


## 6.1 WHAT IS TEXTURE?

The notion of texture admits to no rigid description, but a dictionary definition of texture as “something composed of closely interwoven elements” is fairly apt. The description of interwoven elements is intimately tied to the idea of texture resolution, which one might think of as the average amount of pixels for each discernable texture element. If this number is large, we can attempt to describe the individual elements in some detail. However, as this number nears unity it becomes increasingly difficult to characterize these elements individually and they merge into less distinct spatial patterns. To see this variability, we examine some textures.

Figure 6.1 shows “cane,” “paper,” “coffee beans,” “brickwall,” “coins,” and “wire braid” after Brodatz’s well-known book [Brodatz 1966]. Five of these examples are high-resolution textures: they show repeated primitive elements that exhibit some kind of variation. “Coffee beans,” “brick wall” and “coins” all have obvious primitives (even if it is not so obvious how to extract these from image data). Two more examples further illustrate that one sometimes has to be creative in defining primitives. In “cane” the easiest primitives to deal with seem to be the physical holes in the texture, whereas in “wire braid” it might be better to model the physical relations of a loose weave of metallic wires. However, the paper texture does not fit nicely into this mold. This is not to say that there are not possibilities for primitive elements. One is regions of lightness and darkness formed by the ridges in the paper. A second possibility is to use the reflectance models described in Section 3.5 to compute “pits” and “bumps.” However, the elements seem to be “just beyond our perceptual resolving power” [Laws 1980], or in our terms, the elements are very close in size to individual pixels.



**Fig. 6.1** Six examples of texture. (a) Cane. (b) Paper. (c) Coffee beans. (d) Brick wall. (e) Coins. (f) Wire braid.

The exposition of texture takes place under four main headings:

1. Texture primitives
2. Structural models
3. Statistical models
4. Texture gradients

We have already described texture as being composed of elements of *texture primitives*. The main point of additional discussion on texture primitives is to refine the idea of a primitive and its relation to image resolution.

The main work that is unique to texture is that which describes how primitives are related to the aim of recognizing or classifying the texture. Two broad classes of techniques have emerged and we shall study each in turn. The *structural* model regards the primitives as forming a repeating pattern and describes such patterns in terms of rules for generating them. Formally, these rules can be termed a grammar. This model is best for describing textures where there is much regularity in the placement of primitive elements and the texture is imaged at high resolution. The “reptile” texture in Fig. 6.9 is an example that can be handled by the structured approach. The *statistical* model usually describes texture by statistical rules governing the distribution and relation of gray levels. This works well for many natural textures which have barely discernible primitives. The “paper” texture is such an example. As we shall see, we cannot be too rigid about this division since statistical models can describe pattern-like textures and vice versa, but in general the dichotomy is helpful.

The examples suggest that texture is almost always a property of *surfaces*. Indeed, as the example of Fig. 6.2 shows, human beings tend to relate texture elements of varying size to a plausible surface in three dimensions [Gibson 1950; Stevens 1979]. Techniques for determining surface orientation in this fashion are termed texture *gradient* techniques. The gradient is given both in terms of the direction of greatest change in size of primitives and in terms of the spatial placement of primitives. The notion of a gradient is very useful. For example, if the texture is embedded on a flat surface, the gradient points toward a vanishing point in the image. The chapter concludes with algorithms for computing this gradient. The gradient may be computed directly or indirectly via the computation of the vanishing point.




Fig. 6.2 Texture as a surface property.

## 6.2 TEXTURE PRIMITIVES

The notion of a primitive is central to texture. To highlight its importance, we shall use the appellation *texel* (for texture element) [Kender 1978]. A texel is (loosely) a visual primitive with certain invariant properties which occurs repeatedly in different positions, deformations, and orientations inside a given area. One basic invariant property of such a unit might be that its pixels have a constant gray level, but more elaborate properties related to shape are possible. (A detailed discussion of planar shapes is deferred until Chapter 8.) Figure 6.3 shows examples of two kinds of texels: (a) ellipses of approximately constant gray level and (b) linear edge segments. Interestingly, these are nearly the two features selected as texture primitives by [Julesz, 1981], who has performed extensive studies of human texture perception.

For textures that can be described in two dimensions, image-based descriptions are sufficient. Texture primitives may be pixels, or aggregates of pixels such as curve segments or regions. The “coffee beans” texture can be described by an image-based model: repeated dark ellipses on a lighter background. These models describe equally well an image of texture or an image of a picture of texture. The methods for creating these aggregates were discussed in Chapters 4 and 5. As with all image-based models, three-dimensional phenomena such as occlusion must be handled indirectly. In contrast, structural approaches to texture sometimes require knowledge of the three-dimensional world producing the texture image. One example of this is Brodatz’s “coins” shown in Fig. 6.1. A three-dimensional model of the way coins can be stacked is needed to understand this texture fully.

An important part of the texel definition is that primitives must occur repeatedly inside a given area. The question is: How many times? This can be answered qualitatively by imagining a window that corresponds approximately to our field of view superimposed on a very large textured area. As this window is made smaller, corresponding to moving the viewpoint closer to the texture, fewer and fewer texels are contained in it. At some distance, the image in the window no longer




Fig. 6.3 Examples of texels. (a) Ellipses. (b) Linear segments.

appears textured, or if it does, translation of the window changes the perceived texture drastically. At this point we no longer have a texture. A similar effect occurs if the window is made increasingly larger, corresponding to moving the field of view farther away from the image. At some distance textural details are blurred into continuous tones and repeated elements are no longer visible as the window is translated. (This is the basis for halftone images, which are highly textured patterns meant to be viewed from enough distance to blur the texture.) Thus the idea of an appropriate *resolution*, or the number of texels in a subimage, is an implicit part of our qualitative definition of texture. If the resolution is appropriate, the texture will be apparent and will “look the same” as the field of view is translated across the textured area. Most often the appropriate resolution is not known but must be computed. Often this computation is simpler to carry out than detailed computations characterizing the primitives and hence has been used as a precursor to the latter computations. Figure 6.4 shows such a resolution-like computation, which examines the image for repeating peaks [Connors 1979].

Textures can be hierarchical, the hierarchies corresponding to different resolutions. The “brick wall” texture shows such a hierarchy. At one resolution, the highly structured pattern made by collections of bricks is in evidence; at higher resolution, the variations of the texture of each brick are visible.

### 6.3 STRUCTURAL MODELS OF TEXEL PLACEMENT

Highly patterned textures tessellate the plane in an ordered way, and thus we must understand the different ways in which this can be done. In a regular tessellation the

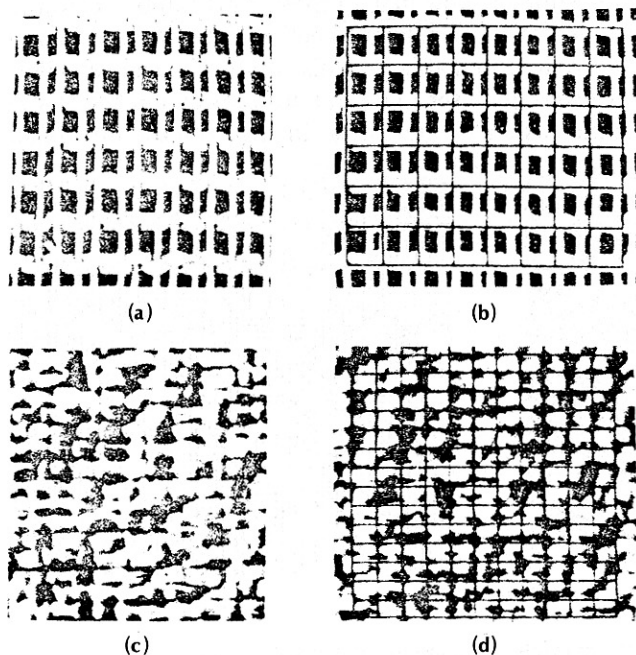




Fig. 6.4 Computing texture resolutions. (a) French canvas. (b) Resolution grid for canvas. (c) Raffia. (d) Grid for raffia.




polygons surrounding a vertex all have the same number of sides. Semiregular tessellations have two kinds of polygons (differing in number of sides) surrounding a vertex. Figure 2.11 depicts the regular tessellations of the plane. There are eight semiregular tessellations of the plane, as shown in Fig. 6.5. These tessellations are conveniently described by listing in order the number of sides of the polygons sur-




(4, 8, 8)




(3, 6, 3, 6)




(3, 4, 6, 4)



(3, 3, 3, 3, 6)



(3, 3, 3, 4, 4)



(3, 3, 4, 3, 4)

**Fig. 6.5** Semiregular tessellations.

rounding each vertex. Thus a hexagonal tessellation is described by (6,6,6) and every vertex in the tessellation of Fig. 6.5 can be denoted by the list (3,12,12). It is important to note that the tessellations of interest are those which describe the *placement* of primitives rather than the primitives themselves. When the primitives define a tessellation, the tessellation describing the primitive placement will be the dual of this graph in the sense of Section 5.4. Figure 6.6 shows these relationships.





Fig. 6.6 The primitive placement tessellation as the dual of the primitive tessellation.

### 6.3.1 Grammatical Models


A powerful way of describing the rules that govern textural structure is through a grammar. A grammar describes how to generate patterns by applying *rewriting rules* to a small number of *symbols*. Through a small number of rules and symbols, the grammar can generate complex textural patterns. Of course, the symbols turn out to be related to texels. The mapping between the stored model prototype texture and an image of texture with real-world variations may be incorporated into the grammar by attaching probabilities to different rules. Grammars with such rules are termed *stochastic* [Fu 1974].

There is no unique grammar for a given texture; in fact, there are usually infinitely many choices for rules and symbols. Thus texture grammars are described as *syntactically ambiguous*. Figure 6.7 shows a syntactically ambiguous texture and two of the possible choices for primitives. This texture is also *semantically ambiguous* [Zucker 1976] in that alternate ridges may be thought of in three dimensions as coming out of or going into the page.

There are many variants of the basic idea of formal grammars and we shall examine three of them: shape grammars, tree grammars, and array grammars. For a basic reference, see [Hopcroft and Ullman 1979]. Shape grammars are distinguished from the other two by having high-level primitives that closely correspond to the shapes in the texture. In the examples of tree grammars and array grammars that we examine, texels are defined as pixels and this makes the



Two choices for primitives:



or




Fig. 6.7 Ambiguous texture.

grammars correspondingly more complicated. A particular texture that can be described in eight rules in a shape grammar requires 85 rules in a tree grammar [Lu and Fu 1978]. The compensating trade-off is that pixels are gratis with the image; considerable processing must be done to derive the more complex primitives used by the shape grammar.

### 6.3.2 Shape Grammars

A shape grammar [Stiny and Gips 1972] is defined as a four-tuple  $\langle V_t, V_m, R, S \rangle$  where:

1.  $V_t$  is a finite set of shapes
2.  $V_m$  is a finite set of shapes such that  $V_t \cap V_m = \phi$
3.  $R$  is a finite set of ordered pairs  $(u, v)$  such that  $u$  is a shape consisting of elements of  $V_t^+$  and  $v$  is a shape consisting of an element of  $V_t^*$  combined with an element of  $V_m^*$
4.  $S$  is a shape consisting of an element of  $V_t^*$  combined with an element of  $V_m^*$ .

Elements of the set  $V_t$  are called terminal shape elements (or terminals). Elements of the set  $V_m$  are called nonterminal shape elements (or markers). The sets  $V_t$  and  $V_m$  must be disjoint. Elements of the set  $V_t^+$  are formed by the finite arrangement of one or more elements of  $V_t$  in which any elements and/or their mirror images may be used a multiple number of times in any location, orientation, or scale. The set  $V_t^* = V_t^+ \cup \{\Lambda\}$ , where  $\Lambda$  is the empty shape. The sets  $V_m^+$  and  $V_m^*$  are defined similarly. Elements  $(u, v)$  of  $R$  are called shape rules and are written  $u \rightarrow v$ .  $u$  is called the left side of the rule;  $v$  the right side of the rule.  $u$  and  $v$  usually are enclosed in identical dashed rectangles to show the correspondence between the two shapes.  $S$  is called the initial shape and normally contains a  $u$  such that there is a  $(u, v)$  which is an element of  $R$ .



A texture is generated from a shape grammar by beginning with the initial shape and repeatedly applying the shape rules. The result of applying a shape rule  $R$  to a given shape  $s$  is another shape, consisting of  $s$  with the right side of  $R$  substituted in  $S$  for an occurrence of the left side of  $R$ . Rule application to a shape proceeds as follows:

1. Find part of the shape that is geometrically similar to the left side of a rule in terms of both terminal elements and nonterminal elements (markers). There must be a one-to-one correspondence between the terminals and markers in the left side of the rule and the terminals and markers in the part of the shape to which the rule is to be applied.
2. Find the geometric transformations (scale, translation, rotation, mirror image) which make the left side of the rule identical to the corresponding part in the shape.
3. Apply those transformations to the right side of the rule.
4. Substitute the transformed right side of the rule for the part of the shape that corresponds to the left side of the rule.

The generation process is terminated when no rule in the grammar can be applied.

As a simple example, one of the many ways of specifying a hexagonal texture  $\{V_t, V_m, R, S\}$  is

$$\begin{aligned} V_t &= \{ \text{hexagon} \} \\ V_m &= \{ \cdot \} \\ R &: \text{hexagon} \rightarrow \text{two hexagons side-by-side}; \text{two hexagons top-to-bottom}; \text{etc.} \\ S &= \{ \text{hexagon} \} \end{aligned} \tag{6.1}$$

Hexagonal textures can be *generated* by the repeated application of the single rule in  $R$ . They can be *recognized* by the application of the rule in the opposite direction to a given texture until the initial shape,  $I$ , is produced. Of course, the rule will generate only hexagonal textures. Similarly, the hexagonal texture in Fig. 6.8a will be recognized but the variants in Fig. 6.8b will not.




Fig. 6.8 Textures to be recognized (see text).

A more difficult example is given by the “reptile” texture. Except for the occasional new rows, a  $(3, 6, 3, 6)$  tessellation of primitives would model this texture exactly. As shown in Fig. 6.9, the new row is introduced when a seven-sided polygon splits into a six-sided polygon and a five-sided polygon. To capture this with a shape grammar, we examine the dual of this graph, which is the primitive placement graph, Fig. 6.9b. This graph provides a simple explanation of how the extra row is created; that is, the diamond pattern splits into two. Notice that the dual graph is composed solely of four-sided polygons but that some vertices are  $(4, 4, 4)$  and some are  $(4, 4, 4, 4, 4)$ . A shape grammar for the dual is shown in Fig. 6.10. The image texture can be obtained by forming the dual of this graph. One further refinement should be added to rules (6) and (7); so that rule (7) is used less often, the appropriate probabilities should be associated with each rule. This would make the grammar stochastic.




Fig. 6.9 (a) The reptile texture. (b) The reptile texture as a  $(3, 6, 3, 6)$  semiregular tessellation with local deformations.

### 6.3.3 Tree Grammars

The symbolic form of a tree grammar is very similar to that of a shape grammar. A grammar

$$G_t = (V_t, V_m, r, R, S)$$

is a tree grammar if

$V_t$  is a set of terminal symbols

$V_m$  is a set of symbols such that

$$V_m \cap V_t = \phi$$

$r : V_t \rightarrow N$  (where  $N$  is the set of nonnegative integers)

is the rank associated with symbols in  $V_t$

$S$  is the start symbol

$R$  is the set of rules of the form

$$X_0 \rightarrow x \quad \text{or} \quad X_0 \rightarrow x$$

$$X_0 \dots X_{r(x)}$$

with  $x$  in  $V_t$  and  $X_0 \dots X_{r(x)}$  in  $V_m$

For a tree grammar to generate arrays of pixels, it is necessary to choose some way of embedding the tree in the array. Figure 6.11 shows two such embeddings.




Fig. 6.10 Shape grammar for the reptile texture.

In the application to texture [Lu and Fu 1978], the notion of pyramids or hierarchical levels of resolution in texture is used. One level describes the placement of repeating patterns in texture windows—a rectangular texel placement tessellation—and another level describes texels in terms of pixels. We shall illus-




Fig. 6.11 Two ways of embedding a tree structure in an array.

trate these ideas with Lu and Fu's grammar for "wire braid." The texture windows are shown in Fig. 6.12a. Each of these can be described by a "sentence" in a second tree grammar. The grammar is given by:

$$G_w = (V_t, V_m, r, R, S)$$

where

$$\begin{aligned}
 V_t &= \{A_1, C_1\} \\
 V_m &= \{X, Y, Z\} \\
 r &= \{0, 1, 2\} \\
 R : X &\rightarrow \begin{array}{c} A_1 \\ \swarrow \quad \searrow \\ X \quad Y \end{array} \quad \text{or } A_1 \\
 Y &\rightarrow \begin{array}{c} C_1 \\ | \\ Z \end{array} \quad \text{or } C_1 \\
 Z &\rightarrow \begin{array}{c} A_1 \\ | \\ Y \end{array} \quad \text{or } A_1
 \end{aligned} \tag{6.2}$$

and the first embedding in Fig. 6.11 is used. The pattern inside each of these windows is specified by another grammatical level:

$$G = (V_t, V_m, r, R, S)$$

where


$$V_t = \{1, 0\}$$

$$V_m = \{A_1, A_2, A_3, A_4, A_5, A_6, A_7, C_1, C_2, C_3, C_4, C_5, C_6, C_7, N_0, N_1, N_2, N_3, N_4\}$$

$$r = \{0, 1, 2\}$$

$$S = \{A_1, C_1\}$$

$R$ :



The application of these rules generates the two different patterns of pixels shown in Fig. 6.13.

### 6.3.4 Array Grammars

Like tree grammars, array grammars use hierarchical levels of resolution [Milgram and Rosenfeld 1971; Rosenfeld 1971]. Array grammars are different from tree grammars in that they do not use the tree-array embedding. Instead, prodigious use of a blank or null symbol is used to make sure the rules are applied in appropriate contexts. A simple array grammar for generating a checkerboard pattern is

$$G = \{V_t, V_n, R\}$$




Fig. 6.12 Texture window and grammar (see text).

where

$V_t = \{0, 1\}$  (corresponding to black and white pixels, respectively)

$V_n = \{b, S\}$

$b$  is a "blank" symbol used to provide context for the application of the rules. Another notational convenience is to use a subscript to denote the orientation of symbols. For example, when describing the rules  $R$  we use

$$0_x b \rightarrow 0_x 1 \quad \text{where } x \text{ is one of } \{U, D, L, R\}$$

to summarize the four rules

$$\begin{array}{l} 0 \rightarrow 0 \\ b \rightarrow 1 \end{array}, \quad \begin{array}{l} b \rightarrow 1 \\ 0 \rightarrow 0 \end{array}, \quad 0b \rightarrow 01, \quad b0 \rightarrow 10$$

Thus the checkerboard rule set is given by

$$R: S \rightarrow 0 \text{ or } 1$$

$$0_x b \rightarrow 0_x 1 \quad x \text{ in } \{U, D, L, R\}$$

$$1_x b \rightarrow 1_x 0$$

A compact encoding of textural patterns [Jayaramamurthy 1979] uses levels of array grammars defined on a pyramid. The terminal symbols of one layer are the start symbols of the next grammatical layer defined lower down in the pyramid. This corresponds nicely to the idea of having one grammar to generate primitives and another to generate the primitive placement tessellations.

As another example, consider the herringbone pattern in Fig. 6.14a, which is composed of  $4 \times 3$  arrays of a particular placement pattern as shown in Fig. 6.14b. The following grammar is sufficient to generate the placement pattern.

$$G_w = \{V_t, V_m, R, S\}$$






Fig. 6.13 Texture generated by tree grammar.

where

$$V_t = \{a\}$$

$$V_n = \{b, S\}$$

$$R: S \rightarrow a$$

$$a_x b \rightarrow a_x a \quad x \text{ in } \{U, D, L, R\}$$

We have not been precise in specifying how the terminal symbol is projected onto the lower level. Assume without loss of generality that it is placed in the upper left-hand corner, the rest of the subarray being initially blank symbols. Thus a simple grammar for the primitive is


$$G_t = \{V_t, V_n, R, S\}$$

	#'	#'	#'	#'
#'	S'	#'	#'	#'
#'	#'	#'	#'	#'
#'	#'	#'	#'	#'

INITIAL ARRAY AT LEVEL 1

$\alpha'$	$\alpha'$	$\alpha'$	$\alpha'$
$\alpha'$	$\alpha'$	$\alpha'$	$\alpha'$
$\alpha'$	$\alpha'$	$\alpha'$	$\alpha'$
$\alpha'$	$\alpha'$	$\alpha'$	$\alpha'$

TERMINAL ARRAY AT LEVEL 1



FINAL ARRAY

Fig. 6.14 Steps in generating a herringbone texture with an array grammar.

where

$$V_t = \{0, 1\}$$

$$V_n = \{a, b\}$$

$$R: \begin{matrix} a & b & b & b \\ b & b & b & b \\ b & b & b & b \end{matrix} \rightarrow \begin{matrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{matrix}$$

## 6.4 TEXTURE AS A PATTERN RECOGNITION PROBLEM

Many textures do not have the nice geometrical regularity of "reptile" or "wire braid"; instead, they exhibit variations that are not satisfactorily described by shapes, but are best described by statistical models. *Statistical pattern recognition* is a paradigm that can classify statistical variations in patterns. (There are other statistical methods of describing texture [Pratt et al. 1981], but we will focus on statistical pattern recognition since it is the most widely used for computer vision purposes.) There is a voluminous literature on pattern recognition, including several excellent texts (e.g., [Fu 1968; Tou and Gonzalez 1974; Fukunaga 1972], and the ideas have much wider application than their use here, but they seem particularly appropriate for low-resolution textures, such as those seen in aerial images [Weszka et al. 1976]. The pattern recognition approach to the problem is to classify instances of a texture in an image into a set of classes. For example, given the textures in Fig. 6.15, the choice might be between the classes "orchard," "field," "residential," "water."

The basic notion of pattern recognition is the *feature vector*. The feature vector  $\mathbf{v}$  is a set of measurements  $\{v_1 \cdots v_m\}$  which is supposed to condense the description of relevant properties of the textured image into a small, Euclidean *feature space* of  $m$  dimensions. Each point in feature space represents a value for the feature vector applied to a different image (or subimage) of texture. The measurement values for a feature should be correlated with its class membership. Figure 6.16 shows a two-dimensional space in which the features exhibit the desired correlation property. Feature vector values cluster according to the texture from which they were derived. Figure 6.16 shows a bad choice of features (measurements) which does not separate the different classes.

The pattern recognition paradigm divides the problem into two phases: training and test. Usually, during a training phase, feature vectors from known samples are used to partition feature space into regions representing the different classes. However, self teaching can be done; the classifier derives its own partitions. Feature selection can be based on parametric or nonparametric models of the distributions of points in feature space. In the former case, analytic solutions are sometimes available. In the latter, feature vectors are *clustered* into groups which are taken to indicate partitions. During a test phase the feature-space partitions are used to classify feature vectors from unknown samples. Figure 6.17 shows this process.

Given that the data are reasonably well behaved, there are many methods for clustering feature vectors [Fukunaga 1972; Tou and Gonzales 1974; Fu 1974].