



Texture Histograms as a Function of Irradiation and Viewing Direction

BRAM VAN GINNEKEN AND JAN J. KOENDERINK

*Helmholtz Institute, Universiteit Utrecht, Buys Ballot Laboratory,
Princeton Plein 5, 3584CC Utrecht, The Netherlands*

bram@isi.uu.nl

j.j.koenderink@phys.uu.nl

KRISTIN J. DANA

Department of Computer Science, Columbia University, New York, NY 10027

dana@cs.columbia.edu

Abstract. The textural appearance of materials encountered in our daily environment depends on two directions, the irradiation and viewing direction. We investigate the bidirectional grey level histograms of a large set of materials, obtained from a texture database. We distinguish important categories, relate the various effects to physical mechanisms, and list material attributes that influence the bidirectional histograms. We use a model for rough surfaces with locally diffuse and/or specular reflection properties, a class of materials that commonly occurs, to generate bidirectional histograms and obtain close agreement with experimental data. We discuss several applications of bidirectional texture functions and histograms. In particular, we present a new approach to texture mapping based on bidirectional histograms. For 3D texture, this technique is superior to standard 2D texture mapping at hardly any extra computational cost or memory requirements.

Keywords: texture, reflection models, BRDF, physics-based vision

1. Introduction

The visual appearance of materials depends on the directions of irradiation and radiation. Moreover, it is a function of *scale*. The scale is essentially *arbitrary*, since the distance of the observer to pertinent objects is unbounded. Any measurement, whether by the human eye or by a CCD camera, is of finite resolution and has therefore an ‘inner scale’, or *grain*. The size of the inner scale determines the area mapped into one image pixel. If the characteristic variations of the material are subpixel, the material is observed as a uniform patch, and its appearance is completely determined by its BRDF, the bidirectional reflection distribution function (Nicodemus et al., 1977). If the characteristic vari-

ations are instead projected on the region of interest, we refer to the image as texture. In other words, texture is a function of the same two directions as the BRDF, the difference being that reflectance is a scalar function, while texture is a matrix function. Therefore we will use, following (Dana et al., 1997), the term BTF, *bidirectional texture function*, analogous to the BRDF. If the radiance of a texture is averaged and divided by the irradiance, the BTF equals the BRDF.

We use BTF data from the CURET database, which is the first comprehensive database of texture images with systematic variation of irradiation and viewing direction (Dana et al., 1996, Dana et al., 1997). This database is freely available on the World Wide Web at www.cs.columbia.edu/cave/curet. The




Fig. 1. Two examples of the dependence of texture on irradiation and viewing direction. On the top row, two images of rough plastic (sample #4 in Figure 2) are shown. The differences show up clearly in the corresponding histograms. The second row shows corduroy (sample #42 in Figure 2). The only difference between the irradiation and viewing geometry on the left and on the right is a 90° rotation of the azimuth angle of the irradiation. Due to the anisotropic surface structure of corduroy, the difference between the images is considerable. Again, this is reflected in the histograms.

main aim of this paper is to identify major material attributes and physical mechanisms that give rise to the wide variety of BTFs that occur in practice. All of these effects are relevant, and should be incorporated in accurate texture models. In this work, we focus on the BTF by considering the bidirectional radiant power distributions, or grey-level histograms of the textures. Clearly, grey-level histograms are not a complete representation of a BTF, since they exclude wavelength dependencies and spatial structure of radiance variations in texture. Nevertheless, the characteristic textural variations in a BTF of a 3D texture are clearly visible and easily identifiable in the histograms. This is illustrated by two examples in Figure 1. The representation of texture by their histograms is an apt one because the important mechanisms show up as separate modes in the histograms: peaks due to shadowing, diffuse reflection and specular reflection. Moreover, bidirectional histograms can be considered a natural extension of BRDFs. The BRDF is just the mean radiant power, divided by the irradiance, histograms make up the full distribution. They can have many applications in computer vision and computer graphics, as is discussed in Section 6. For such applications, accurate models of bidirectional histograms are required. In Section 5 we

present simulated histograms for an important class of materials: Gaussian rough surfaces with locally diffuse or specular reflection properties. This turns out to be an accurate model for many naturally occurring materials.

2. Previous work

Early work on this subject was done by Richards (Richards, 1982) who derived radiant power distributions for textures of Lambertian cylinders and spheres under hemispherical irradiation.

The need for accurate BRDFs is acknowledged in computer graphics and computer vision. Although experimental data is still scarce, many models and representations of BRDFs have been proposed in the literature and are used in applications, see e.g. (Beckmann and Spizzichino, 1963, Torrance and Sparrow, 1967, Phong, 1975, Leader, 1979, Cook and Torrance, 1981, Poulin and Fournier, 1990, Nayar et al., 1991, Tagare and deFigueiredo, 1991, He et al., 1991, Wolff, 1994, Oren and Nayar, 1995, Koenderink et al., 1996, Ginneken et al., 1998). The notion that a BTF is required to deal with the textural dependency on the directions of irradiation and radiation, appears to be novel, and no systematic investigations of this subject have been made.

In computer graphics, synthesis of 3-D textures is an important research area. The literature shows many examples of renderings of complex surfaces, many of which are contained in our database. Examples include cloth, anisotropic aluminium and velvet (Westin et al., 1992), vegetation canopy (Gerstl and Borel, 1992), fur and hair (Kajiya, 1989), metallic patinas (Dorsey and Hanrahan, 1996), human skin and leaves (Hanrahan and Krueger, 1993). It would be most interesting to compare these models with experimental BTF data. One way to do this would be to use bidirectional histograms, as discussed in Section 6.

3. Material and methods

We have used bidirectional texture images from the CURET database. A detailed description of this database and the measurements process used in building this database can be found in (Dana et al., 1996, Dana et al., 1997). The 63 samples in the database are listed and shown in Figure 2. In the remainder of this article, sample numbers from Figure 2 are added when referring to a sample.




Fig. 2. The set of 61 samples. Images are taken from the CUReT database.




Fig. 3. The CUReT database contains images of samples, placed at the origin in this figure, at various orientations (64 in total). In all cases, the irradiation is collimated and from the same direction (the x -axis in this figure). There are seven viewing directions, numbered C1 to C7, separated 22.5° . Sample orientations in the plane of incidence (17 of the total of 64), separated 11.25° , are indicated with a thick dot. For anisotropic samples, two images are available for each sample orientation, the difference being that the samples are rotated around their normal with their orientation kept fixed. For the detector at position C2, the possible sample positions are numbered from 1 to 11. The sample is frontal to the lamp at position 4, is at the specular point at position 6, and is frontal to the camera at position 8. These 11 measurements are shown in subsequent figures, except for Figure 6, where we show measurements for camera positions C1 to C7 with the sample at the specular position.

Histograms were calculated for each image (205 per sample for isotropic samples, 410 for anisotropic samples), using only pixels from the sample area. Spectral information was discarded by averaging red, green and blue intensity values for every pixel. We used 64 bins. The images are photometrically calibrated, so that the histograms show the radiant power distribution of the pixels in the image. The radiant power (x -axis) is scaled to run from 0 (black pixels, pixel value 0) to 1 (white pixels, pixel value 255). The gain of the detector used to obtain the images is set per sample and therefore radiant power distributions for different samples cannot be compared directly. The scaling of the y -axis is arbitrary since we consider only the distribution of pixels over the bins relative to each other (of course the number of pixels within the sample area varies with the viewing direction because of perspective changes).

In the next sections, we show a subset of the full set of 205 or 410 histograms. In most cases we show only measurements in the plane of incidence (that is, irradi-

ation direction, sample normal, and viewing direction lie in the same plane) with the detector at 45° from the irradiation direction. The sample positions in the plane of incidence are separated by 11.25° , and this leads to 11 measurements. This is illustrated in Figure 3.

4. Generic results

We have selected 12 samples, labelled (a) to (l), that represent distinct categories of materials with characteristic histogram changes. This selection includes isotropic samples, (a) to (h), and anisotropic samples, (i) to (l). The results are shown in Figures 4 and 5. A quick inspection of these figures immediately reveals the enormous variety of results. There is one effect that is not visible in the chosen subset of histograms: the rate at which the samples become specular at increasing incidence angles. To illustrate this, we have included the histograms for the 7 specular positions for 9 samples in Figure 6. We will now discuss the results in detail.

4.1. Isotropic samples

We start the discussion with (a) quarry tile (#25) because this sample is approximately smooth at the scale given by the resolution of our images. As a result, the histograms show sharp peaks in all cases, indicating an almost uniform texture. Basically, we are only looking at the BRDF here. Note that this sample is not Lambertian, since in that case the brightness would be independent of the viewing direction and be proportional to the cosine of the incident angle θ_i . In that case, the maximum radiance, that is the histogram peak furthest to the right on the histogram, would be obtained at position 4 where $\theta_i = 0$. Instead, we see a slight increase at the (near) specular positions 5, 6, and 7, indicating that the reflection properties of the tile are diffuse with a small specular component.

The results for (b) crumpled paper (#28) can be neatly explained by thinking of a rough surface that is locally diffuse. There is no increase at the specular position; the maximum average radiance is reached at position 4. Due to roughness we have different amounts of foreshortening at different parts of the sample, and therefore a whole range of grey-values is present in the images. At position 8, the distribution becomes bimodal due to a shadowing effect. From then on, we can distinguish shadowed areas that are only irradiated

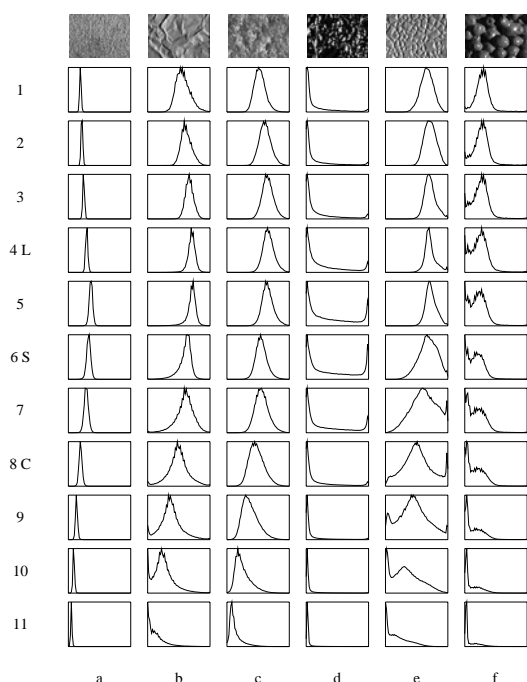


Fig. 4. Sets of 11 histograms for 6 samples shown in the top row. For the meaning of the 11 positions, see Figure 3. The sample is frontal to the lamp at position 4 (indicated with L); is at the specular point at position 6 (indicated with S); and is frontal to the camera at position 8 (indicated with C). The samples shown are (a) quarry tile (#25); (b) crumpled paper (#28); (c) white terrycloth (#3); (d) crumpled aluminium foil (#15); (e) rough plastic (#4); (f) a collection of painted spheres (#35).

through interreflections from directly irradiated areas. The category of rough surfaces that are locally diffuse is quite common in practice.

The results for (c) white terrycloth (#3) are very similar to those for crumpled paper. That is understandable: their surface structures are rough at the resolution of our images and at a microscale paper and terrycloth both consist of small fibers, leading to locally diffuse reflection properties. There is, however, a striking difference between the histograms from position 8 to 11. The crumpled paper clearly shows the effect of shadowing, whereas shadowing is totally absent for the terrycloth. This can be explained by the different nature of surface structures at the relevant scale. The fibers in terrycloth are much larger than those of paper. They are actually visible in the terrycloth image, so one could speak of mesoscale, while for the paper they are microscale. The crumpled paper can be understood in terms of a rough surface with *locally* (due to the cloud-like structure at microscale) diffuse reflec-




Fig. 5. Sets of 11 histograms for 6 samples shown in the top row. For the meaning of the 11 positions, see Figure 3. The sample is frontal to the lamp at position 4 (indicated with L); is at the specular point at position 6 (indicated with S); and is frontal to the camera at position 8 (indicated with C). The isotropic samples shown are (g) Soleirolia, a plant with small leaves (#53); (h) insulation (#27). For the anisotropic samples we show two histograms, obtained by rotating the samples around their normal for 90° or, in case of the linen sample, 45° . The samples are (i) corrugated paper (#38); (j) straw (#40); (k) linen (#44); (l) wood (#54).

tion properties. We refer to such an interface that is rough with no overhang (just one z -value per (x, y)) as a *rough proper surface*. Interfaces that consist of elements (for example small fibers) are, in general, not proper surfaces. For those cases, a description in terms of densities is often more useful. As radiation enters the interface, the density of material gradually increases. We refer to those surfaces as *clouds*. Terrycloth is a cloud surface at the given scale, and also locally diffuse. Light penetrates into the cloud and reaches any area because the cloud is not too dense and the terrycloth has a high albedo. As a result, there is no sharp distinction in the radiance from shadowed and directly irradiated areas and shadowed areas are not markedly visible in the image histograms. For accurate predictions of the reflection properties of cloud surfaces (and consequently their histograms) information about the scattering properties of the particles is required (van de Hulst, 1981). If the density of scatterers is high, mul-

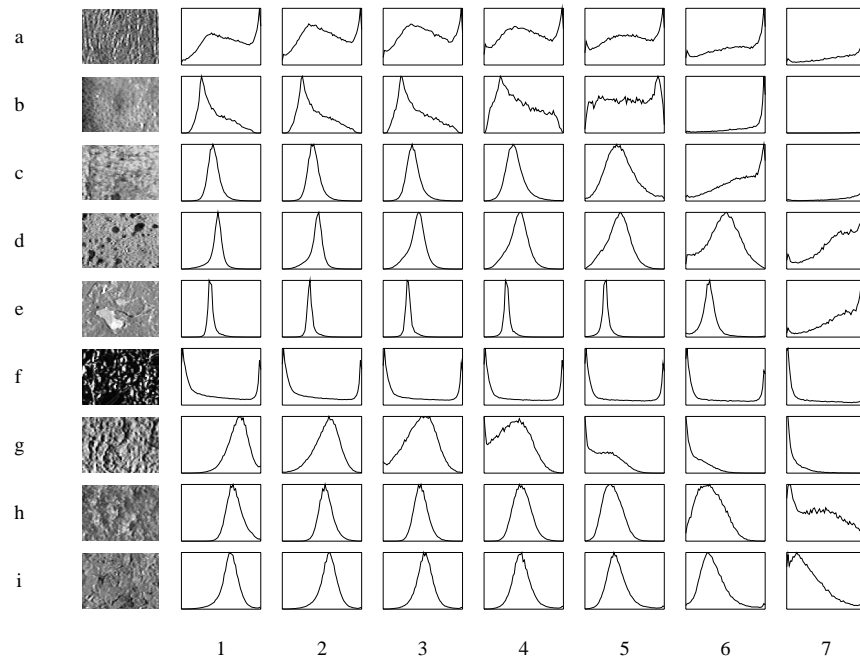


Fig. 6. Histograms of 9 samples shown at the left, at the 7 specular positions. See Figure 3 for an illustration of the measurement geometry. Rows are samples, column are the specular positions for detector positions 1 to 7. The incident angle of irradiation increases from 11.25° at detector position 1 (left column) with steps of 11.25° to 78.75° at detector position 7 (right column). The samples are (a) insulation (#27); (b) orange peel (#55); (c) human skin (#39); (d) sponge (#21); (e) limestone (#36); (f) crumpled aluminium foil (#15); (g) plaster (#11); (h) white terrycloth (#3); (i) salt crystals (#43).

multiple scattering becomes important. This is described by Kubelka-Munk theory (Kubelka and Munk, 1931) and studied by Chandrasekhar (Chandrasekhar, 1960).

The (d) crumpled aluminium foil (#15) has a surface structure that was deliberately prepared to be similar to the crumpled paper. Nevertheless, and not really surprising, the results are totally different. The aluminium foil is locally specular with a negligible diffuse component and the camera has a limited dynamic range. Thus we obtain extreme bimodal histograms with a high radiance peak (pixels at a specular position) and a low radiance peak (pixels *not* at a specular position) and hardly anything in between. The specular peak is strongest at the specular position and quickly decays.

Many surfaces have both a specular and a diffuse reflection component. The sample of (e) rough plastic (#4) is an example. Thus it can be seen as a combination of crumpled aluminium foil and crumpled paper. We observe specular peaks in histograms 7 to 9, together with a shifting Gaussian-like distribution typical of crumpled paper, and a clear shadowing peak at positions 9 to 11. From the fact that the shadowing peak is stronger than for the crumpled paper, we may con-

clude that the surface is rougher, that is, the standard deviation of local slopes will be larger. It might appear surprising that the maximum value of the specular peak is not attained at the specular position. This is due to the fact that this sample, as can be seen from the image in Figure 2, consists of pyramid-like patches, so there are no ‘flat’ parts that are specular at position 6.

Sample (f) is a collection of small balls, painted grey with a thin specular layer of veneer on top of it (#35). The balls are therefore specular and reflect isotropically as specular spheres do. Indeed, in almost any image one can observe highlights (like the image in Figure 2). This illustrates how surface roughness quickly masks a specular component. At a first glance, one might conclude that this sample yields results similar to the crumpled paper: a Gaussian-like distribution together with a strong shadowing peak. However, the shadowing peak is in this case of a fundamentally different nature. It is already visible at position 4, where the source is frontally to the detector so shadowing is impossible for proper surfaces. The low radiation peak comes from cavities between the balls. The peak disappears for histograms 1 to 3 when the sample is viewed from an

oblique angle but irradiated more or less frontally. At those positions the detector only sees the sides of the ball. This effect is purely configurational (due to the surface structure which contains overhang).

The same effect is visible for (g) *Soleirolia*, a plant with tiny round leaves (#53). The background behind the leaves has a low albedo and at most positions we mainly see the background. The result is an overall increase in radiation from position 4 down to 1, which is unusual since the amount of incident light decreases because of foreshortening. As a result, this type of samples has a very unusual BRDF, that can in no way be explained by conventional reflection models. Nevertheless, such texture are not uncommon (e.g. foliage). The requirements for effects like this to occur are a not too dense cloud-like surface structure, so that light penetrates into the material, and a background, or deep structure, with different reflection properties.

The histograms for (h) insulation (#27) are characterized by a large deviation between the specular position and the non-specular positions. This sample consist of a collection of glassy fibers oriented criss-cross in all directions. The fibers themselves are very specular and 'light up' when viewed from a specular position.

4.2. Anisotropic samples

The anisotropic samples provide an excellent way to illustrate the effect of surface structure on texture by comparing the two histograms shown in Figure 5, columns (i) to (l). They are obtained before and after rotating the sample around its normal. The applied rotation was 90° , except for the linen sample, which was rotated 45° . Because it is likely that these samples are not anisotropic at microscale, rotating them will not change local reflection properties. Thus the only change is a different surface structure relative to the irradiation and viewing direction. The results for samples (i), (j), and (l) clearly show the profound influence that surface structure can have on texture.

Sample (i), corrugated (ribbed) paper (#38), is an extreme case of an anisotropic surface. It has sinusoidal height variations in one direction and in the perpendicular direction the height is constant. One of the instances shown in the histogram corresponds to the case where the grooves are perpendicular to the z -axis in Figure 3. In that case there is hardly any structure visible, the texture is almost uniform and the histogram is a single peak. In the other case the grooves are parallel to the

z -axis and a periodical texture of stripes with a high and low radiation appears. In this situation there are always surface normals at a specular position and since the corrugated paper is partly specular (glossy), we see specular peaks in many histograms. The darker peaks (position 7 and higher) are shadow peaks.

The sample of (j) straw (#40) is a more natural version of the same sort of structure. It consists of tubes of straw aligned in roughly the same direction. The 'black' peaks are due to cast shadows, because the surface contains overhang, like the painted spheres in (f). Again there are large changes when the sample is rotated. If the tubes are parallel to the z -axis, the histogram is much broader because the tubes have a dark and a bright side and there is a specular peak.

The (k) linen sample (#44) is not very rough at the inner scale of the image. Therefore the histograms do not show much changes. Sample (l) is a piece of wood (#54). The grain of the wood yields a very anisotropic surface structure.

4.3. Specularities at angles of high incidence

Many surfaces become specular for large angles of incidence¹. This is understandable if the surface is a dielectric, smooth at microscale, since for these optical media the Fresnel coefficient, defined as the ratio of reflected to incident radiation, approaches 1 as $\theta_i \rightarrow 90^\circ$.

The higher the index of refraction, the sooner F approaches 1 for increasing θ_i . This can be seen from the formulas for the Fresnel coefficient F for unpolarized light as a function of θ_i and the index of refraction n (see e.g. (Kortüm, 1969)).

$$F(n, \theta_i) = \frac{1}{2} \left(\frac{\sin^2(\theta_i - \theta_t)}{\sin^2(\theta_i + \theta_t)} + \frac{\tan^2(\theta_i - \theta_t)}{\tan^2(\theta_i + \theta_t)} \right) \quad (1)$$

with $\theta_t = \arcsin\left(\frac{\sin \theta_i}{n}\right)$

Even if the surfaces is not smooth at microscale, it will be specular for grazing angles of incidence, as can be shown by an argument originally from Rayleigh. The rule of the thumb, known as the Rayleigh criterion, is that a surface may be considered smooth for

$$h < \frac{\lambda}{8 \cos \theta}, \quad (2)$$

where h is the characteristic height variation at microscale, λ is the wavelength of the irradiation, and θ is the polar angle of incidence. This effect is due to the wave character of light.

It turns out that the rate at which samples become specular for increasing θ_i , differs from sample to sample in the database and is therefore a characteristic material/surface attribute.

This is illustrated for 9 samples in Figure 6. Samples such as insulation (#27) and orange peel (#55) show specularities for small θ_i and become almost completely white for grazing incidence. For many samples a specular component does not show up for small θ_i , but is very obvious for large θ_i . This is the case for human skin (this sample was moist, which probably accounts for the effect), sponge and limestone. Samples with proper surfaces and high roughness, such as crumpled aluminium foil (#15) and plaster (#11), do not show the effect, of course because of shadowing. Terrycloth (#3) and salt crystal (#43) are two of the few samples which hardly show specularities even at detector position 7 ($\theta_i = 78.75^\circ$).

5. Simulations

In this section we show simulations of the bidirectional radiant power distributions by obtained by physics-based modeling. This allows for a quantitative comparison with the experimental results discussed previously. In the next section we discuss applications that make use of these models. We restrict ourselves to materials that can be described as isotropic surfaces with Gaussian roughness with local reflection properties described by an arbitrary BRDF. In particular we consider diffuse BRDFs, specular BRDFs and combinations of these. This model includes many common surfaces and samples in the database.

Rough surfaces that are locally diffuse, of which the crumpled paper from the previous section was the typical example, are an important category. Figure 7 shows 6 samples that fall into this category, ordered more or less increasing in roughness.

We have also shown examples of rough surfaces with local specular reflection properties and combinations of locally diffuse and specular reflection properties.

The surface roughness can be characterized by a statistical distribution of surface normals (this implies roughness at one scale: the macroscale). The local reflection properties can be characterized by a local BRDF (due to the surface structure and radiometric

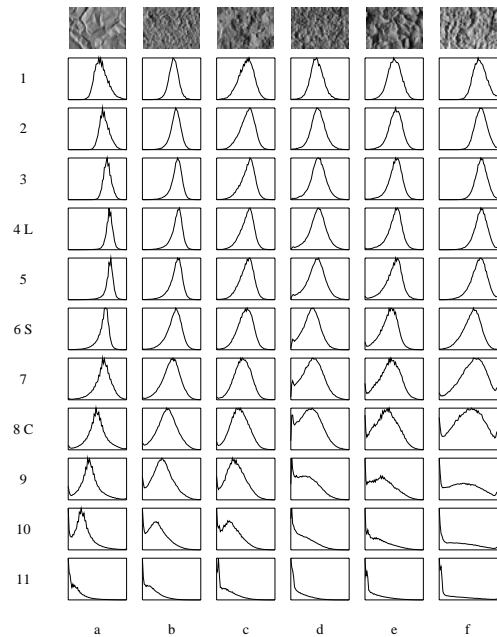


Fig. 7. Sets of 11 histograms for the 6 samples shown at the top. These samples are rough proper surfaces with locally dominantly diffuse reflection properties, ordered more or less increasing in roughness. For the meaning of the 11 positions see Figure 3. The sample is frontal to the lamp at position 4 (indicated with L); is at the specular point at position 6 (indicated with S); and is frontal to the camera at position 8 (indicated with C). The samples shown are (a) crumpled paper (#28); (b) concrete (#49); (c) concrete (#50); (d) pebbles (#8); (e) concrete (#45); (f) wall plaster (#11).

properties at microscale). Radiometric and geometrical effects that, apart from the surface structure and the local reflection properties, determine the histograms for these materials are

- Masking: parts of the surface are invisible from the detector position because they are occluded by other surface parts (masking by intersection) or because they are oriented away from the detector (self-masking).
- Shadowing is essentially the same effect: parts of the surface are not irradiated because of their orientation (self-shadowing) or because the incident beam is intersected (shadowing by intersection).
- Interreflections: parts of the surface irradiate other parts, thus capturing multiple scattering.

We have generated histograms for rough proper surfaces with local reflection properties, including the effects of masking and shadowing. The basic procedure

is as follows:

1. Pick a point with a surface normal, specified by its polar and azimuth angle (θ_a, ϕ_a) according to some distribution.
2. Accept the point only if it is not self-masked. Else go to 1.
3. Accept the point with the probability of it not being masked by intersection. Else go to 1.
4. If the point is self-shadowed, drop it in bin 0 and exit.
5. Else, drop the point in bin 0 and exit, with the probability of it being shadowed by intersection.
6. Else, calculate local angles of incidence and viewing, apply local BRDF and determine radiance. Multiply radiance by a gain factor and drop the point in the appropriate bin.

This procedure is repeated for many points (typically 5×10^4 in our simulations). We used histograms of 64 bins. Finally the histogram is smoothed by convolution with a Gaussian kernel. There are several reasons why smoothing is reasonable (and even necessary). (1) Interreflections will cause a redistribution of light over the sample, thus increasing – in a first order approximation – the radiation from all pixels by a certain amount. This will smooth the histogram. This effect is especially obvious for the shadowed areas. Without smoothing, these will be completely black. (2) In an experimental setting, the irradiation is only uniform within certain bounds. Therefore some parts of the image radiate slightly more than others. As a result, the histogram will be blurred. (3) Spatial albedo variations are ignored in the procedure outlined above, but will always exist in real samples. (4) Detector noise smooths the histograms.

We need to specify the distribution used in step 1. We assume that the height of the surfaces is normally distributed. This Gaussian surface model has been shown to be a good approximation for many surfaces. Using statistical theory, it can be shown that for such surfaces, the probability that a local surface normal lies in a solid angle (θ_a, ϕ_a) is given by

$$P_{d\omega_a}(\theta_a, r) d\omega_a = \frac{C(r)}{\cos^3 \theta_a} \exp\left(\frac{-\tan^2 \theta_a}{2r^2}\right) d\omega_a, \quad (3)$$

where $C(r)$ is a normalization constant and r is the root mean square slope, a parameter that indicates roughness. The distribution of ϕ_a is uniform, since we con-

sider isotropic surfaces, therefore ϕ_a does not appear in Equation 3. For derivation and details see (Ginneken et al., 1998) and (Middleton, 1960).

The (approximated) probability that a point is not shadowed by intersection, given that it is not self-shadowed, is given by

$$P_{ill}(\theta_i, r) = \frac{1}{1 + \Lambda(r, \theta_i)} \quad (4)$$

where

$$\Lambda(r, \theta_i) = \frac{r}{\sqrt{2\pi} \cot |\theta_i|} \exp\left(\frac{-\cot^2 \theta_i}{2r^2}\right) - \frac{1}{2} \operatorname{erfc}\left(\frac{\cot |\theta_i|}{\sqrt{2}r}\right). \quad (5)$$

The corresponding probability for the point not being masked by intersection is obtained by replacing θ_i by θ_r . For derivation and details see (Ginneken et al., 1998) and (Smith, 1967). Note that in reality, the probabilities for masking and shadowing by intersection are not independent, as is assumed here.

As local BRDF (step 6) we use either a constant value (Lambertian or perfect diffuse reflection) or the BRDF from the Torrance-Sparrow model (Torrance and Sparrow, 1967) with a small σ value. This corresponds to the BRDF of an almost perfectly smooth mirror².

Figure 8 shows the results of generated histograms for surfaces of increasing roughness that are locally Lambertian. We see Gaussian-like histograms, which become broader for oblique incidence and finally shadowed areas appear and start to dominate. This happens earlier and quicker with increasing roughness. One should compare the generated histograms with the histograms in Figure 7.

Figure 9 shows generated histograms in (a), (c), and (e), side by side with histograms from the database for (b) concrete (#49), (d) crumpled aluminium foil (#15) and (f) rough plastic (#4).

In order to model concrete, we used only local diffuse reflection properties. We empirically selected an appropriate roughness value r and gain and smoothing factor. So we used only 3 parameters (of which the gain is physically unimportant) for all histograms and yet obtain excellent agreement between simulations and experimental data. Note that the agreement would probably have been even slightly better had we used a non-linear fitting procedure to determine the optimal parameters.

Using only locally specular reflection, we have empirically selected parameter values for roughness r , the

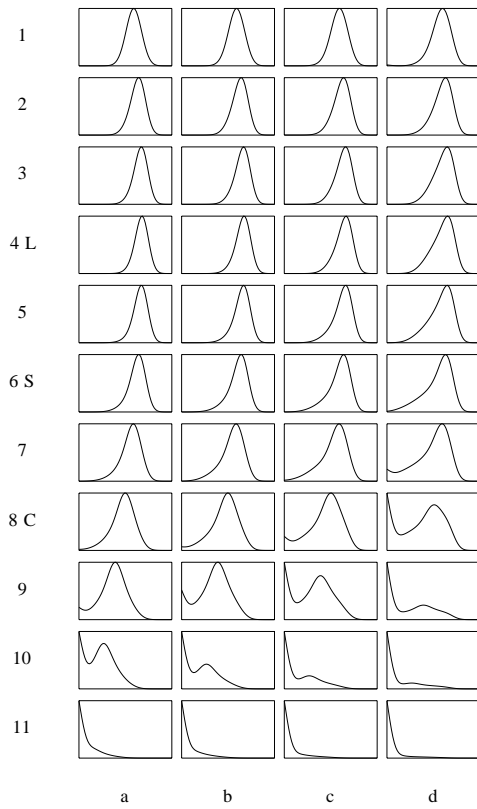


Fig. 8. Sets of generated histograms for rough surfaces that are locally diffuse. For the meaning of the 11 positions, see the Figure 3. The sample is frontal to the lamp at position 4 (indicated with L); is at the specular point at position 6 (indicated with S); and is frontal to the camera at position 8 (indicated with C). Smoothing and gain values used are the same for all histograms. The roughness values are (a) $r = 0.33$; (b) $r = 0.4$; (c) $r = 0.5$; (d) $r = 0.6$.

σ value in the Torrance-Sparrow model, gain and the width of the smoothing kernel that yield histograms comparable to the sample of crumpled aluminium foil.

By combining local specular reflection and diffuse reflection for a certain roughness value and appropriate values for gain and smoothing, we obtain histograms that are close to those of rough plastic. Deviations occur at the specular position because rough plastic does not show a maximum in the specular peak at that position, as a surface with normal height distribution would. This is explained by the fact that the rough plastic sample consists of pyramid-like patches, and thus is not a Gaussian surface. Still, the major changes in the histograms are qualitatively well predicted. In both the simulated and experimental data the histograms change from unimodal (diffuse peak) to bimodal (a

specular and a diffuse peak), become even trimodal (a shadowing peak appears) and unimodal again (only the shadowing peak remains).

6. Applications and implications

In this section, we discuss possible applications of BTFs, and bidirectional histograms in particular.

6.1. Estimation of material parameters

The simulations in Figure 9 show that by fitting a model to bidirectional histogram data, one can estimate mate-

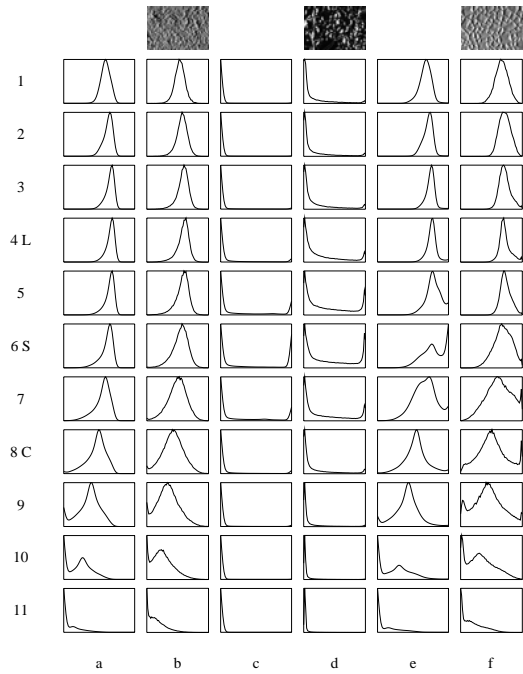


Fig. 9. Comparison of generated (columns without image) and measured (columns with sample image shown at the top) histograms. For the meaning of the 11 positions, see Figure 3. The sample is frontal to the lamp at position 4 (indicated with L); is at the specular point at position 6 (indicated with S); and is frontal to the camera at position 8 (indicated with C). (a) shows generated histograms with $r = 0.4$ and locally diffuse reflection. This roughness value was chosen empirically to obtain results similar those of concrete (#49), which are shown in column (b). (c) shows generated histograms with $r = 0.4$ and locally specular reflection, using the Torrance-Sparrow model with $\sigma = 0.13$. These values were empirically determined to obtain results similar to those of crumpled aluminium foil (#15), shown in column (d). (e) shows histograms with $r = 0.4$ and locally a linear combination of diffuse and specular reflection empirically selected to obtain results similar to rough plastic (#4), shown in column (f).

rial parameters such as roughness and local reflection properties. This is comparable to the common application of fitting a BRDF model to BRDF data. The obtained parameters can be useful as such, for instance to estimate the roughness of the material, or they can be used to predict the BRDF of the material. Since bidirectional histograms contain more information than BRDF data, we may expect better results from fitting to histograms, in the sense that the estimate of physical parameters will be more accurate and that the results are more robust.

6.2. Verification of texture models

In computer vision and computer graphics practice, texture is often treated as 2-D (flat texture; just albedo variations) and only perspective mapping is assumed or applied. When 3-D surface structure is approximately accounted for, the techniques employed do not always give rise to physically plausible results. A discussion of the pitfalls of common 3-D texture rendering schemes is given in (Koenderink and van Doorn, 1996). In the field of computer graphics, rendering textures with a high degree of realism is an important and challenging topic. Many attempts that use physics-based approaches have been made. (Westin et al., 1992) render velvet and cloth using ray-tracing with a specific geometric model to describe the material at the scale of the rendered texture and a BRDF to describe local reflection properties. In fact, this is very similar to the technique we used in our simulation of Gaussian rough surfaces, except for the fact that we used analytical expressions for the probability of masking and shadowing instead of performing ray-tracing. (Hanrahan and Krueger, 1993) render human skin and leaves as cloud densities, using linear transport theory to approximate the scattering of light in the cloud. They choose the properties of the clouds that make up skin and plant leaves heuristically, and spatially vary these properties to obtain textures which have the desired appearance. (Kajiya, 1989) models clouds of fur and hair using texels, which allows one to use ray-tracing techniques and local reflection models with a cloud surface. (Dorsey and Hanrahan, 1996) render metallic patinas, using a physics-based model of their surface structure and a BRDF based on the Kubelka-Munk model with empirically determined parameters. In the field of remote sensing, (Gerstl and Borel, 1992) show that vegetation canopy can be modeled as a collection of circular disks

with randomized offsets and can be rendered using radiosity methods.

Although some of the results of these methods look convincing, no attempts have been made to compare the results with measurements. (Gerstl and Borel, 1992) and (Westin et al., 1992) show BRDF data of their results, by averaging their textures and compare these to existing BRDF models. But in order to verify texture models, they should not be compared with BRDF data, but with BTF data. A practical problem is that the BTF is a matrix function and there is no evident metric to compare two textures (no two instances of a texture are the same). A good alternative is to compare the bidirectional histograms of model and experiment, as we have done in our simulations. Clearly, the verification that a texture model accurately predicts the radiant power distribution as a function of irradiation and viewing angle is much stronger than the fact that the mean intensity (BRDF) is in accordance with experimental data or known models. Our texture database can serve as a benchmark to test and compare proposed methods for texture rendering.

6.3. Texture recognition

Texture recognition is an important and active research area (see e.g. (Bovik et al., 1990) and the references therein). The ability to classify image parts according to their textural appearance also opens the way for segmentation methods based on texture recognition. We have shown that for many common materials, textural appearance depends on the irradiation and viewing geometry. This notion is important for the construction of texture classifications schemes that are invariant to these effects.

6.4. Estimation of irradiation and viewing geometry (Shape from texture)

The estimation of material parameters using histogram models is essentially a non-linear fitting problem: given a data set, values for the unknowns (the parameters in the histogram model) are computed for which the difference between model prediction and measured data is minimal. Such a procedure can be used equally well to estimate parameters characterizing the viewing geometry.

This is an example of shape from texture, a variation on shape from shading or photometric stereo, where

one generally uses BRDF data to convert radiance values into shape cues (Horn and Brooks, 1989). In this case, we convert histograms into shape parameters. The shape from shading problem is generally ill-posed. Given the BRDF of a material and source and detector position, there are usually many sample orientations that yield the same radiance measurement. Moreover, without an absolute measurement of the irradiation it is not possible to compute an orientation, or set of possible orientations from a single measurement. Instead, from the radiant power distribution of a single texture image of a material with known surface structure and local reflection properties, we can already estimate the sample normal direction. Consider the histograms in Figure 7, they are bimodal with a diffuse peak and a shadowing peak. The strength of both modes relative to each other, and the width of both modes show a clear correlation with the sample orientation.

For 2D texture, the bidirectional histograms do not show structural changes, but merely scale. In that case, fitting histogram models will be no more effective than using plain BRDF data. In other applications of shape from texture, the orientation is usually determined by estimating the perspective distortion of the texture image, e.g. (Aloimonos, 1988), (Gårding and Lindeberg, 1996). This is fruitful for 2D textures, and could be extended to 3D textures by combining it with fitting to bidirectional histograms. Alternatively, one could use other texture characteristics that depend on irradiation and viewing geometry, such as the power spectrum (Dana et al., 1997).

6.5. Texture rendering

We show an example of how the bidirectional radiant power distributions can be used for realistic texture rendering. In standard 2D texture mapping, a texture image undergoes a perspective transformation and its brightness is adjusted to account for illumination foreshortening. Essentially, this means that a Lambertian reflection model is used. If BRDF data is available, it can be used to obtain a more accurate brightness adjustment. Still, for 3D textures this procedure will not give satisfactory results.

Instead of adjusting the brightness of the whole texture patch by a constant factor, equal for each pixel, we propose to use a monotonic grey-level transformation that will ensure that the histogram of the transformed texture equals that of measured (or modeled) histogram data. This grey-level transformation can be obtained as

follows. First transform the texture image to account for a change in viewing direction. Let i denote image intensity (pixel value) and $p(i)$ the normalized radiant power distribution (histogram). The cumulative histogram is defined as $c(i) = \int_0^i p(i') di'$. Let $c_1(i)$ be the cumulative histogram of the transformed texture and $c_2(i)$ that of the desired histogram. Now change the pixel value i of every pixel in the transformed texture to i^* so that $c_1(i^*) = c_2(i)$. After this transformation the cumulative histograms are equal, and therefore the histograms as well. This transformation is indeed monotonic, since the cumulative histograms are monotonically increasing by definition.

Figure 10 shows some results using this method with a texture image of concrete (#45). We use an image that is illuminated under 45° and apply grey level transformations so that the histograms become equal to the corresponding images illuminated under 22.5° and 67.5° degrees. For comparison, the results that would have been obtained with standard 2D texture mapping where the brightness is adjusted according to measured BRDF data is also shown. The results of the grey-level transformation are superior to 2D texture mapping. The rendering of the image illuminated under 22.5° has a realistic “flat” appearance owing to the absence of shadowing. The rendering of the image illuminated under 67.5° has the same characteristic alternation of bright areas and dark shadowed regions as the actual image.

In Figure 11, we transform an image of terrycloth (#3) illuminated under 22.5° . Even the bottom-right image resembles the real data remarkably, although the irradiation has changed from almost frontal to near grazing incidence.

In practical applications, it is often not viable to obtain images for many irradiation and viewing geometries for all textures present in the scenes to be rendered. If detailed information about the geometry and local reflection properties is available together with accurate models regarding the image formation process, one may be able to generate these images. But even then, it may be impractical to store and retrieve all these image patches during rendering. In the method proposed here, one could store a small number of texture patches and a set of look-up tables that transform the textures according to the proper bidirectional radiant power distributions. The extra amount of memory required for the storage of look-up tables is negligible compared to storing extra images. One can use standard 2D texture mapping, followed by a change to the appropriate look-up table to account for 3D texture




Fig. 10. Texture rendering using look-up table manipulation based on bidirectional histograms. Left column shows texture images of concrete (#45) taken from the database viewed frontally, and illuminated from 45° (top), 22.5° (middle) and 67.5° (bottom) respectively. Middle column shows 2D texture mapping of the top left texture image so that it resembles the middle-left and bottom-left texture image. Brightness of the images is adjusted according to BRDF data. Right column shows the results of using the grey-level transformations of the top-left texture image, as described in the text. The middle-right and middle-left image, and the bottom-right and bottom-left images have identical histograms.

changes. Since look-up table manipulation is implemented in standard graphics hardware and thus very fast, this rendering scheme can easily be incorporated in real-time applications that now use 2D texture mapping.

Of course, adjusting the histograms and perspective of a texture cannot account for all the changes in visual appearance of texture. Detailed investigation of the rendered and measured textures will reveal differences in terms of spatial structure. In materials that contain specular parts, the exact location of highlights may be wrong, the relative positions and orientations of shadowed and directly irradiated areas in rough surfaces may have changed. Yet, to a casual observer both

patches will look about the same. The effect is not unlike techniques commonly used by visual artists. Their paintings of for example treetops can be highly realistic, although no branches or leaves have been painted. Apparently, differences in the spatial structure are not that important to give the right visual appearance, as long as a texture patch contains the proper size of textural elements (the grains) and distribution of intensities. This is exactly what our proposed method ensures, through a perspective transformation and a histogram adjustment, respectively.

More complex irradiation fields, with multiple sources and interreflections, can be handled by using a interpolating the appropriate look-up tables. If ex-




Fig. 11. Texture rendering using look-up table manipulation based on bidirectional histograms. Left column shows texture images of white terrycloth (#3) viewed frontally and illuminated from 22.5° (top), 45° (middle) and 67.5° (bottom) respectively. Middle column shows 2D texture mapping of the top left texture image so that it resembles the middle-left and bottom-left texture image. Brightness of the images is adjusted according to BRDF data. Right column shows the results of using the grey-level transformations of the top-left texture image, as described in the text. The middle-right and middle-left image, and the bottom-right and bottom-left images have identical histograms.

perimental data is not available, one could estimate material parameters such as roughness and local reflection properties and use a model like we have used in our simulations to approximate the bidirectional histograms.

7. Discussion and conclusions

The texture database convincingly shows the need to incorporate the dependency of texture on irradiation and viewing geometry in computer graphics and computer vision applications. We have shown that changes in irradiation and viewing directions induce large changes

in radiant power distributions. Summarizing the previous results, we may identify the following material attributes that play a crucial role:

- *Spatial (albedo) variations*: We have ignored these in our discussion, although the database contains some examples of samples with spatially varying pigmentation and samples that are made up of different components. Spatial variations yield important textures in practice. But there is no change in the histograms as a function of irradiation and viewing directions or these textures can be understood in terms of their parts or components.

- *Surfaces roughness*: If the surface structure is smooth at the resolution of the texture the BTF essentially reduces to the BRDF. Roughness increases the effects of foreshortening, masking, and shadowing. If the surface is very rough, the parts which are visible depend strongly on the viewing direction: For normal viewing angles the “low” parts are visible, for grazing angles only the “top” parts can be seen. If these parts have different reflection properties, large changes in the histogram can be expected. This may sound like a contrived example but such surfaces are in fact quite common. Examples from the database are the painted spheres (#35), plant leaves (#53), grass (#13), velvet (#7) and rugs (#18, #19).
 - *Local reflection properties*: The radiometric properties of the material below the inner scale of the texture are captured in the local reflection properties. Major dichotomy is between diffuse reflection (Lambertian reflection in the extreme case) and specular reflection. Combinations of diffuse and specular reflection are also common. In fact, one could specify a complete BRDF for the local reflection properties that, together with the surface structure at scale of resolution of the texture images, determines the BTF.
 - *Proper surfaces versus density clouds*: Surface which are cloud instead of ‘bulk material’ have different textural appearance and form a separate class. The distinction between clouds and rough surfaces with overhang is not sharp and depends on scale. Irradiation penetrates into clouds, which makes interreflections much more important. In general, the histograms of clouds show hardly any shadowing peaks. The precise textural appearance of materials in this class will be determined by the density of the cloud, the orientation of the parts and their geometry and nature. This geometry may be characterized by their aspect ratio (blobs, platelets, elongated structures). With nature we mean the way in which these part scatter irradiation. Examples of clouds from the database are rug (#18, #19), grass (#19), plant leaves (#53), the collection of painted spheres (#35), lamb-swool (#22) and rabbit fur (#24). In a way foams are *negative clouds* and belong to the same class. Examples of foams are white bread (#52), brown bread (#48), sponge (#21) and styrofoam (#21).
 - *Isotropic versus anisotropic surfaces*: Materials with an anisotropic surfaces essentially “change surface structure” when rotated around their surface normal. Since surface structure is such an important attribute for textural appearance, the impact of these changes can be huge, as is shown by the histograms of anisotropic samples such corrugated paper (#38), straw (#40), corn husk (#51) and wood (#54, #56).
 - *The rate at which the material becomes specular for grazing incidence*: Various physical mechanism account for this effect which turns out to be an important attribute that distinguishes the BTF of different materials.
 - *Albedo*: For bulk materials a high albedo increases the importance of interreflections relative to primary radiation. If the albedo is wavelength dependent, this may induce colour changes as a function of irradiation and viewing geometry. But the effect on grey-level histograms is not so clearly visible. For cloud materials, the albedo is a more important parameter that has a great effect on the effectiveness of the distribution of radiation through the cloud (the albedo is determined by the scattering properties of the parts of the cloud). In clouds that consist of high albedo parts, radiation penetrates further and this may make the difference between a shadowing peak or no shadowing peak in the histogram. For example, when the angle between irradiation and viewing direction is small, the samples sponge (#21), white bread (#52) and brown bread (#48) (all density clouds) do not show shadowing and are very different from rough bulk materials such as wall plaster. But at large angles between irradiation and viewing direction, the radiation has to travel further to reach shadowed areas. In this case brown bread (lower albedo, shadow peak in histogram) looks similar to wall plaster (#11) and very different from sponge and white bread (higher albedo, no shadow peak in the histogram).
- The boundaries between the categories presented in this discussion are far from sharp. We have merely enumerated material attributes that influence the BTF. The exact form of the BTF depends on the interplay of all these attributes. An accurate, widely applicable model will have to deal quantitatively with all these effects. We have demonstrated that these effects are clearly visible if one ignores the spatial structure of the textures and considers bidirectional radiant power distributions. Therefore, accurate histogram models will have to incorporate these effects as well.
- We have shown that a model that assumes Gaussian roughness and locally diffuse and/or specular reflection, and takes into account masking and shadowing effects, is able to accurately predict the quite compli-

cated changes in histograms as a function of irradiation and viewing geometry with only a few parameters. We have discussed applications of such models and proposed a new method of texture rendering using bidirectional histograms.

Acknowledgements

We thank Shree Nayar who participated in the development of the texture database and Astrid Kappers for carefully reading the manuscript. This research was supported by the ESPRIT Basic Research Action Realise of the European Community.

Notes

1. Even a piece of paper, often mentioned as an example of a Lambertian reflector, becomes specular at grazing incident angles.
2. We cannot use the expression for a perfect mirror in these simulations, because then we will *never* see any specularity because the local angles of incidence and viewing will never be *exactly* equal to the specular angles.

References

- Aloimonos, Y. 1988. Shape from texture. *Biological Cybernetics*, 58: 345–360.
- Beckmann, P. and Spizzichino, A. 1963. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon, New York.
- Bovik, A. C., Clark, M. and Geisler, W. S. 1990. Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(1): 55–73.
- Chandrasekhar, S. 1960. *Radiative Transfer*. Dover: New York.
- Cook, R.L. and Torrance, K.E. 1981. A reflectance model for computer graphics. *ACM Computer Graphics*, 15(3): 307–316.
- Dana, K.J., van Ginneken, B., Nayar, S.K. and Koenderink, J.J. 1996. Reflectance and texture of real-world surfaces. Columbia University, Technical Report CUCS-048-96.
- Dana, K.J., van Ginneken, B., Nayar, S.K. and Koenderink, J.J. 1997. Reflectance and texture of real-world surfaces. *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 151–157.
- Dorsey, J. and Hanrahan, P. 1996. Modeling and rendering of metallic patinas. *Proc. ACM SIGGRAPH*, pp. 387–396.
- Gårding, J. and Lindeberg, T. 1996. Direct computation of shape cues based on scale-adapted spatial derivative operators. *International Journal of Computer Vision*, 17(2): 163–192.
- Gerstl, S.A.W. and Borel, C.C. 1992. Principles of the radiosity method versus radiative transfer for canopy reflectance modeling. *IEEE Trans. on Geoscience and Remote Sensing*, pp. 271–275.
- Ginneken, B. van, Stavridi, M. and Koenderink, J.J. 1998. Diffuse and specular reflection from rough surfaces. *Applied Optics*, 37(1): 130–139.
- Hanrahan, P. and Krueger, W. 1993. Reflection from layered surfaces due to subsurface scattering. *Proc. ACM SIGGRAPH*, pp. 165–174.
- He, X.D., Torrance, K.E., Sillion, F.X. and Greenberg, D.P. 1991. A comprehensive physical model for light reflection. *ACM Computer Graphics*, 25(4): 175–186.
- Horn, B.K. and Brooks, M.J. (eds) 1989. *Shape from shading*. Massachusetts Institute of Technology.
- Kajiya, J.T. 1989. Rendering fur with three-dimensional textures. *Proc. ACM SIGGRAPH*, 23(3): 271–280.
- Koenderink, J.J. and van Doorn, A.J. 1996. Illuminance texture due to surface mesostructure. *Journal of the Optical Society of America A*, 13: 452–463.
- Koenderink, J.J., van Doorn, A.J. and Stavridi, M. 1996. Bidirectional reflection distribution functions expressed in terms of surface scattering modes. *European Conference on Computer Vision*, pp. 28–39.
- Kortüm, G. 1969. *Reflectance Spectroscopy*. Springer-Verlag, Berlin, Heidelberg, New York.
- Kubelka, P. and Munk, F. 1931. Ein Beitrag zur Optik der Farbenstriche. *Z. Tech. Phys*, 12: 593.
- Leader, J.C. 1979. Analysis and prediction of laser scattering from rough-surface materials. *Journal of the Optical Society of America*, 69(4): 610–628.
- Middleton, D. 1960. *Introduction to Statistical Communications Theory*. McGraw-Hill, New York.
- Nayar, S.K., Ikeuchi, K. and Kanade, T. 1991. Surface reflection: physical and geometrical perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13: 611–634.
- Nicodemus, F.E., Richmond, J.C. and Hsia, J.J. 1977. *Geometrical Considerations and Nomenclature for Reflectance*. National Bureau of Standards. Monograph No. 160.
- Oren, M. and Nayar, S.K. 1995. Generalization of the Lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14: 227–251.
- Phong, B.T. 1975. Illumination for computer generated pictures. *Communications of the ACM*, 18(6): 311–317.
- Poulin, P. and Fournier, A. 1990. A model for anisotropic reflection. *ACM Computer Graphics*, pp. 273–282.
- Richards, W.A. 1982. Lightness scale from image intensity distributions. *Applied Optics*, 21(14): 2569–2582.
- Smith, B.G. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation*, 15(5): 668–671.
- Tagare, H.D. and deFigueiredo, R.J.P. 1991. A theory of photometric stereo for a class of diffuse non-Lambertian surfaces. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 13(2): 133–152.
- Torrance, K. and Sparrow, E. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9): 1105–1114.
- van de Hulst, H.C. 1981. *Light scattering by small particles*. Dover: New York.
- Westin, S., Arvo, J.R. and Torrance, K.E. 1992. Predicting reflectance functions from complex surfaces. *Proc. ACM SIGGRAPH*, 26: 255–264.
- Wolff, L.B. 1994. Diffuse-reflectance model for smooth dielectric surfaces. *Journal of the Optical Society of America A*, 11(11): 2956–2968.