PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Shape from Local and Global Analysis of Texture

James V. Stone

Phil. Trans. R. Soc. Lond. B 1993 **339**, 53-65
doi: 10.1098/rstb.1993.0004

Shape from local and global analysis of texture

JAMES V. STONE

Cognitive and Computing Studies, University of Sussex, Brighton BN1 9QY, U.K.

SUMMARY

We demonstrate methods for estimating shape from texture†, for textures which are neither isotropic nor homogeneous. Whereas earlier work assumed either isotropy or homogeneity, we make a more general assumption: the local tangent distribution of the surface texture is invariant with respect to position. We then present a taxonomy of texture types, and a set of methods that can be used to recover surface orientation for each texture type.

We define two classes of methods for estimating shape from texture: invariance-seeking methods, and value-seeking methods. These definitions are used to account for the performance of new methods on different classes of texture. The approach outlined in this paper provides a general strategy for generating methods for obtaining shape from texture for known classes of texture.

1. INTRODUCTION

Established methods for estimating shape from texture assume homogeneity (Aloimonos 1988; Kanatani & Chou 1989), or they assume isotropy (Witkin 1981; Kube 1986; Blake & Marinos 1990; Brown & Shvayster 1990), but not both. (A homogeneous texture has the same 'amount of texture' per unit surface area. A texture that is isotropic has the same proportion of texture elements at each orientation.) Methods that rely on isotropy use data from, and provide estimates for, only a single image region of a given image. It should be noted that these methods cannot determine the sign of surface tilt. Such methods cannot take advantage of the information afforded by gradients of texture (Gibson 1979) observed under perspective projection. In contrast, methods that rely on homogeneity depend on the presence of gradients of image texture to estimate surface orientation, and do not, in general, require an isotropic surface texture. These methods can integrate data from different regions of an image in order to provide an estimate of surface orientation for each of those regions. In contrast to those methods that assume isotropy, these methods provide estimates which include the sign of surface tilt.

In § 1*a*,*b* we describe two sets of methods for estimating shape from texture. The first set consists of methods that assume homogeneity, and the second set consists of methods that assume isotropy. In § 1*c* we show that these sets are special cases of two classes which utilise more general assumptions than homogeneity or isotropy. In subsequent sections new methods associated with these two classes are described and tested on images of planar textured surfaces.

† The term 'shape from texture', as it is commonly used in the literature, is somewhat misleading. It is used, as we do here, to describe methods for estimating the orientations of planar surfaces (see, for example, Aloimonos (1988); Kube (1986); Brown & Shvayster (1990)) and it is assumed that arbitrary shapes can be approximated by collections of planar facets.

(*a*) *Homogeneity*

Textural homogeneity has played a critical role in the development of ideas (Gibson 1979) and computational methods (Aloimonos 1988; Kanatani & Chou 1989; Stone 1990, 1992; Stone & Isard 1993) for deriving estimates of surface orientation under perspective projection. There are at least two distinct working definitions of homogeneity. In Aloimonos (1988) it is defined in terms of the length of lines detected per unit surface area, while Kanatani & Chou (1989) gives a more abstract definition of which length per unit surface area is a special case.

If a surface texture is not homogeneous then it is difficult to separate the inhomogeneities of an image texture that are due to perspective projection, and those due to textural inhomogeneity on the surface. The basic strategy for taking advantage of textural homogeneity is to find surface orientation parameters that minimize the inhomogeneity of the back-projected surface texture. Clearly, this approach fails when the surface texture is not approximately homogeneous.

(*b*) *Isotropy and the local tangent distribution*

One of the earliest demonstrations of estimating surface orientation from texture is given in Witkin (1981). Witkin assumed that the surface texture is isotropic and that the image is formed by orthographic projection. The basic strategy used by Witkin and all subsequent methods in this class is as follows. The projection of an isotropic surface texture generates an anisotropic image texture (if the surface is not fronto-parallel). Under orthographic projection, the degree of image anisotropy is a function only of the surface orientation. Therefore, the estimated surface orientation is given by those surface orientation parameters which minimise a measure of the anisotropy of the back-projected image texture.

Phil. Trans. R. Soc. Lond. B (1993) **339,** 53–65
Printed in Great Britain

The basic strategy used by Witkin has been re-implemented many times by making use of different metrics of anisotropy, such as the two-dimensional autocorrelation function (Brown & Shvayster 1990), or the output of Gabor filters (Kube 1986; Turner *et al.* 1991). All of these methods attempt to find surface parameters which generate a minimally anisotropic estimated surface texture, and all (except the method described in Gårding (1991)) assume orthographic projection.

(*c*) *Invariance-seeking and value-seeking methods*

The methods described in the last two sub-sections exemplify two distinct strategies for estimating shape from texture. The first (exemplified in § 1*a*) assumes that the surface texture is invariant with respect to a specified property (e.g. density). This strategy is consistent with the approach popularized by Gibson (1979) for utilizing changes‡ in two-dimensional images in order to estimate corresponding values of physical invariances (e.g. surface orientation). We define methods that use this strategy as invariance-seeking methods, because they only require that a specified property of a three-dimensional scene remains invariant.

In contrast, the second strategy (exemplified in § 1*b*) does not require that the surface texture is invariant with respect to a specified property. Instead, it assumes that the value of a specified property (e.g. anisotropy) is known. We define methods that use this strategy as value-seeking methods. Methods that use this strategy adjust their estimate of a specified parameter (e.g. surface orientation) in order to maximize the agreement between the assumed value of a property (e.g. anisotropy = 0) and the estimated value of that property based on a set of image data and the current estimate of the specified parameter (e.g. surface orientation). For example, in Witkin (1981) it is assumed that the degree of textural anisotropy is zero. Thus, zero is the value assumed to be associated with the degree of anisotropy of the surface texture. The method adjusts the estimated surface orientation so as to maximize the agreement between the assumed degree (= zero) of surface texture anisotropy and the estimated degree of surface anisotropy (where the latter is a function of the current estimate of the surface orientaton). Such methods fall outside of the framework constructed by Gibson (1979) for making use of physical invariances (e.g. textural homogeneity) associated with the gradients of imaged texture. In contrast, in Aloimonos (1988) it is assumed that the amount of texture per unit surface area is invariant with respect to position on the surface. In this case, density is the property that is assumed to be invariant, and the method should work equally well (in principle) for any value of density of the surface texture.

The methods described in subsequent sections are presented in an order which is inversely related to the amount of 'knowledge' each method requires regarding the nature of surface texture to be analysed. Accordingly, we begin by describing a set of methods which are value-seeking, and then a larger set of methods which are invariance-seeking. The methods become increasingly general (across and within these two sets of methods) in terms of the types of texture they can successfully analyse. Additionally, images of textured surfaces are presented in an order which is approximately proportional to the number of degrees of freedom associated with the surface texture in each image. Thus, all methods perform well on the 'constrained' texture of figure 1, an increasing proportion of the images are successfully analysed by subsequent methods, and only the final method (IMTD-MV) described performs well on the 'unconstrained' texture shown in figure 8.

‡ Either within a single image, or changes associated with a single image position across temporally contiguous images.

(*d*) *Test images*

Before defining more formally a taxonomy of textures and computational methods associated with each, table 1 lists the properties (defined in later sections) of each of the test images used here. Note that each of the images was generated from a planar surface using perspective projection. The apparent non-planarity of a given surface is due only to the statistics of the texture on that planar surface.

In each results table the estimated surface orientation is given as a pair of gradient space parameters (P,Q) (see Appendix 1), and the error is given as the angle, in degrees, between the estimated and actual surface normal. The actual surface orientation of all surfaces shown here is $(P,Q) = (-0.287, 0.789)$, and the focal length is equal to the image height.

For each of the methods described below the estimated surface orientation is given by a (P,Q) pair which minimizes a particular cost function. For each method, evaluation of the associated cost function involves back-projecting image data onto the estimated surface (see Appendix 2). The cost function is minimized by making use of the 'amoeba' search method described in Nelder & Mead (1965).

2. METRICS OF ANISOTROPY

A set of edge orientations, or tangents, can be plotted as a histogram with the abscissa representing edge orientation α and the ordinate representing the number of edges, or length S, at each orientation. An isotropic set of tangents defines a histogram that

Table 1. *Properties of textures shown in figures 1–8*

figure	isotropic	weakly isotropic	homotropic	weakly homotropic	homogeneous
1,2	yes	yes	yes	yes undef	yes
3,4	no	no	yes	yes	yes
5	yes	yes	yes	yes	no
6	no	no	yes	yes	no
7	no	no	no	yes	yes
8	no	no	no	yes	no

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

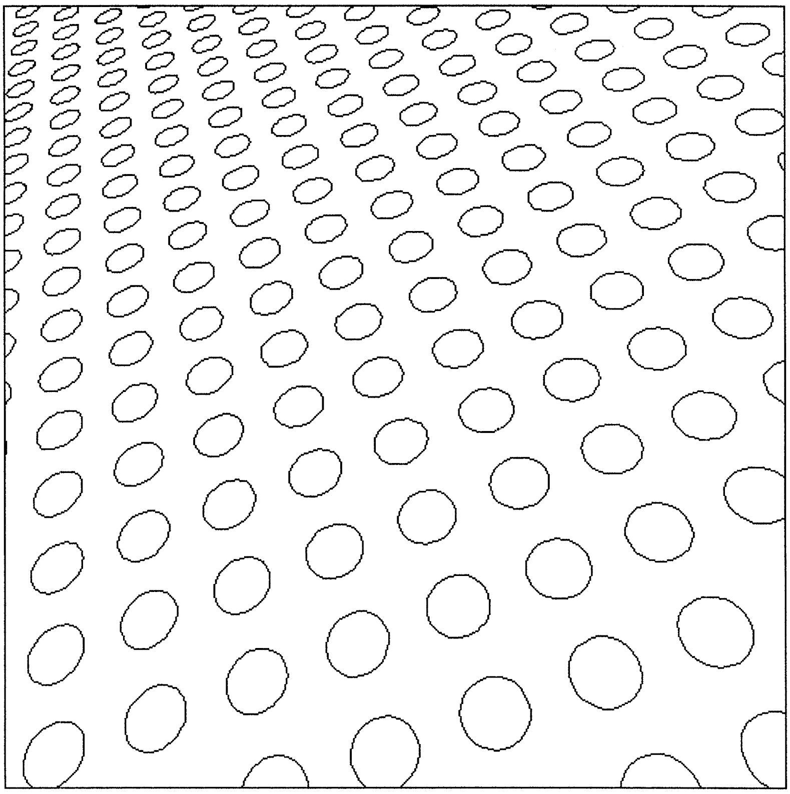

Figure 1. Isotropic, homogeneous texture.

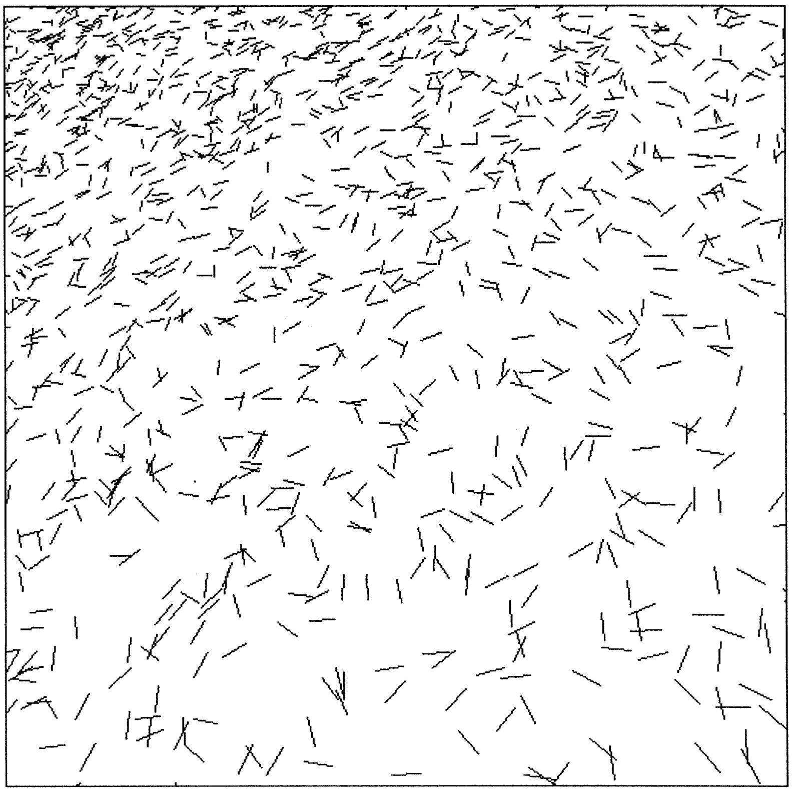

Figure 2. Isotropic, homogeneous texture.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

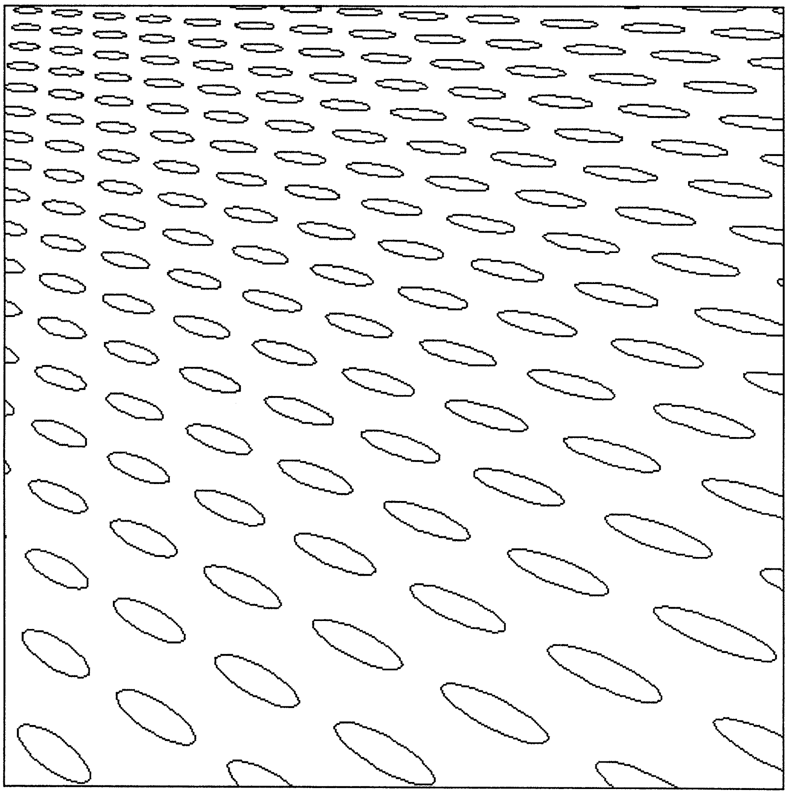

Figure 3. Homotropic, homogeneous texture.

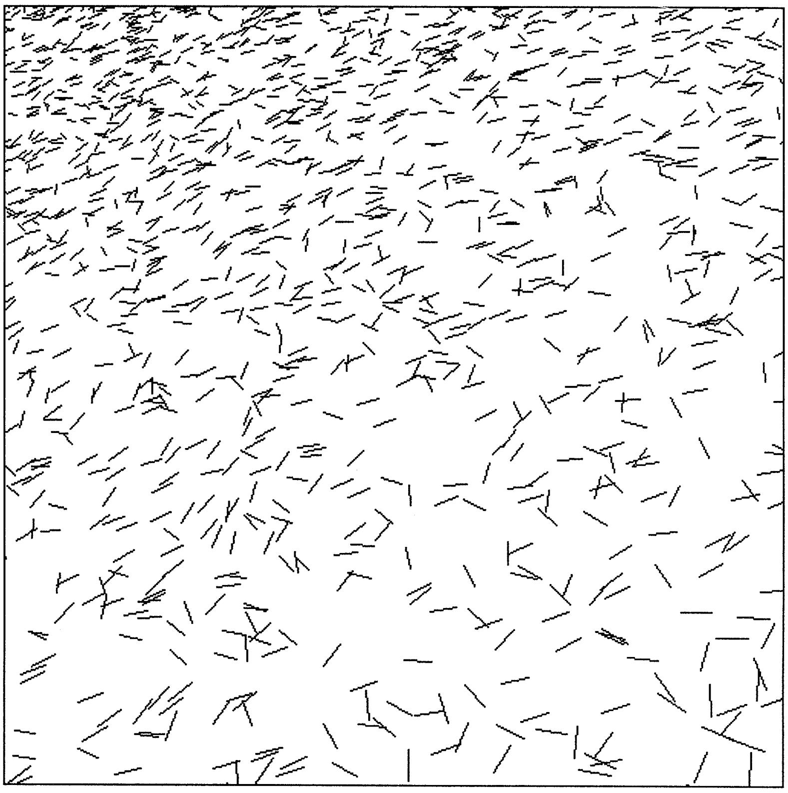

Figure 4. Homotropic, homogeneous texture, with elliptical tangent distribution.

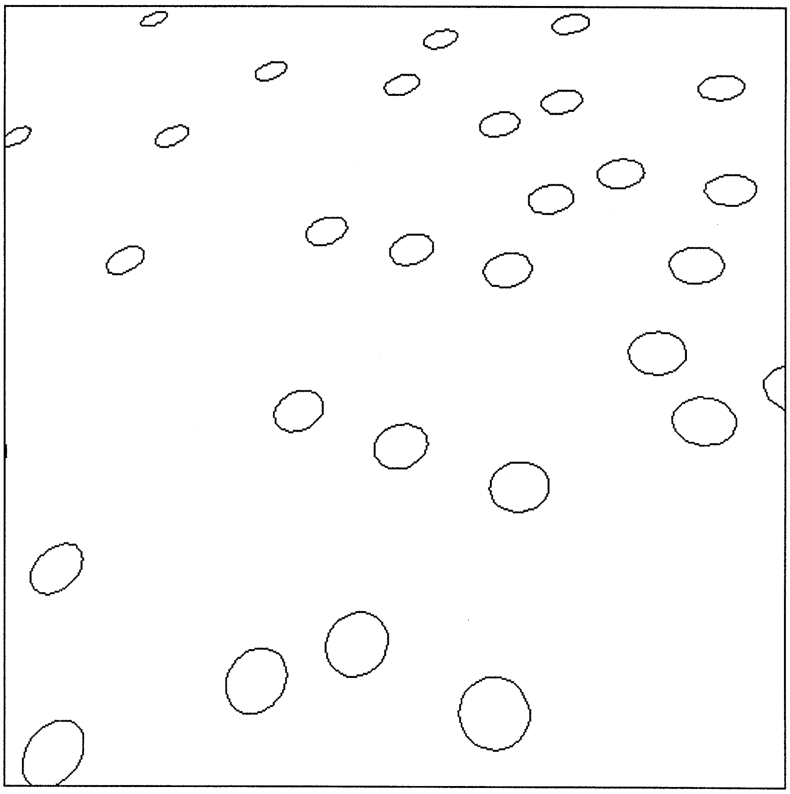

Figure 5. Isotropic, inhomogeneous texture.

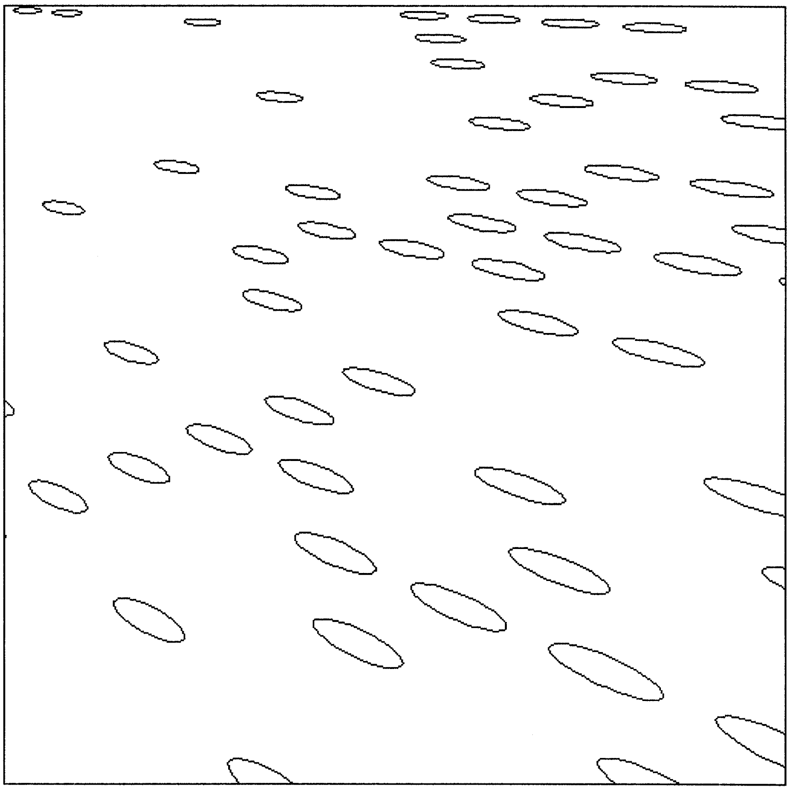

Figure 6. Anisotropic, inhomogeneous, homotropic texture.

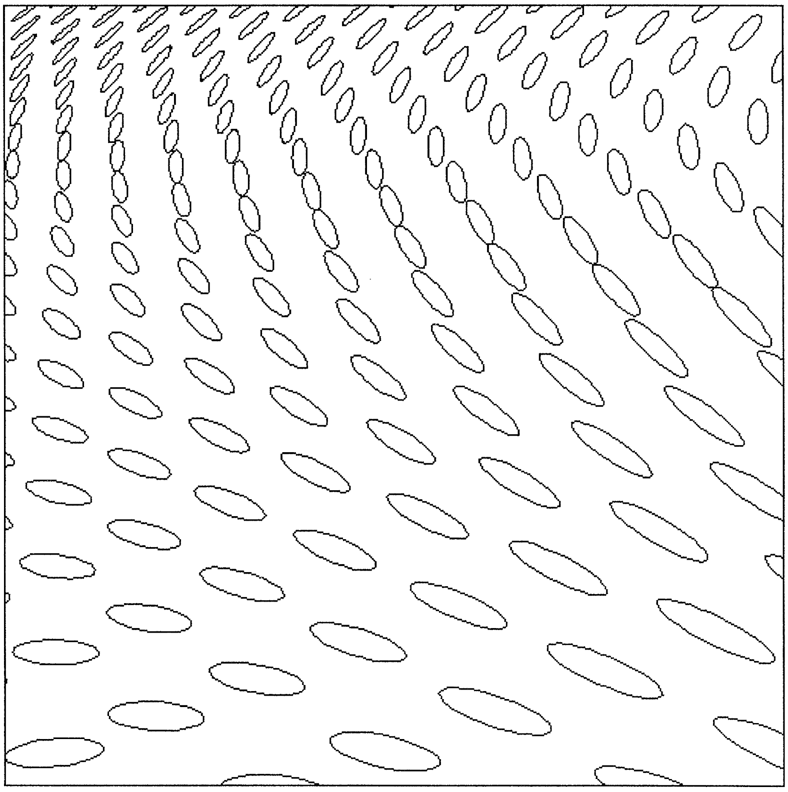

Figure 7. Anisotropic, homogeneous, weakly homotropic texture.

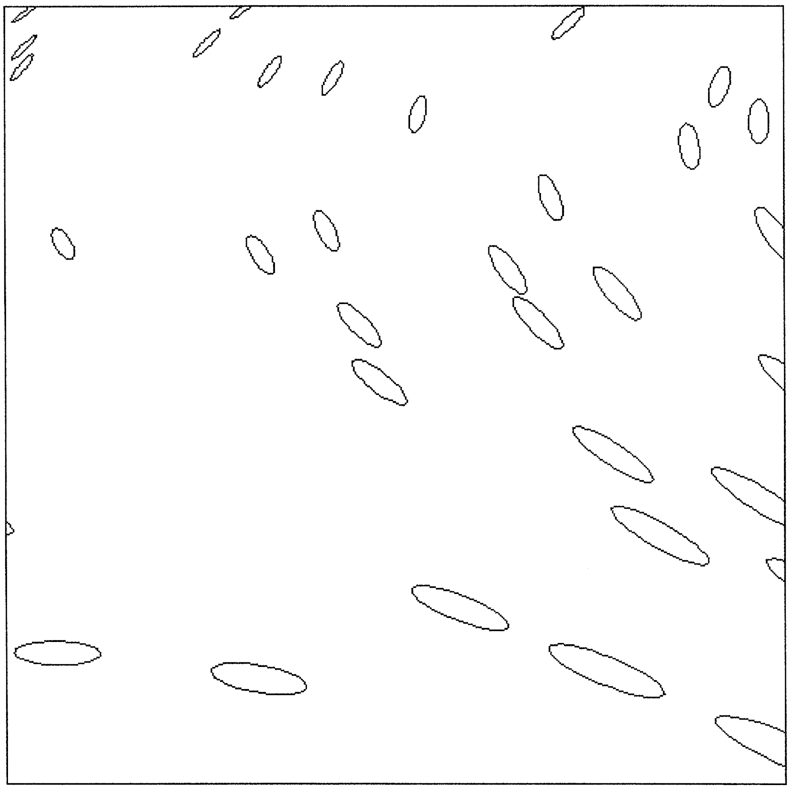

Figure 8. Anisotropic, inhomogeneous, weakly homotropic texture.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

consists of a horizontal line. Deviations from isotropy produce corresponding deviations in such a line. However, it is not obvious how such deviations could be used to index the degree of anisotropy of a given set of edges.

A more intuitively appealing representation can be obtained by plotting a tangent set in polar coordinates S,α. Thus a point (S_1,α_1) is plotted S_1 units from the origin, and describes a line at α_1 degrees from horizontal. If an isotropic set of tangents is plotted in polar coordinates then a semi-circle is observed. This is because there are only 180 degrees of orientation. This can be remedied by re-plotting the original data set in the range 181–360 degrees to produce a circular polar plot. Alternatively, the length at each orientation, α can be plotted with direction 2α (Gårding 1991). A circular polar plot can only be obtained from a data set that is isotropic, and deviations from isotropy produce corresponding changes in the shape of the polar plot. Re-plotting a histogram in polar coordinates suggests ways of defining metrics for the degree of anisotropy which are not apparent using the conventional histogram plot. Two of these definitions are described next.

For a given continuous tangent distribution G the degree of anisotropy can be measured as the area of the corresponding polar plot:

$$J = \int_0^\pi G(\alpha)^2\,\mathrm{d}\alpha. \tag{1}$$

Where $G(\alpha)$ is the length of components at orientation α. J is a measure of the degree of anisotropy associated with a set of tangents, and is minimal if these tangents are distributed isotropically. However, the magnitude of J depends on the total length of components in G. This is undesirable if the relative degree of anisotropy of two histograms with different areas (integrals) is required§. To modifify J so that it does not depend on the total length of components in J we first define a normalized distribution F, such that:

$$F(\alpha) = \frac{G(\alpha)}{\int_0^\pi G(\alpha)\,\mathrm{d}\alpha}, \tag{2}$$

and a corresponding normalized version of J:

$$I = \int_0^\pi F(\alpha)^2\,\mathrm{d}\alpha. \tag{3}$$

For isotropic textures $I = 1/\pi$. As before, any deviation from isotropy is reflected by an increase in I. I is independent of the amount of texture used to construct a histogram, and therefore of the size of the image region from which the texture was derived. Finally it should be noted that I is invariant with respect to scale, so that certain of the methods (e.g. ITD, IMTD) described below may be applicable to fractal-like textures.

§ If the degree of anisotropy associated with two image regions which contain different amounts of texture is required, for instance.

In contrast to I, weak isotropy is defined in (Gårding 1991) in terms of the mean vector (MV) $V = (C,S)$ of a tangent distribution. For a distribution F the components of its MV are:

$$C = \int_0^\pi F(\alpha)\cos 2\alpha\,\mathrm{d}\alpha,$$

$$S = \int_0^\pi F(\alpha)\sin 2\alpha\,\mathrm{d}\alpha. \tag{4}$$

With $F(\alpha)$ being subject to the same normalization as in equation (2). The length $|V| = (C^2 + S^2)^{1/2}$ of V is normalized in equation (4) so that $(0 <= |V| <= 1)$. If $|V| = 0$ then the texture is defined in (Gårding 1991) as being weakly isotropic.

Note that all textures that are isotropic are also weakly isotropic. Textures that consist of regular polygons are not isotropic, but they are weakly isotropic. Both I and $|V|$ are measures of the degree of anisotropy of a tangent distribution; both have a minimal value for isotropic distributions, and a maximal value of distributions that consist of a single non-zero length component. Even though the quantities I and $|V|$ are both minimal for isotropic distributions, a plot of I and $|V|$ versus 'degree of anisotropy' would not, in general, display two parallel curves. The quantity $|V| = 0$ for any distribution F in which every component length $F(\alpha) = F(\alpha + \pi/2)$. This is because, in evaluating the discrete version of equation (4), two such components yield $F(\alpha)\sin(2\alpha) + F(\alpha + \pi/2)\sin(2\alpha + \pi) = 0$, and $F(\alpha)\cos(2\alpha) + F(\alpha + \pi/2)\cos(2\alpha + \pi) = 0$, as contributions to C and S, respectively. In contrast, no such cancellation of terms occurs with I, so that it is possible for I to be non-minimal for distributions which make $|V| = 0$. The converse does not occur. Results (not reported here) indicate that substituting I for $|V|$in methods described in this paper usually yields only a small difference in the estimated surface orientation.

(*a*) *Using metrics of anisotropy to estimate surface orientation*

We can use metrics of anisotropy to estimate surface orientation either by using the strategy described by Witkin for data from a single image region, or by finding a 'best-fit' solution for data from multiple image regions. Both of these strategies give rise to methods which are of the value-seeking class of methods.

(i) *Isotropy and the MLA methods*

These methods represent implementations of the strategy originally described in (Witkin 1981). For a given image region the estimated surface orientation is given by the value of surface orientation parameters (P,Q) that minimizes $|V|$, or I, of the back-projected image distribution. These methods do not require data from multiple image regions, and the 'image region' for which results are shown here is the whole image. Minimizing $|V|$ in order to estimate surface orientation was originally described in Gårding (1991), although the method used to minimize $|V|$

Table 2. *Results for the* MLA2 *method*

figure	result (P,Q)	error/deg.
1	−0.294, 0.789	0.294
2	−0.330, 0.766	2.143
3	0.158, 1.640	26.709
4	−0.380, 1.173	11.072
5	−0.294, 0.773	0.646
6	0.144, 1.670	26.751
7	0.540, 0.324	44.473
8	0.616, 0.570	41.681

here is substantially different from that used by Gårding.

In the case of an isotropic texture and a single image region there are two (P,Q) pairs that minimize $|V|$ and I, for the same reason that there are two surface circles that could have generated any image ellipse. These two solutions are symmetric, so that having found one solution, the other is easy to locate. The second solution differs from the first only inasmuch as the direction of tilt associated with it has a different sign. We refer to methods which minimize I and $|V|$ as the minimal local anisotropy (MLA) methods. MLA1 minimizes I, and MLA2 minimizes $|V|$.

For both methods the estimated surface tangent distribution is obtained by back-projecting image edges onto the estimated surface, as described in Appendix 2. The quantities I and $|V|$ are computed by using this discrete version of equations (3) and (4), respectively. The final solutions are found by minimizing $|V|$ (I) with respect to (P,Q).

It has been demonstrated in (Gårding 1991; Stone 1991) that surface orientation parameters that minimize $|V|$ and I (respectively) provide accurate estimates for isotropic textures. Note that, even though these methods assume perspective projection, they require image data from, and provide two surface orientation estimates for, each image region analysed. These methods sit squarely in the Witkin (value-seeking) framework, in that they require a particular type of tangent distribution to be present on a surface, and do not take account of any invariances on the surface (such as textural density).

Results for the MLA2 method are shown in table 2. Results for MLA1 are similar to those for MLA2 and are not shown. As expected the MLA2 method performs well on isotropic textures (figures 1, 2 and 5), even if these textures are not homogeneous (see figure 5).

Table 3. *Results for the* MGA2 *method*

figure	result (P,Q)		error/deg.
1	−0.282,	0.787	0.214
2	−0.318,	0.731	2.630
3	0.192,	1.660	27.570
4	−0.411,	1.252	12.895
5	−0.292,	0.794	0.256
6	0.222,	1.726	28.765
7	−0.583,	−1.365	87.751
8	−1.184,	−1.487	86.570

(ii) *Isotropy and the* MGA *methods*

A unique solution for isotropic textures can be obtained under perspective projection by making use of data from multiple image regions. This is achieved by measuring the degree of anisotropy associated with each of a number of image regions $r_1 .. r_n$ and then finding those surface orientation parameters that minimize the total anisotropy over the corresponding set of surface regions. This strategy can be implemented using I or $|V|$ as a measure of anisotropy. If we use I then minimizing:

$$E_I = \sum_{i=1}^{n} I_i, \tag{5}$$

yields an estimate of the surface orientation. Each I_i is associated with a different region R_i. The method that minimizes equation (5) is called MGA1. Alternatively, if we use $|V|$ then minimizing:

$$E_{|V|} = \sum_{i=1}^{n} |V_i|, \tag{6}$$

yields an estimate of the surface orientation. The method that minimizes equation (6) is called MGA2. These methods minimize the total anisotropy over a number of surface regions, and are referred to as the minimal global anisotropy (MGA) methods. Minimizing equation (6) forces each region to have the same estimated surface orientation as all of the other regions. For non-planar surfaces it is possible to regularize (6) by introducing a smoothing term, and allowing degrees of surface smoothness to be 'traded' for degrees of textural anisotropy. However, for the planar surfaces considered here this is not necessary.

The MGA methods make use of more than one instance of a tangent distribution so that it might be expected that these methods provide a more accurate estimate of surface orientation than the MLA methods. Additionally, the MGA methods 'trade off' estimated amounts of anisotropy in one surface region with amounts of anisotropy in another surface region. These methods are therefore likely to produce different results from those obtained by simply minimizing anisotropy for a single surface region. Note that, even though these methods use data from more than one region, this is done in order to find a minimally anisotropic surface texture, and not because these methods make use of any invariance of the surface texture.

Results for method MGA2 are shown in table 3. Results for MGA1 are similar to those for MGA2 and are not shown here. Results for the isotropic texture of figure 1 are similar to result using the MLA2 method. Like the MLA2 method, the MGA methods perform well only on the isotropic texture of figures 1, 2 and 5.

3. METRICS OF TEXTURAL INVARIANCE

The methods described so far do not assume that the surface texture is invariant, but only that it is isotropic. These methods are not the invariance-seeking mechanisms implied by Gibson. Rather, they are mechanisms which depend on the presence of a known (isotropic) texture. In this section we introduce methods which depend only on the presence of particular types of invariances on a surface.

The basic strategy is as follows. Define a metric M (e.g. textural density, the MV of a tangent distribution) such that measures, $T_1 = T_2 = .. T_n$, of M can reasonably be expected to be invariant on a textured surface. Define a set $r_1 .. r_n$ of image regions. This set of regions defines a corresponding set $R_1 .. R_n$ of surface regions. For a given estimate (P', Q') of the surface orientation (P, Q) we can then compute an estimate T'_i of the actual surface quantity T_i associated with each R_i. Each T'_i is computed from the back-projected texture of an image region r_i. The surface orientation is then given by those surface orientation parameter valves which make elements of the set $T'_i .. T'_n$ equal to each other.

This strategy is implemented for all remaining methods (except Aloimonos' method) described in this paper as follows. The image is divided into a number of square image regions (see Appendix 3). Given a current estimate of (P, Q), data from each region is back-projected onto the estimated surface (see Appendix 2). The back-projected data from each image region is used to construct an edge-orientaton histogram (tangent distribution) for that surface region. This tangent distribution represents the discrete analogue of $F(\alpha)$, and can be used to compute other statistics associated with the local surface tangent distribution. These statistics contribute to the cost-function associated with a given method, allowing the current estimate of (P, Q) to be evaluated in terms of this cost-function.

A standard method which assumes textural invariance (homogeneity) is described in (Aloimonos 1988)). Results for this method (see table 5) provide a baseline for the invariance-seeking class of methods described in the rest of this paper. This method depends only on an assumption of textural homogeneity, as indexed by the mean deviation|| in the estimated surface texture density associated with a set of image regions¶. The method used to find surface orientation parameters which minimise the mean deviation of the surface texture density is the 'amoeba' method described in (Nelder & Mead 1965). Aloimonos' method differs from that described in the previous paragraph in two respects only. First, the tangent distribution for a given region is not required. Second, a method for estimating the surface area associated with each image region is required in order to compute the local texture density on the surface (see Aloimonos 1988). As expected, the method of Aloimonos gives inaccurate results only for surface textures that are not homogeneous (figures 5, 6 and 8).

(*a*) *Homotropy*

We define homotropy to provide a measure of the variability in the shapes and principal orientations of a set of tangent distributions. However, before providing a definition of homotropy it is necessary to define the shape and principal orientation of a tangent distribution. The shape of a distribution in polar coordinates is given by the closed curve described by F between 0 and 2π. The principal orientation of a tangent distribution is the orientation of maximum variation in F (in polar coordinates). It is the orientation of the longest eigenvector of the covariance matrix of the locus of points defined by F. This can be visualized by fitting an ellipse to F in polar coordinates, and identifying the orientation of the major axis of the ellipse. The principal orientation is not defined only for isotropic textures, because the corresponding fitted ellipses are circles.

If tangent distributions have the same shape and principal orientation (but are not necessarily isotropic) everywhere on a surface then there exist surface orientation parameters for which the back-projection of a set of image regions minimizes a measure of variability (of shape and principal orientation) over the corresponding set of surface tangent distributions. If the tangent distributions associated with a set of surface regions all have the same shape and share a common principal orientation then the texture is homotropic. One measure of the degree of inhomotropy of a set of distributions can be defined in terms of the mean vectors of that set. If a texture is homotropic then $V_1 = V_2 = .. V_n$ for any set $R_1 .. R_n$ of surface regions, where V_i is the MV associated with region R_i.

Note that it is possible for V to be the same everywhere on a textured surface, but for the shape of the associated tangent distributions to vary.

(i) *Homotropic textures and value-seeking methods*

The perspective images of homotropic, anisotropic textures are shown in figures 3, 4 and 6 (figures 1–6 are all homotropic, but only figures 3, 4 and 6 are also anisotropic). Clearly, attempts to analyse an image of this type of texture using Witkin's strategy, or the MLA methods, will fail. They will fail because each of these methods attempts to find surface parameter values which back-project the image tangent distribution to an isotropic surface distribution. A similar argument exists for the MGA methods, which attempt to minimize the summed degree of anisotropy (as indexed by I) of a set of back-projected tangent distributions. The degree to which these methods fail is a positive

|| The definition of mean deviation used in this paper is

$$(1/n)\sum_{i=1}^{n} \mathrm{abs}(\bar{x} - x_i),$$

where x is the estimated surface texture density, for example.

¶ See Appendix 3 for a description of how these regions are defined.

Table 4. *Results for Aloimonos' method*

figure	result (P,Q)		error/deg.
1	−0.276,	0.813	1.054
2	−0.338,	0.772	2.349
3	−0.289,	0.799	0.376
4	−0.420,	0.693	7.200
5	0.753,	0.397	50.205
6	0.871,	0.398	53.936
7	−0.297,	0.791	0.404
8	0.759,	−0.002	61.526

function of the degree of anisotropy of the surface texture.

Note that, if the degree of surface anisotropy is known, then these methods can be made to work by minimizing the difference between the degree of anisotropy of the back-projected tangent distribution and the known degree of anisotropy of the surface texture. However, such knowledge is not usually available, and methods which embody more general assumptions (as described next) can be used in these cases.

(iii) *Homotropic textures and the* ITD *method*

By extending the strategies associated with homogeneity and the global analysis of texture (e.g. Aloimonos 1988) to deal with quantities that are usually associated with the local analysis of texture (e.g. tangent distribution) we can obtain exact solutions for homotropic (and even non-homotropic) surface textures. If a surface texture is homotropic then the MVs of surface regions are identical (the converse does not, in general, hold). Therefore surface orientation parameter values that make all of the estimated mean surface vectors identical provide an estimate of surface orientation. It is unrealistic to expect the MV associated with each surface region to be identical to all other MVs, so that a metric of the variability of the estimated MVs is required. Minimizing such a metric provides a 'best' estimate of surface orientation. We define the metric of variability of a set of MVs as the total distance of all MVs from the average vector of the set:

$$E_V = \sum_{i=1}^{n} |\bar{V} - V_i|. \qquad (7)$$

Where each V_i is associated with a surface region R_i. This method minimizes a measure of the variability in the set of surface tangent distributions, and is referred to as the invariant tangent distribution (ITD) method.

The performance of this method, which is based on the assumption of an invariance (homotropy) of the surface texture, cannot (in general) be equaled by methods that depend only on the local analysis of textures which are assumed to be isotropic.

Even though the ITD method depends on the presence of an invariance on the surface, that invariance is defined with respect to the local tangent distribution. Thus, unlike other methods that depend on a textural invariance (e.g. Aloimonos 1988; Kanatani & Chou 1989), the particular invariance required is not defined in terms of the density of the surface texture. It is therefore possible to analyse textured surfaces that are neither isotropic nor homogeneous.

Results for the ITD method are shown in table 5. As expected the method works well for all textures which are homotropic (figure 1–6), even if they are not homogeneous. This is because the estimated tangent distributions are normalized so that $(0 \leqslant |V| \leqslant 1)$, irrespective of the local density of texture. In this respect the ITD method is at odds with human perceptions of these images. The non-homogeneous, homotropic texture of figure 6 does not give a compelling impression of a planar surface. It is as if the ITD method, in order to correctly estimate the orientation of the surface in figure 6, invokes assumptions about the surface texture that human observers do not.

The large errors for figures 7 and 8 are not surprising given that this method relies on an assumption of homotropy, whereas the textures shown in these figures are only weakly homotropic (defined in the next sub-section). That is, the ITD method assumes that the tangent distributions in a given image have the same shape and principal orientation, whereas the tangent distributions of textures shown in figures 7 and 8 have invariant shapes, but not invariant principal orientations. Not only are the principal orientations variable, they vary systematically across each image so that they appear to mimic a pattern of principal orientations which would be generated by a curved surface. Therefore it is not surprising that a method (such as ITD) which relies on the invariance of the principal orientation of local tangent distributions performs poorly on such images.

With the ITD method we have demonstrated a method that produces unintuitive results, inasmuch as this method provides accurate estimates of surfaces which human observers see as being curved. Moreover, these estimates are associated with minimal values for E_x (where x specifies a method), implying that the data are consistent with planar surfaces, and not just that these methods give good results only because they assume planarity; that is, methods which attempted to fit these data to a curved surface model would have to violate the minimal assumptions upon which the methods presented here are based.

(*b*) *Weak homotropy*

We can further weaken constraints associated with a texture by specifying that only the shape of polar plots (not necessarily the principal orientations) of a set of tangent distributions is invariant. We define such a texture to be weakly homotropic. Weakly homotropic textures are shown in figures 7–8. It is worth reminding the reader that if a texture is homogeneous then we do not have to rely on the rather minimal assumption of weak homotropy in order to recover surface orientation. Thus Aloimonos' methods provides an accurate result for the homogeneous, weakly homotropic texture of figure 7, but not the inhomogeneous, weakly homotropic texture of figure 8.

Table 5. *Results for the* ITD *method*

figure	result (P,Q)	error/deg.
1	−0.280, 0.796	0.430
2	−0.297, 1.039	7.757
3	−0.295, 0.823	1.190
4	0.311, 0.892	3.415
5	−0.291, 0.793	0.202
6	−0.370, 0.511	12.106
7	0.771, −1.724	101.997
8	0.906, −1.488	99.503

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Table 6. *Results for the* IMTD-MV *method*

figure	result (*P*,*Q*)	error/deg.
1	−0.307, 0.825	1.393
2	−0.224, 0.958	6.471
3	−0.316, 0.783	1.30
4	−0.449, 0.988	7.957
5	−0.455, 0.845	6.730
6	−0.576, 0.797	11.541
7	−0.276, 0.909	4.097
8	−0.264, 0.905	4.155

(i) *Weakly homotropic textures and the* IMTD *methods*

Both methods described here depend on the invariance of measures of the moments of tangent distributions, and are thus referred to as invariant moment of tangent distribution (IMTD) methods.

An estimate of surface orientation can be obtained by minimizing the mean deviation in the estimated lengths of surface MVs:

$$E_{\sigma} = (1/n)\sum_{i=1}^{n} \text{abs}\,(\overline{|V|} - |V_i|). \tag{8}$$

Where each V_i is associated with an image region. The method which minimizes this function is called the IMTD-MV method. An estimate of surface orientation may also be obtained by minimizing the mean deviation in the estimated Is of a set of tangent distributions:

$$E_I = (1/n)\sum_{i=1}^{n} \text{abs}\,(\bar{I} - I_i). \tag{9}$$

The method that minimizes E_I is called the IMTD-I method. Results for the IMTD-MV method are shown in table 6. Results not reported here (and results for images similar to those in figures 1–8 reported in Stone (1991)) suggest that the performances of the IMTD-MV and IMTD-I methods are similar. The assumption of weak homotropy ensures that the IMTD-MV method works well on all of the textures shown. Once more note that the IMTD methods rely on the invariance of a surface quantity, rather than on the assumption of a known surface quantity.

Finally, we have not included a method which follows logically from the ordering of methods described so far. This is a method that assumes only that the principal orientation of surface tangent distributions is invariant. This method (the invariant principal orientations (IPO) method) was implemented, but was found not to be effective for the types of textures shown here.

4. DISCUSSION

Throughout this paper we have attempted to draw a strong distinction between two classes of methods. The value-seeking methods assume a known quantity (e.g. minimal degree of anisotropy) on the surface, whereas the invariance-seeking methods assume only that there is an invariant quantity (e.g. texture density) on the surface.

Methods which require that the value of a surface metric is known are likely to prove less general than methods which require only that the value of a surface metric is invariant. For instance, the assumption of isotropy can be used to obtain accurate estimates of surface orientation only if a surface texture is isotropic. In contrast, the assumption of homotropy can be used for textures on which the shape and orientation of the local tangent distribution is invariant with respect to position, irrespective of the form of the local tangent distribution. We propose that invariance-seeking methods are likely to provide more plausible models of human perceptual processes than the more conventional value-seeking methods.

Characterizing a method as invariance-seeking or value-seeking does not predict its behaviour for textures that violate the assumptions on which the method is based. Thus, some methods fail gracefully, and others do not fail so gracefully, when their assumptions are violated. In this paper we have attempted to demonstrate methods which work well on textures which are consistent with the methods' assumptions, and how such methods fail when presented with textures which violate their assumptions.

An intriguing side-effect of this work has been to generate images which appear to depict non-planar surfaces, even though certain statistics of the surface texture are known to be consistent with planarity. In particular, in figure 8, the changing orientation of the local surface tangent distribution gives a strong impression of a curved surface. In contrast, the anisotropic, but fixed, in-plane orientation of the local surface tangent distribution of figure 3 gives an impression of a planar surface. This implies that, if the in-plane orientation of a planar surface tangent distribution is changing, then human observers tend to see a curved surface. This, in turn, suggests that homotropy is assumed by the visual system. These observations, whilst plausible, are only speculative, because we are not aware of published work specifically on the effects of anisotropy on the perception of perspective images of textured surfaces.

5. CONCLUSION

We have demonstrated that it is possible to obtain accurate estimates of surface orientation for textures that are neither homogeneous nor isotropic. Moreover, we have demonstrated a set of methods that are more general than those which depend only on assumptions of homogeneity (Aloimonos 1988; Kanatani & Chou 1989), isotropy (Witkin 1981), or weak isotropy (Gårding 1991).

Future work will apply these techniques to images of real surfaces using the adaptive filtering technique described in (Stone 1990, 1991; Stone & Isard 1993). Additionally we intend to use local analysis to obtain initial estimates of surface orientations for a curved surface, and then to use an invariance-seeking strategy to iteratively update these estimates, with continuity constraints on the estimates of adjacent image regions.

Thanks to Stephen Isard (Edinburgh) and Raymond Lister (Brisbane) for comments on drafts of this paper. Thanks also

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

to Peter North and Alistair Bray (Sussex). This work was undertaken as part of a D.Phil. at the Department of Experimental Psychology, University of Sussex, and was supported by the Department of Cognitive and Computing Sciences, University of Sussex.

REFERENCES

Aloimonos, J. 1988 Shape from texture. *Biol. Cyber.* **58**, 345–360.

Blake, A. & Marinos, C. 1990 Shape from texture: estimation, isotropy and moments. *Artif. Intell.* **45**, 323–380.

Browne, L.G. & Shvayster, H. 1990 Surface orientation from projective foreshortening of isotropic texture autocorrelation. *IEEE PAMI* **12** (6), 584–588.

Gårding, J. 1991 Shape from surface markings. Ph.D. thesis, Department of Numerical Analysis and Computing Science, University of Stockholm.

Gibson, J.J. 1979 *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Kanatani, K. & Chou, T. 1989 Shape from texture: general principle. *Artif. Intell.* **38** (1), 1–49.

Kube, P. 1986 Using frequency and orientation tuned channels to determine surface slant. In *Eighth Annual Conference of the Cognitive Society, Amherst, Massachusetts*, pp. 235–244.

Nelder, J.A. & Mead, R. 1965 Described in Press, W.H., Flannery, B.P., Teukolsky, S.A. & Vetterling W.T. Numerical recipes in C, pp. 305–309 (1988). Cambridge University Press.

Stone, J.V. 1990 Shape from texture: textural invariance and the problem of scale in perspective images of textured surfaces. *British Machine Vision Conference, Oxford, U.K.* pp. 181–187.

Stone, J.V. 1991 Shape from texture: a computational analysis. D.Phil. thesis, Department of Experimental Psychology, University of Sussex.

Stone, J.V. 1992 The adaptive bisector method: separating slant and tilt in estimating shape from texture. *British Machine Vision Conference. Leeds, U.K.*, pp. 177–186.

Stone, J.V. & Isard, S. 1993 Adaptive Multi-Scale Filtering: A General Method for Obtaining Shape From Texture. *IEEE PAMI.* (Submitted.)

Turner, M.R., Gerstein, G.L. & Bajcsy, R. 1991 Underestimation of visual texture slant by human observers: a model. *Biol. Cyber.* **65**, 215–226.

Witkin, A.P. 1981 Recovering surface shape and orientation from texture. *Artif. Intell.* **17**, 17–45.

APPENDIX 1. THE COORDINATE SYSTEM

The image (x,y) and the world (X,Y,Z) coordinate system share a common origin, with the x and y axes being parallel to the X and Y axes, respectively. The Z-axis points away from the viewer. This defines a left-handed coordinate system. The central axis of the imaging system in the Z-axis. The focal point is at the origin, and the image plane is $Z=1$, which defines a focal length f of unit length.

The orientation of a plane is defined in terms of the gradient space parameters:

$$P=-(\partial Z/\partial X), \tag{10}$$

$$Q=-(\partial Z/\partial Y). \tag{11}$$

The surface normal of a plane has direction $(P,Q,1)$. The pair (P,Q) defines a set of parallel planes:

$$-Z=PX+QY-K, \tag{12}$$

where K is a constant, the value of which defines a single plane.

Using perspective projection with $f=1$, a point (X,Y,Z) on a plane projects to a point (x,y) in the image:

$$x=X/Z, \tag{13}$$

$$y=Y/Z. \tag{14}$$

APPENDIX 2. ESTIMATING A SURFACE TANGENT DISTRIBUTION

For a given estimate (P,Q) of surface orientation the surface tangent distribution of a given surface region is required. Each surface region is the back-projection of a corresponding square image region. The surface tangent distribution is computed by back-projecting each line in the image region to a line in the corresponding surface region. The estimated length and orientation of the surface line are also computed. Having estimated the length and orientation of every line in a given surface region it is then a simple matter to construct an edge orientation histogram of these surface lines. This edge orientation histogram is the discrete version of the tangent distribution.

Note: in the simulations for which results are shown here the orientations of surface lines were quantized into 20° bins. The image region sizes were 120×120 pixels.

Define a line mapping function *lmf* which maps an image length s_i with tangent β to a corresponding surface length S_i:

$$S_i=lmf(x_0,y_0,P,Q,\beta,s_i). \tag{15}$$

It can be shown (Stone 1990, 1991) that:

$$S_i=\int_0^{s_i}(K/W)(Qx_0\sin(\beta)-\cos(\beta)(1+Qy_0))^2+(\sin(\beta)(1+Px_0)-Py_0\cos(\beta))^2+(P\cos(\beta)+Q\sin(\beta))^2)^{1/2}\mathrm{d}s. \tag{16}$$

Where $W=(1+Px_0+Qy_0+s(P\cos(\beta)+Q\sin(\beta))^2$.

Where (x_0,y_0) is the image position of one end of the 'line' being measured, and $P=-(\partial Z/\partial)$, $Q=-(\partial Z/\partial Y)$.

The corresponding orientation α of the surface line may be obtained from a tangent mapping function *tmf* defined as:

$$\alpha=tmf(x_0,y_0,P,Q,\beta). \tag{17}$$

This is evaluated as follows. The orthogonal projection of a surface tangent at (X,Y,Z) with orientation α onto a front-parallel plane yields $T=\mathrm{d}Y/\mathrm{d}X$. The projection of the tangent T into the image then yields the image tangent $\tan(\beta)$.

Combining these transformations, it can be shown (Stone 1991) that:

$$\tan(\alpha)=\tan(\arctan(T)-\tau)\cos(\sigma). \tag{18}$$

Where τ is the tilt (arctan (P/Q)), and σ is the slant (arctan $((P^2 + Q^2)^{1/2})$) of the surface, and

$$T = \mathrm{d}Y/\mathrm{d}X$$

$$= \frac{(1 + Px)\tan(\beta) - Py)}{((1 + Qy) - Qx\tan(\beta))}.$$

APPENDIX 3. IMPLEMENTATION DETAILS

For all methods described in this paper we use a general search technique, the 'amoeba simplex', desribed in (Nelder & Mead 1965) to find that set of orientation parameters (P,Q) which minimizes a given cost-function (E_x).

Each image has unit area and the focal length is unity. The regions used to obtain image data are defined by a grid. Each element of this grid has area $= 0.166 \times 0.166$. To increase the number of regions a second set of regions is defined by shifting the first grid by (0.166/2,0.166/2). For this second set, regions which do not fall entirely within the image are discarded. The set of regions referred to in this paper consists of both the first and (smaller) second set defined here.

All lines in textures consisting of line segments have the same length on the surface. Textures consisting of circles and ellipses are represented as individual line segments. Each circle or ellipse was synthesized as 36 line segments. The surface orientation of all surfaces shown here is $(P,Q) = (-0.287,0.789)$.

Received 11 March 1992; accepted 13 August 1992