Bidirectional Texture Contrast Function

Sylvia C. Pont and Jan J. Koenderink

Helmholtz Institute, Universiteit Utrecht, Princetonplein 5, 3584CC Utrecht,
The Netherlands,
s.c.pont@phys.uu.nl,
http://www.phys.uu.nl/~wwwpm/HumPerc/

Abstract. We consider image texture due to the illumination of 3D
surface corrugations on globally smooth curved surfaces. The same
surface corrugations give rise to different image textures depending on
illumination and viewing geometries. We study surfaces that are on the
average approximately Lambertian. The surface roughness gives rise
to luminance modulations of the global shading pattern. The extreme
values of the luminance depend on simple geometrical factors such as
whether surface micro facets exist that squarely face the light source or
are in shadow. We find that a simple microfacet-based model suffices
to describe textures in natural scenes robustly in a semi-quantitative
manner. Robust statistical measures of the texture yield the parameters
for simple models that allow prediction of the BRDF. Thus texture
analysis allows the input parameters for inverse rendering and material
recognition to be estimated.

Keywords. Texture, shading, and colour. Image features. Surface ge-
ometry. Inverse rendering.

1 Introduction

The term texture is used in a number of categorically different senses. "Texture"
as it appears in "shape from texture" [8,10] and "texture mapping" techniques
refers to image structure due to flat patterns on smooth surfaces, such as wall-
papered surfaces, tattooed bodies, or painted decoration. We will refer to this
type of texture as wallpaper texture. Techniques like shape from texture and
texture mapping depend on the fact that the texture itself has no 3D structure,
but reflects the structure of the decorated surface. Thus the only property of
interest is the fact that texture provides landmarks. On a closer analysis one
finds that perhaps the majority of textures in images of natural scenes are of a
different variety. Such textures are due to the fact that most surfaces are rough
on scales that can be resolved by the camera. These "3D textures" generate
image structure due to a variety of optical effects. Such textures exist even if the
surface has a uniform albedo, such as a white plastered wall.

Wallpaper texture provides landmarks and thus allows one to draw infer-
ences concerning the shape of the global surface. In contradistinction, 3D tex-
ture does not provide landmarks in a strict sense, because the image structure

depends strongly on both the illumination and viewing geometry [20]. Although 3D texture is perhaps less useful to generate landmarks, it is a rich source of information concerning the structure of the local light field and the nature of the surface corrugations. In some sense there is a continuous spectrum that runs from object shape, over 3D texture, to the domain of smooth surfaces. In the case of wallpaper texture one ignores the 3D micro structure of the surface.

In this paper we assume that surfaces have uniform albedo throughout, that is to say, we assume that all texture is 3D texture. In the case of 3D texture one acknowledges that the surface is 3D structured, but chooses to ignore the shapes of individual protrusions or indentations. As a summary description one may use such measures as the probability density of heights, the autocorrelation function of heights and the probability density of orientations of local micro facets. For example "bump mapping" techniques [7] in computer graphics regard only the distribution of orientations and ignore differences in height. Indeed much of the image structure generated by 3D texture is due to the fact that surface microfacets differ in orientation and thus receive different illuminances according to Lambert's law [18]. However, the height distribution is also important because it causes such important effects as vignetting, shadowing and occlusion on the micro scale.

Although 3D texture has not received much attention in computer vision, its importance has long been recognized by visual artists. An early scientific application of 3D texture is Galileo's work on the moon. In 1610 Galileo published his observations of the Moon (Collection Biblioteca Nazionale Centrale, Florence) showing that the Moon's surface wasn't smooth but covered with mountains, valleys and craters. He calculated the height of the mountains on the basis of measurements of the shadows and his drawings clearly show that the contrast is highest at the shadow terminator. Grazing illumination methods are used in art historical research to bring out minute relief in paint layers that may perhaps be due to pentimenti by the artist or later additions by third parties. ([2,4, 12,15,24]). Photographers also use illumination at very shallow angles to reveal contrast in rough subjects, and note different treatments for matte and specular scattering on the micro scale [14].

3D texture was formalized by means of the Bidirectional Texture Function or shortly BTF [6], as the dependency of texture (the local variation of pixel intensity for projections of the characteristic surface corrugations onto several image pixels) on viewing and illumination directions. As opposed to BTFs for 2D albedo variations on an object surface, BTFs for 3D geometrical variations of the surface mesostructure cannot be derived from single values of this BTF through simple texture or bump mapping, because there are also effects of local vignetting and shadowing [16].

In this paper we study the variation of contrast as a function of viewing and illumination directions. This function we call, by analogy with the BTF, the Bidirectional Texture Contrast Function (BTCF). The texture contrast gradients identified here are robust and extremely common (tree bark, foliage, architectural

detail, cloth, skin, common building materials, . . .) regularities that still await to be exploited.

2   Theoretical Considerations

### 2.1   Objects in Light Fields

The texture mapping from computer graphics applies only to "wallpaper texture". That is to say, the texture consists of spatial albedo variations on a globally smooth surface. In this case the local illuminance in the image of the texture is independent of the way the surface is illuminated. The luminance is simply the product of the local illuminance and albedo. For the type of texture considered here this doesn't apply at all. We consider the case that the local albedo is uniform and all the texture in the image is due to variations in local illuminance. Variations in local illuminance arise from a number of factors: the attitude effect (Lambert's law), or shading proper, vignetting, that is occultation of the source by the object itself, mutual interreflections between surface facets, . . .

For a collimated beam such as direct sunlight vignetting is all or none and one usually uses the terms body shadow and cast shadow instead of vignetting. When a surface facet is illuminated, its illuminance is due to the attitude effect, when not illuminated it registers as black in the image (we ignore interreflections in this paper).

For a diffuse beam such as illumination by the overcast sky vignetting is gradual and can hardly be distinguished from the attitude effect [9]. Thus conventional shape from shading methods [13] do not apply. Diffuse beams can still be strongly directional. A generic example is the beam due to an infinite luminous pane at one side of the scene. Such a "uniform hemispherical beam" illuminates every point on a convex object such as a sphere.

In the case of illumination with a collimated beam a sphere will have one hemisphere in total darkness (the "body shadow"), whereas the illuminated hemisphere has an illuminance distribution $\cos\theta$ where θ is the angle from the pole facing the source. This dependence is due to the surface attitude effect (Lambert's Law). In finding the contrast due to 3D surface corrugations we may simply assume a range of slants of micro facets at any point of the sphere. The range will be centered on the fiducial slant and we may suppose it to extent to an amount $\Delta\theta$ on either side. In order to find the contrast we need to look for the micro facet slants that minimize and maximize the irradiance. Near the pole facing the source it will be possible to find micro facets that face the source fully and thus receive the maximum illuminance ($\cos 0 = 1$). Farther from the pole (when $\theta > \Delta\theta$) this will not be possible and the maximum illuminance will be $\cos(\theta - \Delta\theta)$. Near the terminator $\theta = \pi/2$ it will always be possible to find facets that face away from the source and thus receive zero illuminance ("shadow"). Nearer the illuminated pole (when $\theta < \pi/2 - \Delta\theta$) any facet will receive at least some radiation, the minimum irradiance being $\cos(\theta + \Delta\theta)$. When we define the contrast as the difference of the maximum and minimum illuminance divided by

twice the fiducial illuminance (see methods), we find that the contrast monotonically increases from the illuminated pole to the terminator, actually explodes near the terminator. In practice there will always be some ambient light and the contrast maximum will be found near $\theta = \pi/2$ and be arbitrarily large when the ambient level is low. This is illustrated in figures 1, 2 and 3.
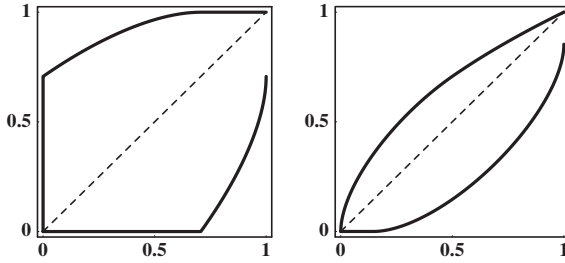

Fig. 1. Theoretical plots of the "minimum" (5% percentile) and "maximum" (95% percentile) illuminance against the fiducial (median) illuminance, left for a collimated beam, right for a uniform hemispherical diffuse beam. A range of micro facet slopes ($\Delta\theta$) up to 45° from the nominal normal was assumed. In this calculation only the attitude effect (Lambert's Law) is taken into account.

Consider the case of very diffuse beams. A typical example is the uniform hemispherical diffuse beam. In computer graphics this is often called "point source and ambient illumination" for no better reason than that the illuminance of a sphere is proportional with $(1 + \cos\theta)/2$, where θ denotes the angle from the pole that faces the source. The expression is interpreted as the sum of two terms, one $(1/2 \cos\theta)$ due to a "point source" (at infinity), and the other $(1/2)$ due to an "ambient" (fully isotropic) beam. This appears to work because the illumination on the hemisphere facing a collimated beam is proportional with $\cos\theta$. However, in the latter case the other hemisphere is in darkness (in graphics the "negative illuminance" is clipped). In the former case the ambient term conveniently cancels the "negative contribution" of the "point source".

In reality there is neither a point source, nor an ambient beam. "Negative" illumination is a nonentity. Notice that $(1+\cos\theta)/2 = \cos(\theta/2) \times \cos(\theta/2)$, where we write the square as a product of two terms. One term is due to vignetting, that is the fact that a surface facet is typically only illuminated by part of the source, the other parts being occluded by the object itself. The other term is due to the fact that the surface facet will be at an oblique attitude with respect to the effective source direction. The latter is caught by "bump mapping", the former has not led to approximate graphics methods, but is apparently of equal importance. We elaborate this formal structure because surface corrugations can affect either term of the (physically significant) expression whereas the "point source with ambient illumination" formulation leads to nonsense because it is physically meaningless.
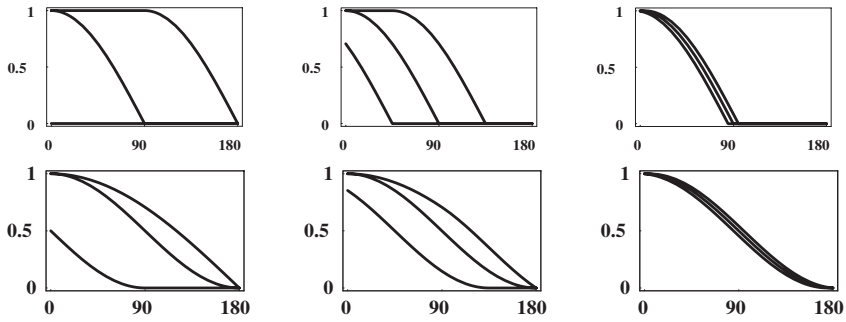
Fig. 2. Theoretical "minimum" (5% percentile), "maximum" (95% percentile) and median illuminances as a function of the nominal direction of illumination with a collimated beam (upper row) or with a uniform hemispherical diffuse beam (lower graphs). Micro facet slope variations ($\Delta\theta$) of 60°, 30° and 15° (left to right) assumed. Typically (*e.g.*, see the center plot) the minimum reaches the zero level before the median does: From that point on micro shadows exist. In the upper row we can see that in the case of illumination with a collimated beam the maximum stays constant over a range where the median is steadily decreasing: In this regime micro facets exist that face the source head on. Notice that the maximum extends far into the shadow region. That is the case because the height distribution was disregarded in this (simplest) model. In realistic cases the maximum can't extend very far into the shadow region because the tops will soon fail to stick out of the cast shadow region. The lower graphs show that with diffuse illumination the maximum drops much slower than the median, both are steadily decreasing as one moves away from the illuminated pole.
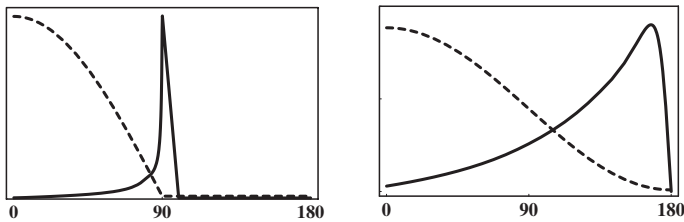
Fig. 3. The fiducial illuminance (dashed curve) and the 3D texture contrast (drawn curve) as a function of the nominal direction of illumination, as predicted from the simplest possible model (only Lambert's Law taken into account, micro facet slope variations up to 10° from the nominal surface normal assumed). On the left for illumination with a collimated beam, on the right for a uniform hemispherical diffuse beam with small ambient component. For collimated illumination 3D texture contrast peaks at the terminator, for diffuse illumination at the dark pole.

Consider how a surface facet that slants with respect to the fiducial slant (of the surface of the true sphere) may affect its illuminance: It may vignet the source itself, leading to decreased effective power of the source, and it may change the obliquity term which may lead to either increased or decreased illuminance.

In finding the contrast due to 3D surface corrugations we may simply assume a range of slants of micro facets at any point of the sphere. The range will be centered on the fiducial slant and we may suppose it to extend to an amount $\Delta\theta$ on either side. In order to find the contrast we need to look for the micro facet slants that minimize and maximize the irradiance. Near the pole that faces the source it will be possible to find facets that squarely face the effective source, thus leading to an illuminance of $\cos(\theta/2)$ (not $1/2$, because there will always be vignetting by the global object). When such a facet cannot be found (which is the case when $\theta > 2\Delta\theta$) the maximum illuminance is $\cos(\theta/2)\cos(\theta/2 - \Delta\theta)$. Near the pole that faces away from the source it will be possible to find facets that vignet the source completely, leading to zero illuminance ("shadow"). At other locations (when $\theta < 2\Delta\theta$) the minimum illuminance will be $\cos(\theta/2)\cos(\theta/2 + \Delta\theta)$. The contrast monotonically increases from $1/4$ at $\theta = 0$ to infinity at $\theta = \pi$. In real scenes there will always be some ambient light (real ambient light!) and the contrast maximum will thus be limited to large values. When we prepare the same format plots as in the case of collimated illumination we find the right graph in figure 1, the lower graphs in figure 2 and the right graph in figure 3.

2.2   Modelling the Optical Effects Due to Surface Corrugations

The analysis of the previous section is the simplest possible model for 3D surface corrugations. We find that a more realistic model yields essentially the same results. In a more realistic model we calculate the histogram of radiance at the eye or camera, and define the minimum, maximum and fiducial radiance as the 5% 95% and 50% percentiles. In order to find the histograms one needs a statistical model of the surface.

Much of the physics is purely due to the fact that surface micro facets differ in slant from the fiducial slant of the nominal surface, the only significant exception being the protrusions that stick out of the cast shadow volume and lead to illuminated patches in the nominal "body shadow" area. This occurs in a narrow strip at the terminator where one needs some information concerning the heights of the protrusions. Let's say the maximum protrusion height is h and the radius of the object is R. It is simple to derive that the strip is limited by an angle $\lambda \approx \sqrt{2h/R}$ from the shadow terminator. For instance, protrusions with a height of only 1% of the radius ($h/R = .01$) lead to illuminated patches in the body shadow area up to .14 radians or 8°. This effect depends upon the distribution of heights.

Most of the image structure can be understood from a description that only recognizes the statistics of the orientations of surface micro facets though, disregarding the height distribution completely. This is roughly the kind of texture that can be generated through classical "bump mapping". Thus we are led to consider micro facet theory as has been developed in the theory of the BRDF[21] of rough surfaces [1]. The calculations are elementary, taking the probability density of facet orientations and the effect of foreshortening of their visible areas into account. The one non–trivial problem is that of shadowing and occlusion:

We need to estimate the probability that a facet be both illuminated and visible. This very difficult problem has only been solved exactly for some specific cases [17] and even its statistical treatment (in the case of Gaussian surfaces for instance) is problematic [11,23]. We have used the simple and convenient approximative treatment advocated by Ashikhmin[1]. We find that the results of contrast estimations are rather insensitive to the details of the model (even the coarsest possible treatment disregarding occlusion and vignetting on the microscale completely is not very different from more precise modelling). To include elaborate models for the shadowing function is apparently overkill since the effect is hardly sensitive to it. This is an important observation, because one will typically lack precise surface models in actual applications. When we use a simple tangent–normal probability density for the surface micro facet distribution we obtain results that very closely mimic our empirical results as can be seen in the sequel.

3   Empirical Examples

3.1   Methods

We present examples of 3D texture measurements both from the laboratory under strictly controlled conditions and from natural scenes obtained via informal methods.

In the laboratory we produced collimated illumination via a Xenon arc source and a small aperture. Objects were mounted on a robot arm and both the illuminator and the camera could be rotated about the object very precisely. The camera was a calibrated Leaf scanning camera (the camera has a dynamic range of 12 bits). Measurements were done in a dark room with black painted walls. Diffuse illumination was produced via a large uniformly backlighted diffusing screen. The object was placed very close to the screen and was much smaller than the screen. Thus this approximates uniform hemispherical illumination closely.

The images obtained from natural scenes were taken in direct sunlight under blue sky conditions using a hand-held camera. The camera has a dynamic range of 8 bits preceded by a gamma transformation. For the analysis we corrected for the gamma.

In this paper we concentrated on the measurement of texture contrast, ignoring most of the spatial structure and precise shape of the intensity histograms. This is about the simplest analysis that makes sense. A number of distinct notions of "contrast" are in common use, and each has its uses. In case one has symmetric modulations on a steady level perhaps Michelson's definition (commonly used in optics, [3]) may be the most reasonable: Michelson Contrast is the difference of the minimum and maximum levels divided by their sum. For isolated peaks on a steady level it is often more reasonable to use the peak amplitude (as measured from the steady level) divided by the value of the steady level. In our applications we deal with the 5% ("minimum"), 95% ("maximum") and 50% (median illuminance) levels, which are the robust measures required for

natural images, whereas there is also a notion of the "fiducial illuminance" in the case of simple objects like (approximate) cylinders and spheres. Although the median illuminance may often be used as an estimate of the fiducial illuminance, it may well be biased, especially near the terminator. In cases of interest here the Michelson contrast is inconvenient since it "saturates" at unity and tends to hide structure that is of interest. We use the difference of the 95% and 5% percentiles divided by twice the median value. This measure sometimes "explodes" near (or beyond) the terminator, but we find that it typically yields useful insight into the structure of the illuminance pattern. In fact, even an "explosion" itself is relevant. When necessary (for instance when the median is zero) we simply specify the three percentiles, this disambiguates contrast measures sufficiently in practice. The next "step up" would be to specify the full histogram, then the degree of approximation can be precisely specified. A contrast measure is essentially a measure that captures the relative width of the histogram.

3.2   Results

The Curet database [5] contains BTFs of about 60 natural surfaces. Each BTF is characterized through 205 images for different illumination and viewing angles. This gave a fine opportunity to test the generality of this thesis. For this purpose we looked at the following samples: 1 (felt), 2 (polyester), 3 (terrycloth), 4 (rough plastic), 5 (leather), 6 (sandpaper), 8 (pebbles), 11 (plaster), 12 (rough paper), 15 (aluminum foil), 28 (crumpled paper) and 45 (concrete). We computed minimum (5% percentile), maximum (95% percentile) and median pixelvalues for grayscaled pieces of textures that were cut out of all 205 photo's per sample. For samples 6 (sandpaper), 8 (pebbles), 11 (plaster), and 15 (aluminum foil) we show the results in figures 4 and 5. Figure 4 can be compared with figure 3 (left graph). The data in the upper row of figure 5 are directly comparable to the results in figure 1 (left graph). One should keep in mind that we interpolated (with 5° resolution) from the discrete original data and that there are no images in the Curet data for the shadow region. Still, it is clear that contrast is highest near the terminator and the maximum contrast values and the scatterplots for plaster, pebbles and sandpaper also clearly show the relation with the roughness of the samples. It is evident from this figure that this is a result of three dimensional surface undulations of the samples; the textures on the left are arbitrarily chosen images of the samples and cannot be transformed into each other via "texture mapping". The scatterplots in figure 5 show in the upper row the 5% and 95% percentiles against the 50% percentiles for all 205 photo's per sample in the database. The lower row shows the contrast against the median luminance. These plots show patterns in line with the very simple models explained earlier (figures 1 and 2). Apparently the simple model accounts already for most of the structure in the data.

Crumpled aluminum foil shows quite different properties than the other samples. The contrast immediately rises to very large values which is caused by the specular reflection of the facets. This material indeed violates the model assumption of Lambertian scattering surface micro facets. Oversimplified, the aluminum
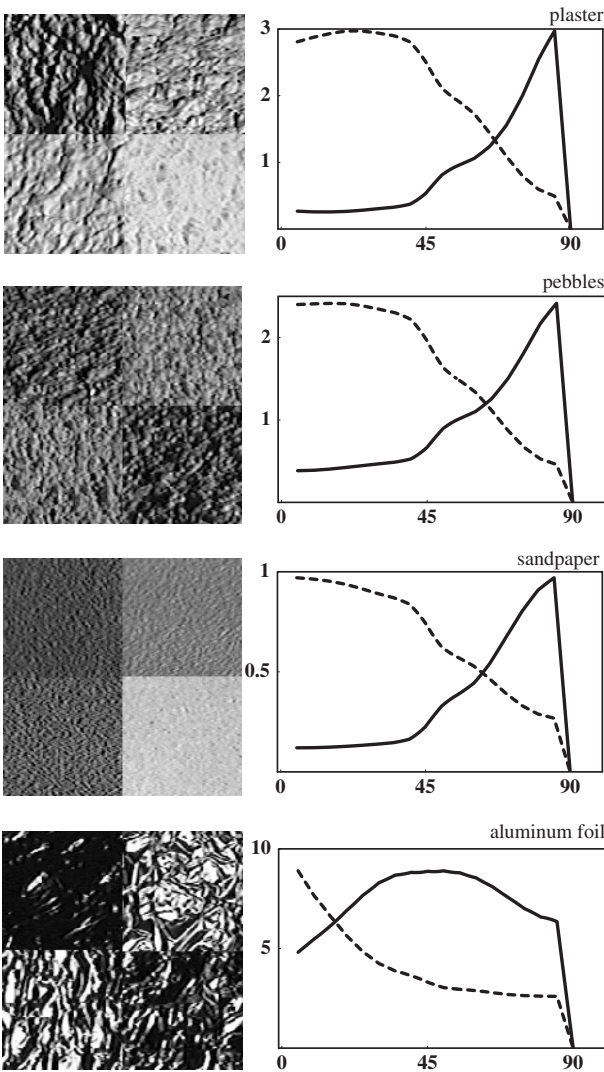
Fig. 4. Samples 6 (sandpaper), 8 (pebbles), 11 (plaster), and 15 (aluminum foil) from the CURET database. On the left a few sample images. Notice that the textures have distinct spatial structures and cannot be transformed into each other via "texture mapping". On the right plots of the contrast (drawn curve) and median luminance (dashed curve) as a function of the lighting angle in the plane of incidence (interpolated from the discrete CURET data with 5° resolution). The scale applies to texture contrast, the median luminance has been scaled to the same maximum value as the contrast. Notice that the curves do not extend into the shadow region because no measurements exist for this region. The graphs can be compared with figure 3 (left graph).
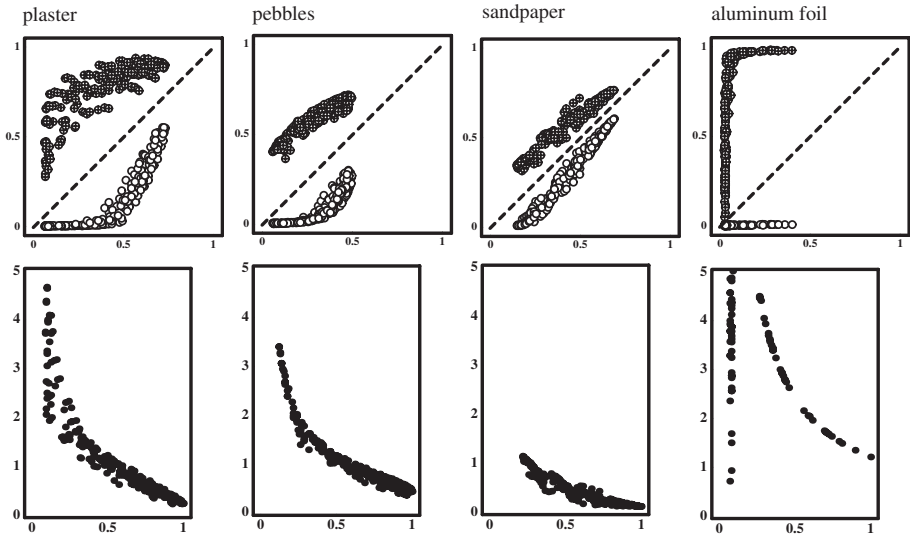
Fig. 5. Samples 6 (sandpaper), 8 (pebbles), 11 (plaster), and 15 (aluminum foil) from the CURET database. The upper row shows scatterplots of the minimum (5% percentiles) and maximum (95% percentiles) against the median luminance. The lower graphs show the contrast against median luminance. This figure can be compared with figure 1

foil images can be seen as binary textures; facets are either highlighted because they reflect the light towards the camera or completely dark because they scatter the light in a different direction. The exact patterns of course depend on the distribution of slopes. Various samples showed results that are similar to the aluminum foil texture though less outspoken and mixed with the typical pattern seen for matte surfaces. The "aluminum foil effect" shows up most clearly in the outliers due to specular reflection (rough plastic, crumpled paper).

In figure 6 we show the results for a candle in the cases of collimated illumination (lower part of the figure) and diffuse illumination (upper part). From the photographs it is clear that the three dimensional contrast is highest at the shadow terminator in the case of collimated illumination and at the pole that faces away from the diffuse source, closely resembling the predictions of the model. The upper and lower graphs can be compared to the right and left graphs in figure 3, respectively. Other spherical and cylindrical objects, for instance cylinders which were covered with linen and knitting, showed qualitatively similar results. Numerical results differed in line with the variations in the roughness of the samples.

In figures 7 and 8 the results are shown for two natural scenes. Again we find that contrast explodes near the shadow terminator. It is clear that this analysis is very robust for conditions of illumination, deviations from the exact cylindrical
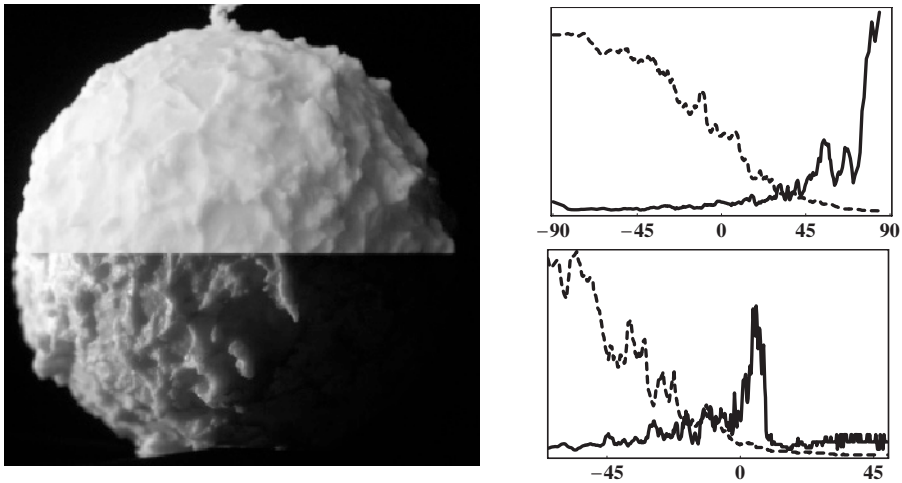
Fig. 6. On the left two images of an "artistic" candle (roughly spherical with very rough finish, uniform white) illuminated with a diffuse beam (upper image) or a collimated beam of a Xenon arc source (lower image). Notice the very strong and coarse 3D texture contrast, in the upper image increasing up to the right edge and for the case of the collimated beam extending into the nominal "body shadow" area. On the top–right graphs of the median radiance (dashed curve) and contrast (drawn curve) on arbitrary linear scales. On the abscissa the obliquity with respect to the viewing direction. The minimum irradiance for the upper image is at 90°. The shadow terminator in the lower graph is located at the origin. The "noise" on the data is due to the rough 3D texture of the object. The upper and lower graphs can be compared to the right and left graphs in figure 3, respectively.

geometry of the sample and the exact lay-out of the measurement. The graphs can be compared with the model predictions in figures 2 and 3.

3.3   Robust Semi-quantitative Estimation of Surface Roughness Parameters

Microfacet models require knowledge of the distribution of the orientations of microfacets and possibly the distribution of heights. From observation of the texture we are able to estimate some overall measures that can be used to do rough microfacet modelling. An estimation of the effective heights of surface protrusions can be found from the distance from the terminator at which parts of the surface that are in the cast shadow region still manage to catch a part of the incident beam.

An estimate of the range of orientations of surface microfacets orientations can be obtained in two independent ways. The maximum luminance of the 5% percentile divided by the maximum illuminance of the 95% percentile equals the cosine of the largest local deviation. Alternatively we find the location where the 95% percentile starts to decrease and compare it with the location of the
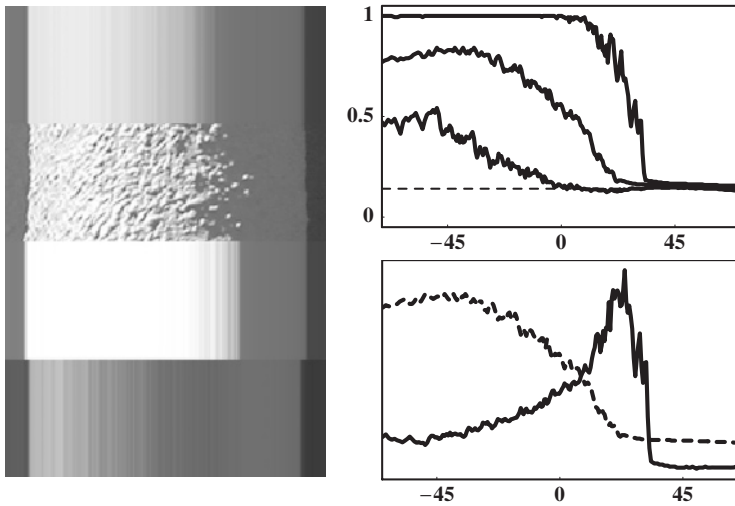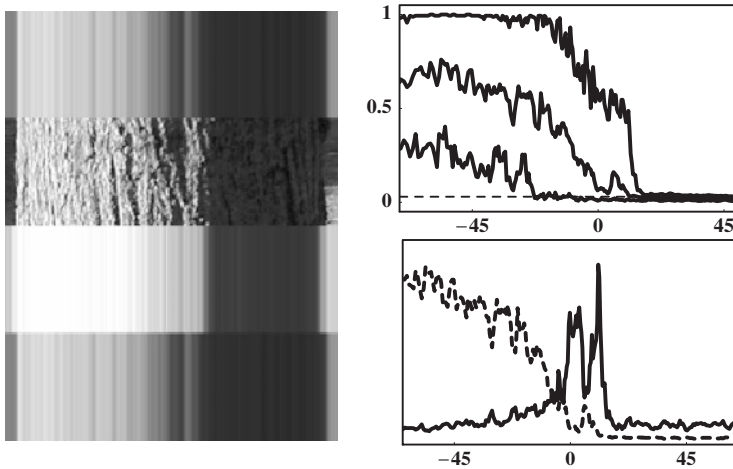
Fig. 7. On the left a collage of (from top to bottom) the median radiance, actual image, maximum (95% percentile) and minimum (5% percentile) radiances. Notice the clear gradient in texture contrast visible in the photograph and the light islands well within the "body shadow" area. On the upper right the minimum, maximum and median intensities are plotted as a function of the illumination angle. The terminator is at the origin. Notice the appreciable ambient component (dashed level). The lower graph shows the median illuminance (dashed curve) and the 3D texture contrast (drawn curve), which is highest at the terminator. This was a pillar coated with coarse white plaster on a San Francisco house illuminated by direct sunlight. The graphs can be compared with figures 2 and 3.

terminator. This again lets us estimate the maximum local deviation from the nominal surface normal.

The width of a typical surface modulation is easily obtained from the width of the autocorrelation function of the pixel intensity along a scan line parallel to the cylinder axis.

In the candle example (figure 6) the texture is very coarse, the width of the autocorrelation function being 6% of the radius. Illuminated parts can be found up to roughly 10° from the terminator, suggesting a height of 1.5% of the radius. From the luminances at the frontally illuminated side we find that slopes up to 43° occur, whereas the drop-off point of the 95% percentile suggest 37°.

In the pillar example (figure 7) the luminances at frontal illumination reveal that slopes up to 62° occur. From the onset of drop-off in the 95% percentile we find 70°. For the pillar illuminated spots are found 10° from the terminator suggesting that protrusions with heights up to 1.5% of the radius occur. The autocorrelation function reveals that the width of the protrusions is 2.4% of the radius.

Fig. 8. On the left a collage of (from top to bottom) the median radiance, actual image, maximum (95% percentile) and minimum (5% percentile) radiances. On the right the minimum, maximum and median intensities are plotted as a function of the illumination angle. The terminator is at the origin. Notice the appreciable ambient component (dashed level). The lower graph shows the median illuminance (dashed curve) and the 3D texture contrast (drawn curve). This was a tree in a Cambridge (UK) park illuminated by direct sunlight. The "noise" is due to the very coarse structure of the bark. The graphs can be compared with figures 2 and 3.

In the tree example (figure 8) the width of the autocorrelation function is 3.0% of the radius. Illuminated patches are found up to 15° from the terminator, from which we find that protrusions with heights up to 3.4% of the radius exist. This agrees rather well with the range of surface orientations. From the maximum luminance of the 5% percentile we find slopes up to 73°, whereas the onset of the drop-off of the 95° percentile suggests the existence of slopes up to 70°.

It is apparent that the various estimates are mutually highly consistent. Thus these comparatively crude, but very robust measures suffice to obtain quite reasonable estimates of the characteristic parameters of the 3D texture.

4   Conclusions

3D texture is characterized by the probability density and the autocorrelation function of heights and the probability density of the slopes of surface micro facets. The image structure can be characterized by the spatial autocorrelation function and the intensity histogram. The image structure reflects the structure of the 3D texture although the transformation is a complicated one involving both viewing and illumination parameters. On curved surfaces one has a sampling of different viewing and illumination parameters and thus sufficient information to draw strong inferences regarding the 3D texture.

In this paper we have concentrated on the most simple analysis. Such an analysis already allows us to estimate the range of surface orientations. For most of the images a simple description taking only the distribution of normals of micro facets into account amply suffices to predict the main features of 3D texture. This means that observation of the 3D texture contrast will typically only yield estimates of the spread of surface normals about the average. In addition, the observation of the spatial structure of the texture allows inferences concerning the width of the autocorrelation function of heights. The distribution of the heights themselves can only indirectly be inferred from the width of the autocorrelation function and the spread of normals. The simplest model for the texture on the dark side of the terminator involves the distribution of heights directly though and allows an estimate of the peak heights.

Shape from shading depends upon knowledge of the (usually unknown) BRDF of the surface. The 3D texture (and perhaps specularity) observations allow one to guesstimate the BRDF [1,11,16,22]. These models depend on only a few parameters, the spread of slopes of microfacets being the most important. Thus more precise shape from shading inferences are possible when the 3D texture information is taken into account. Another field for which this is of interest is that of inverse rendering. We find that our estimates are qualitatively consistent with a simple microfacet-based model, for collimated as well as diffuse illumination of cylindrical and spherical shapes with 3D textures of a wide variety under laboratory conditions, for cylindrical shapes which were photographed ad lib in the field and for discrete sets of textures from the Curet database. The Bidirectional Texture Contrast Function might therefore be a good start in a bootstrap procedure for inverse rendering and BRDF guesstimation. After this first step on the basis of the general character of surface reflectance, the resulting functions can be finetuned on the basis of, for instance, the exact shape of the histograms, autocorrelation functions of the textures, etc.

Specularities yield similar information to that provided by 3D texture. The width of the specular patch is a direct measure of the angular spread of normals (global curvature and spread of the illumination assumed known) and the patchiness is a direct measure for the width of the height autocorrelation function [19]. Thus the structure and width of the specularity and the nature of the 3D texture are closely related and any inference concerning surface micro structure should regard both. This is a possible fruitful area of further research.

In conclusion, we find that there is much more to texture than what applies to the wallpaper type. Most texture is due to 3D surface structure and can be considered a rich source of information on both the local light field and the geometry of surfaces on the microscale.

References

1. Ashikhmin, M., Premoze, S., Shirley, P.: A microfacet-based BRDF generator. Proceedings ACM SIGGRAPH, New Orleans, 2000, 65–74
2. Bomford, D., Kirby, J., Leighton, J., Roy, A.: *Art in the Making Impressionism.* National Gallery Publications Limited, Great Britain, 1990

3. Born, M., Wolf, E.: *Principles of Optics.* Cambridge University Press, Cambridge, 1998
4. Brady, M.: *Ancient Documents.*
   `http://www.robots.ox.ac.uk/~jmb/ancient.html`
5. Curet: *Columbia–Utrecht Reflectance and Texture Database.*
   `http://www.cs.columbia.edu/CAVE/curet`
6. Dana, K.J., Ginneken B. van: Reflectance and texture of real-world surfaces. Proceedings IEEE Computer Science Conference on Computer Vision and pattern Recognition, 1977
7. Foley, J. D., Dam, A. van, Feiner, S. K. and Hughes, J. F.: *Computer Graphics, Principles and Practice.* Addison–Wesley Publishing Company, Reading, Massachusetts, 1990
8. Gårding, J.: Direct estimation of shape from texture. IEEETrans. Pattern Anal. and machine Intell. **15**, 1993, 1202–1208
9. Gershun, A.: The Light Field. Transl. by P. Moon and G. Timoshenko. J.Math.Phys. **18(51)**, (1939)
10. Gibson, J.: *The perception of the visual world.* Houghton Mifflin Company, Boston, 1950
11. Ginneken, B.van, Stavridi M., Koenderink J.J.: Diffuse and specular reflection from rough surfaces. Applied Optics **37(1)**, (1998), 130–139
12. Hogarth, B.: *Dynamic Light and Shape.* Watson–Guptil, New York, 1981
13. Horn, B. K. P., Brooks, M. J.: Shape from Shading. The M.I.T. Press, Cambridge Massachusetts, 1989
14. Hunter, F., Fuqua, P.: *Light, Science & Magic: An Introduction to Photographic Lighting.* Focal Press, Boston, 1990
15. Jacobs, T. S.: *Light for the Artist.* Watson–Guptill Publications, New York, 1988
16. Koenderink, J. J., Doorn, A. J. van: Illuminance texture due to surface mesostructure. J.Opt.Soc.Am. A **13(3)**, (1996), 452–463
17. Koenderink, J. J., Doorn, A. J. van, Dana, K. J., Nayar, S.: Bidirectional reflection distribution function of thoroughly pitted surfaces. International Journal of Computer Vision **31** (2/3), (1999), 129–144
18. Lambert, J. H.: *Photometria Sive de Mensure de Gradibus Luminis, Colorum et Umbræ.* Eberhard Klett, Augsburg, 1760
19. Lu, R.: *Ecological Optics of Materials.* Ph.D. thesis Utrecht University, 2000
20. Leung, T., Malik, J.: On perpendicular texture: Why do we see more flowers in the distance? IEEE Conf. Computer Vision and Pattern Recognition, Puerto Rico, June 1997
21. Nicodemus, F. E, Richmond, J. C., Hsia, J. J.: *Geometrical Considerations and Nomenclature for Reflectance.* Natl.Bur.Stand., (U.S.), Monogr. 160, 1977
22. Nayar, S. K., Oren, M.: Visual appearance of matte surfaces. Science **267**, (1995), 1153–1156
23. Smith, B.G., Geometrical shadowing of a random rough surface. IEEE Transactions on Antennas and Propagation **AP−15(5)**, (1967), 668–671
24. Wallert, A.: *Still Lifes: Techniques and Style, the examination of paintings from the Rijksmuseum.* Waanders Publishers, Zwolle, 1999