the 1980s as a serious alternative to the co-occurrence matrix approach. It is as well to note that alternative methods that are potentially superior have also been devised—see for example the local rank correlation method of Harwood *et al.* (1985),¹⁸ and the forced-choice method of Vistnes (1989)¹⁹ for finding edges between different textures which apparently has considerably better accuracy than the Laws approach. Vistnes's (1989)¹⁹ investigation concludes that the Laws approach is limited by (a) the small scale of the masks which can miss larger-scale textural structures, and (b) the fact that the texture energy smoothing operation blurs the texture feature values across the edge. The latter finding (or the even worse situation where a third class of texture appears to be located in the region of the border between two textures) has also been noted by Hsiao and Sawchuk (1989)^{20,21} who applied an improved technique for feature smoothing; they also used probabilistic relaxation for enforcing spatial organisation on the resulting data.

1.8. Problems with Texture Segmentation

As noted in the previous section, when texture analysis algorithms such as the Laws' method are used for texture segmentation, inappropriate regions and classifications are frequently encountered. These arise because the statistical nature of textures means that spatial smoothing has to be done at a certain stage of the process: as a result, the transition region between




Fig. 1.10. Problems with texture segmentation. (a) depicts an original texture where darkness indicates the density of peaks near each pixel. (b) shows where spurious classifications can occur between (in this case) one pair of regions and also between a triplet of regions.

textures may be classified as a totally different texture. For example, if one texture T_1 has n_1 intensity peaks over a smoothing area A , and another texture T_2 has n_2 such peaks over a corresponding area, then a transition region can easily have $(n_1 + n_2)/2$ intensity peaks over an intermediate smoothing area. If this is close to the number of peaks expected for a third texture T_3 on which the classifier has been trained, the transition region will be classified as type T_3 .

The situation will be even more complicated where three textures T_1 , T_2 , T_3 come together at a point P . Not only can the texture regions between pairs of textures be segmented and classified erroneously, but a further texture region may also appear around P . Figure 1.10 shows a possible scenario for this. In this case, if there are n_1, n_2, n_3 peaks over the corresponding smoothing areas, T_4 appears if $n_4 \simeq (n_1 + n_2)/2$ and T_5 appears if $n_5 \simeq (n_1 + n_2 + n_3)/3$. Note that such situations will not arise if the classifier has not been trained on the additional textures T_4 and T_5 . However, the output of the smoothing area will still change gradually from n_1 to n_2 on moving from T_1 to T_2 , and similarly for the other cases. (Note that the scenario depicted in Fig. 1.10 is not the worst possible, as additional textures could appear in all three transition regions T_1 - T_2 , T_2 - T_3 , T_3 - T_1 , and also in the triple transition region T_1 - T_2 - T_3 .)

Fortunately, these types of misclassification scenario should appear less often than indicated above. This is because we have assumed that just one microfeature is being measured: but in fact, as Fig. 1.9 shows, there will normally be nine or more microfeatures, and each should lead to different




Fig. 1.11. Texture segmentation with directional textures. This figure shows that with perfect pattern structure, particularly good segmentation can be performed—though this may well not apply for less artificial textures or when the boundaries between textures are at all crinkly.

ways in which textures such as T_1 and T_2 differ. There will therefore be less likelihood that spurious regions such as T_4 will occur. However, regions such as T_5 have a higher probability of occurring, because there are so many combinations of textures that can arise in the many sampling regions surrounding P . Again, generalisation is difficult because a lot depends on the exact training that the classifier has been subjected to.

Overall, we can say that while any individual microfeature may allow a texture to be misclassified in a particular transition region, there is much less likelihood that a number of them will act coherently in this way, though the possibility is distinctly enhanced where three textures meet at a point.

Finally, we enquire whether it is ever impossible for such spurious regions to occur. Take, for example, the case of a set of three directional textures meeting at a point, as in Fig. 1.11. Analysis of the patterns could then yield strict boundaries between them without the possibility of spurious regions being introduced. However, such accurately constructed patterns lack the partial randomness characteristic of a texture; thus it is difficult to envisage smoothing areas not being needed for real textures. For further enlightenment on such points, see Vistnes (1989)¹⁹ and Hsiao and Sawchuk (1989).^{20,21}

1.9. An X-Ray Inspection Application

The application outlined in this section relates to the inspection of bags of frozen vegetables such as peas, sweetcorn or stir-fry (Patel *et al.* 1996).²²

In the past it has been usual to use X-rays for this purpose, as ‘hard’ contaminants such as pieces of metal can be located in the images by global thresholding. However, such schemes are very poor at locating ‘soft’ contaminants such as wood, plastic and rubber; in addition, they are often unable to detect small stones even though these are commonly classed as hard contaminants. One basic problem is the high level of intensity variation in the X-ray image of the substrate vegetable matter: the fact that several layers of vegetables contribute to the same image means that the latter appears highly textured, and it is rather ineffective to apply simple thresholding.

In this application, no assumptions can be made about the individual foreign objects that might occur, so the usual algorithms for locating defects cannot be used. In particular, shape analysis and simple measures of intensity are mostly inappropriate, and thus it is necessary to recognise foreign objects by the fact that they disrupt the normal (textural) intensity pattern of the substrate. *A priori*, it might have been thought that a set of feedforward artificial neural networks (ANNs), each adapted to detect a particular foreign object, would be useful. However, so many types of foreign object with so many possible shapes and sizes can occur that this is not a viable approach except for certain crucial contaminants.

To solve this problem Laws’ approach to texture analysis was adopted. The reasons for this choice were (a) ease of setting up and (b) the fact that Laws’ approach is well adapted to hardware implementation as it employs small neighbourhood convolutions to obtain a set of processed images. Summing the ‘textural energies’ in these images permits any foreign objects to be detected by thresholding coupled with a minor amount of further processing.

Following Ade’s (1983)¹⁴ modification of the Laws’ schema, it was found that sensitivity is enhanced by making use of principal components analysis (PCA). However, instead of using conventional diagonalisation procedures, the Hebbian type of ANN (Oja 1989)²³ was adopted. A major advantage of the Hebbian approach is that it permits PCA to be applied without the huge computational load that would be expected when dealing with large matrices.

Finally, by adopting a statistical pattern recognition approach, it is possible to classify the images into three regions—background region, food-bag region, and any foreign object regions. Thus it is unnecessary to have a preliminary stage of bag location—with the result that the whole inspection algorithm becomes significantly more efficient.




Fig. 1.12. Foreign object detection system. The input image comes in on the left, and the output classification (which is not an image) emerges on the right. © World Scientific 2000.

1.9.1. Further details of the algorithm

The considerations mentioned above lead to a system design of the form shown in Fig. 1.12. In particular, the initial acquisition stage is followed by a preprocessing stage, a feature extraction stage and a decision stage. For simplicity, the Hebbian training paths are not shown in this figure, which just includes the data paths for normal testing of the input images.

A major part of Fig. 1.12 that has not been covered by the earlier discussion is the preprocessing stage. In fact, this has several components. The first is the log transform, which compensates for the non-linearity of the image acquisition process, thereby making the occupation levels of the grey-levels more uniform and the subsequent processing more reliable. Rank-order filtering provides further capabilities for preprocessing. In particular, local intensity minimisation operations have been found valuable for expanding small dark foreign objects in order to make them more easily discernible. In some cases the same operation has also been found useful for enhancing the contrast between soft contaminants and the food substrate. Finally, thresholding is added to the texture analysis scheme, both to provide the capability for locating contaminants directly and as the final decision-making stage of the texture analysis process.

It was found to be both effective and computationally efficient to use Laws' masks of size 3×3 to form the microfeatures, and, following absolute value determination, to use smoothing masks of size 5×5 to obtain the texture energy macrofeatures. The tests were made with 1 lb. bags of frozen sweetcorn kernels into which foreign objects of various shapes, sizes and origins were inserted: specifically, foreign objects consisting of small pieces of glass, metal and stone and larger pieces of plastic, rubber and wood were used for this purpose (Fig. 1.13).




Fig. 1.13. Foreign object detection using texture analysis. (top left) Original X-ray image of a packet of frozen sweetcorn. (top right) An image in which any foreign objects (here a splinter of glass) have been enhanced by texture analysis. (bottom left and right) The respective thresholded images. Notice that false alarms are starting to arise in the bottom left, whereas in the bottom right there is much increased confidence in the detection of foreign objects. © MCB University Press 1995.

1.10. Other Approaches to Texture Analysis

1.10.1. *Fractal-based measures of texture*

An important new approach to texture analysis that arose in the 1980s was that of fractals. This incorporates the observation due to Mandelbrot (1982)²⁴ that measurements of the length of a coastline (for example) will vary with the size of the measuring tool used for the purpose, since details smaller than the size of the tool will be missed. If the size of the measuring

tool is taken as λ , the measured quantity will be $M = n\lambda^D$, where D is known as the fractal dimension and must in general be larger than the immediate geometric dimension if correct measurements are to result (for a coastline we will thus have $D > 2$). Thus, when measurements are being made of 2D textures, it is found that D can take values from 2.0 to at least 2.8 (Pentland, 1984).²⁵ Interestingly, these values of D have been found to correspond roughly to subjective measures of the roughness of the surface being inspected (Pentland, 1984).²⁵

Since the fractal approach was put forward by Pentland (1984),²⁵ other workers have expressed certain problems with it. For example, reducing all textural measurements to the single measure D clearly cannot permit all textures to be distinguished (Keller *et al.*, 1989).²⁶ Hence there have been moves to define further fractal-based measures. Mandelbrot himself brought in the concept of lacunarity and in 1982 provided one definition, while Keller *et al.* (1989)²⁶ and others provided further definitions.

Finally, note that Gärding (1988)²⁷ found that fractal dimension is not always equivalent to subjective judgements of roughness: in particular he found that a region of Gaussian noise of low amplitude superimposed on a constant grey-level will have a fractal dimension that approaches 3.0—a rather high value, which is contrary to our judgement of such surfaces as being quite smooth. (An interpretation of this result is that highly noisy textures appear exactly like 3D landscapes in relief!)

1.10.2. Markov random field models of texture

Markov models have long been used for texture synthesis, to help with the generation of realistic images. However, they have also proved increasingly useful for texture analysis. In essence a Markov model is a 1D construct in which the intensity at any pixel depends only upon the intensity of the previous pixel in a chain and upon a transition probability matrix. For images this is too weak a characterisation, and various more complex constructs have been devised. Interest in such models dates from as early as 1965 (Abend *et al.*, 1965),²⁸ and during the 1980s a considerable amount of further work was being published (e.g. Geman and Geman, 1984; Derin and Elliott, 1987).^{29,30} Space does not permit details of these algorithms to be given here. Suffice it to say that by 1987 impressive results for texture segmentation of real scenes were being achieved using this approach (Derin and Elliott, 1987).³⁰

1.10.3. *Structural approaches to texture analysis*

It has already been remarked that textures approximate to a basic textural element or primitive that is replicated in a more or less regular manner. Structural approaches to texture analysis aim to discern the textural primitive and to determine the underlying gross structure of the texture. Early work (e.g. Pickett, 1970)³¹ suggested the structural approach, though little research on these lines was carried out until the late 1970s—e.g. Davis (1979).³² An unusual and interesting paper by Kass and Witkin (1987)³³ shows how orientated patterns from wood grain, straw, fabric and fingerprints, and also spectrograms and seismic patterns can be analysed: the method adopted involves building up a flow coordinate system for the image, though the method rests more on edge pattern orientation analysis than on more usual texture analysis procedures. A similar statement may be made about the topologically invariant texture descriptor method of Eichmann and Kasparis (1988),³⁴ which relies on Hough transforms for finding line structures in highly structured textiles. More recently, pyramidal approaches have been applied to structural texture segmentation (Lam and Ip, 1994).³⁵

1.10.4. *3D shape from texture*

This is another topic in texture analysis that developed strongly during the 1980s. After early work by Bajcsy and Liebermann (1976)³⁶ for the case of planar surfaces, Witkin (1981)³⁷ significantly extended this work and at the same time laid the foundations for general development of the whole subject. Many papers followed (e.g. Aloimonos and Swain, 1985; Stone, 1990)^{38,39} but there is no space to cover them all here. In general, workers have studied how an assumed standard texel shape is distorted and its size changed by 3D projections; they then relate this to the local orientation of the surface. Since the texel distortion varies as the cosine of the angle between the line of sight and the local normal to the surface plane, essentially similar ‘reflectance map’ analysis is required as in the case of shape-from-shading estimation. An alternative approach adopted by Chang *et al.* (1987)⁴⁰ involves texture discrimination by projective invariants. More recently, Singh and Ramakrishna (1990)⁴¹ exploited shadows and integrated the information available from texture and from shadows.

1.10.5. More recent developments

Recent developments include further work on automated visual inspection (e.g. Davies, 2000; Pun and Lee, 2003),^{42,43} medical, remote sensing and other applications. The paper by Pun and Lee is specifically aimed at rotation-invariant texture classification but also aims at scale invariance. Other work (Clerc and Mallat, 2002)⁴⁴ is concerned with recovering shape from texture via a texture gradient equation, while Ma *et al.* (2003)⁴⁵ are particularly concerned with person identification based on iris textures. Mirmehdi and Petrou (2000)⁴⁶ describe an in-depth investigation of colour texture segmentation. In this context, the importance of ‘wavelets’^f as an increasingly used technique of texture analysis with interesting applications (such as human iris recognition) should be noted (e.g. Daugman, 2003).⁴⁸ Note that they solve in a neat way the problems of Fourier analysis that were noted in Sec. 1.3 (essentially, they act as local Fourier transforms).

Finally, in a particularly exciting advance, Spence *et al.* (2003)⁴⁹ managed to eliminate texture by using photometric stereo to find the underlying surface shape (or ‘bump map’), following which they were able to perform impressive reconstructions, including texture, from a variety of viewpoints; McGunnigle and Chantler (2003)⁵⁰ have shown that this sort of technique is also able to reveal hidden writing on textured surfaces, where only pen pressure marks have been made. Similarly, Pan *et al.* (2004)⁵¹ have shown how texture can be eliminated from ancient tablets (in particular those made of lead and wood) to reveal clear images of the writing underneath.

1.11. Concluding Remarks

This chapter started by exploring the meaning of texture—essentially by asking “What is a texture and how is a texture formed?” Typically, a texture starts with a surface^g that exhibits local roughness or structure, which is then projected to form a textured image. Such an image exhibits both regularity and randomness to varying degrees: directionality and orientation will also be relevant parameters in a good many cases. However, the essential feature of randomness means that textures have to be characterised by statistical techniques, and recognised using statistical

^fWavelets are directional filters reminiscent of the Laws edges, bars, waves and ripples, but have more rigorously defined shapes and envelopes, and are defined in multiresolution sets (Mallat, 1989).⁴⁷

^gNaturally, textures also arise inside solid bodies, seen through the medium of X-rays.

classification procedures. Techniques that have been used for this purpose have been seen to include autocorrelation, co-occurrence matrices, texture energy measures, fractal-based measures, Markov random fields, and so on. These aim both to analyse and to model the textures. Indeed, it can be said that workers in this area spend much time striving to achieve ever-improved models of the textures they are working with in order to better recognise and segment them. Failure to model accurately in the end means failure to perform the requisite classification tasks. And, as elsewhere in vision, modelling is the key: we need to be able to generate accurate look-alike scenes in order to succeed with classification. Nevertheless, an additional ingredient is necessary—the ability to infer the parameters that permit the currently viewed scene to be modelled. In fact, using different techniques, different representations and procedures will be necessary in order to perform the optimisations, and, again, workers have to strive to make their own techniques work well. The early success with PCA (cf. the Unser and Ade enhancement of the Laws approach) reflects this, but at the same time this approach has its limitations. This is why so many other methods are described by the authors of the later chapters in this book. Not only do we find Markov random field models, but also local statistical operators, ‘texems’, hierarchical texture descriptions, bidirectional reflectance distribution functions, trace transforms, structural approaches, and more. The developing methodology is so wide and so varied^h that it seems difficult to consider texture analysis as a mature subject: but yet, in terms of the ideas outlined above, it is clear that each of the authors has been able to find generic statistical approaches that match some important subset of the wide and complex range of textures that exist in the real world.

Acknowledgements

The work on this chapter has been supported by Research Councils UK, under Basic Technology Grant GR/R87642/02. Tables 1.1 and 1.2, and Figs. 1.5, 1.6, 1.7, 1.8 and 1.9, and some of the text are reproduced from Chapter 26 of: E.R. Davies *Machine Vision: Theory, Algorithms, Practicalities* (3rd edition, 2005), with permission from Elsevier. Figure 1.12 and some of the text are reproduced from Chapter 11 of: E.R. Davies *Image Processing for the Food Industry* (2000), with permission from World Scientific. Figure 1.13 is reproduced from: D. Patel, E.R. Davies, and

^hSee also the recent book by Petrou and Sevilla (2006).⁵²

I. Hannah *Sensor Review* 15(2):27–28 (1995), with permission from MCB University Press.

References

1. R.K. Bajcsy. Computer identification of visual surfaces. *Computer Graphics and Image Processing*, 2:118–130, October 1973.
2. W.B. Shaming. Digital image transform encoding, 1974. RCA Corp. paper no. PE-622.
3. C.R. Dyer and A. Rosenfeld. Fourier texture features: Suppression of aperture effects. *IEEE Trans. Systems, Man and Cybernetics*, 6:703–705, 1976.
4. J.S. Weszka and A. Rosenfeld. An application of texture analysis to materials inspection. *Pattern Recognition*, 8(4):195–200, October 1976.
5. H. Kaizer. *A Quantification of Textures on Aerial Photographs*. MS thesis, Boston University, 1955.
6. A. Rosenfeld and E.B. Troy. Visual texture analysis, 1970. Computer Science Center, Univ. of Maryland Techn. Report TR-116.
7. A. Rosenfeld and E.B. Troy. Visual texture analysis. In *Conf. Record for Symposium on Feature Extraction and Selection in Pattern Recogn. IEEE Publication 70C-51C, Argonne, Ill., Oct.*, pages 115–124, 1970.
8. R.M. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *IEEE Trans. Systems, Man and Cybernetics*, 3(6):610–621, November 1973.
9. R.W. Connors and C.A. Harlow. A theoretical comparison of texture algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2(3):204–222, May 1980.
10. K.I. Laws. Texture energy measures. In *Proc. Image Understanding Workshop, November*, pages 47–51, 1979.
11. K.I. Laws. Rapid texture identification. *SPIE*, 238:376–380, 1980.
12. K.I. Laws. *Textured Image Segmentation*. PhD thesis, University of Southern California, LA, 1980.
13. M. Pietikäinen, A. Rosenfeld, and L.S. Davis. Experiments with texture classification using averages of local pattern matches. *IEEE Trans. Systems, Man and Cybernetics*, 13:421–426, 1983.
14. F. Ade. Characterization of textures by ‘eigenfilters’. *Signal Processing*, 5:451–457, 1983.
15. M. Unser. Local linear transforms for texture measurements. *Signal Processing*, 11:61–79, July 1986.
16. M. Unser and M. Eden. Multiresolution feature extraction and selection for texture segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(7):717–728, July 1989.
17. M. Unser and M. Eden. Nonlinear operators for improving texture segmentation based on features extracted by spatial filtering. *IEEE Trans. Systems, Man and Cybernetics*, 20(4):804–815, 1990.
18. D. Harwood, M. Subbarao, and L.S. Davis. Texture classification by lo-

- cal rank correlation. *Computer Vision Graphics and Image Processing*, 32(3):404–411, December 1985.
19. R. Vistnes. Texture models and image measures for texture discrimination. *International Journal of Computer Vision*, 3(4):313–336, November 1989.
 20. J.Y. Hsiao and A.A. Sawchuk. Supervised textured image segmentation using feature smoothing and probabilistic relaxation techniques. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(12):1279–1292, December 1989.
 21. J.Y. Hsiao and A.A. Sawchuk. Unsupervised textured image segmentation using feature smoothing and probabilistic relaxation techniques. *Computer Vision Graphics and Image Processing*, 48(1):1–21, October 1989.
 22. D. Patel, E.R. Davies, and I. Hannah. The use of convolution-operators for detecting contaminants in food images. *Pattern Recognition*, 29(6):1019–1029, June 1996.
 23. E. Oja. A simplified neuron model as a principal component analyzer. *J. Math. Biol.*, 15:267–273, 1982.
 24. B.B. Mandelbrot. *The Fractal Geometry of Nature*. Freeman, 1982.
 25. A.P. Pentland. Fractal-based description of natural scenes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6(6):661–674, November 1984.
 26. J.M. Keller, S.S. Chen, and R.M. Crownover. Texture description and segmentation through fractal geometry. *Computer Vision Graphics and Image Processing*, 45(2):150–166, February 1989.
 27. J. Gårding. Properties of fractal intensity surfaces. *Pattern Recognition Letters*, 8:319–324, December 1988.
 28. K. Abend, T.J. Harley, and L.N. Kanal. Classification of binary random patterns. *IEEE Trans. Information Theory*, 11(4):538–544, October 1965.
 29. S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the bayesian restoration of images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6(6):721–741, November 1984.
 30. H. Derin and H. Elliott. Modelling and segmentation of noisy and textured images using Gibbs random fields. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 9(1):39–55, January 1987.
 31. R.M. Pickett. Visual analysis of texture in the detection and recognition of objects. In *Picture Processing and Psychopictorics*, pages 289–308, 1970.
 32. L.S. Davis. Computing the spatial structures of cellular texture. *Computer Graphics and Image Processing*, 11(2):111–122, October 1979.
 33. M. Kass and A.P. Witkin. Analyzing oriented patterns. *Computer Vision Graphics and Image Processing*, 37(3):362–385, March 1987.
 34. G. Eichmann and T. Kasparis. Topologically invariant texture descriptors. *Computer Vision Graphics and Image Processing*, 41(3):267–281, March 1988.
 35. S.W.C. Lam and H.H.S. Ip. Structural texture segmentation using irregular pyramid. *Pattern Recognition Letters*, 15(7):691–698, July 1994.
 36. R.K. Bajcsy and L.I. Lieberman. Texture gradient as a depth cue. *Computer Graphics and Image Processing*, 5(1):52–67, 1976.
 37. A.P. Witkin. Recovering surface shape and orientation from texture. *Artificial Intelligence*, 17(1-3):17–45, August 1981.

38. Y. Aloimonos and M.J. Swain. Shape from texture. In *Proc. IJCAI*, pages 926–931, 1985.
39. J.V. Stone. Shape from texture: textural invariance and the problem of scale in perspective images of textured surfaces. In *Proc. British Machine Vision Assoc. Conf., Oxford, 24–27 Sept.*, pages 181–186, 1990.
40. S. Chang, L.S. Davis, S.M. Dunn, A. Rosenfeld, and J.O. Eklundh. Texture discrimination by projective invariants. *Pattern Recognition Letters*, 5:337–342, May 1987.
41. R.K. Singh and R.S. Ramakrishna. Shadows and texture in computer vision. *Pattern Recognition Letters*, 11:133–141, 1990.
42. E.R. Davies. Resolution of problem with use of closing for texture segmentation. *Electronics Letters*, 36(20):1694–1696, 2000.
43. C.M. Pun and M.C. Lee. Log-polar wavelet energy signatures for rotation and scale invariant texture classification. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(5):590–603, May 2003.
44. M. Clerc and S. Mallat. The texture gradient equation for recovering shape from texture. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(4):536–549, April 2002.
45. L. Ma, T.N. Tan, Y. Wang, and D. Zhang. Personal identification based on iris texture analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(12):1519–1533, December 2003.
46. M. Mirmehdi and M. Petrou. Segmentation of colour textures. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(2):142–159, 2000.
47. S.G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(7):674–693, July 1989.
48. J.G. Daugman. Demodulation by complex-valued wavelets for stochastic pattern recognition. *Int. J. of Wavelets, Multiresolution and Information Processing*, 1(1):1–17, 2003.
49. A. Spence, M. Robb, M. Timmins, and M. Chantler. Real-time per-pixel rendering of textiles for virtual textile catalogues. In *Proc. INTEDEC, Edinburgh, 22–24 Sept.*, 2003.
50. G. McGunnigle and M.J. Chantler. Resolving handwriting from background printing using photometric stereo. *Pattern Recognition*, 36(8):1869–1879, August 2003.
51. X.B. Pan, M. Brady, A.K. Bowman, C. Crowther, and R.S.O. Tomlin. Enhancement and feature extraction for images of incised and ink texts. *Image and Vision Computing*, 22(6):443–451, June 2004.
52. M. Petrou and P.G. Sevilla. *Image Processing: Dealing with Texture*. Wiley, Chichester, UK, 2006.