

Reflectance and Texture of Real-World Surfaces

Kristin J. Dana
Shree K. Nayar
Department of Computer Science
Columbia University
New York, NY 10027
dana@cs.columbia.edu
nayar@cs.columbia.edu

Bram van Ginneken
Jan J. Koenderink
Department of Physics
Utrecht University
3508 TA Utrecht, the Netherlands
b.vanginneken@fys.ruu.nl
j.j.koenderink@fys.ruu.nl

Abstract

In this work, we investigate the visual appearance of real-world surfaces and the dependence of appearance on imaging conditions. We present a BRDF (bidirectional reflectance distribution function) database with reflectance measurements for over 60 different samples, each observed with over 200 different combinations of viewing and source directions. We fit the BRDF measurements to two recent models to obtain a BRDF parameter database. These BRDF parameters can be directly used for both image analysis and image synthesis. Finally, we present a BTF (bidirectional texture function) database with image textures from over 60 different samples, each observed with over 200 different combinations of viewing and source directions. Each of these unique databases has important implications for a variety of vision algorithms and each is made publicly available.

1 Introduction

Characterizing the appearance of real-world surfaces is important for many computer vision algorithms. The appearance of any surface is a function of the scale at which it is observed. When the characteristic variations of the surface are subpixel, all local image pixels have the same intensity determined by the surface *reflectance*. The variation of reflectance with viewing and illumination direction is captured by the BRDF (*bidirectional reflectance distribution function*). If the characteristic surface undulations are instead projected onto several image pixels, there is a local variation of pixel intensity, referred to as image *texture*. The dependency of texture on viewing and illumination directions is described by the BTF (*bidirectional texture function*). This taxonomy is illustrated in Figure 1.

In this work we measure the BRDF of over 60 samples of rough, real-world surfaces. Although BRDF models have been widely discussed and used in vision (see [11],[17],[20],[8],[13]) the BRDFs of a

large and diverse collection of real-world surfaces have never before been obtained. Our measurements comprise a comprehensive BRDF database (the first of its kind) that is now publicly available at www.cs.columbia.edu/CAVE/curet. Exactly how well the BRDFs of real-world surfaces fit existing models has remained unknown as each model is typically verified using a small number (2 to 6) of surfaces. Our large database allows us to evaluate the performance of known models. Specifically, the measurements are fit to two existing analytical representations: the Oren-Nayar model [13] for surfaces with isotropic roughness and the Koenderink et al. decomposition [8] for both anisotropic and isotropic surfaces. Our fitting results form a concise BRDF parameter database that is also publicly available at www.cs.columbia.edu/CAVE/curet. These BRDF parameters can be directly used for both image analysis and image synthesis. In addition, the BRDF measurements can be used to evaluate other existing models [11],[17],[20] as well as future models.

While obtaining BRDF measurements, images of each real-world sample are recorded. These images prove valuable since they comprise a texture database, or a BTF database, with over 12,000 images (61 samples with 205 images per sample). Current literature deals almost exclusively with textures due to albedo and color variations on planar surfaces (see [19],[3],[7]). In contrast, the texture due to surface roughness has complex dependencies on viewing and illumination directions. These dependencies cannot be studied using existing texture databases that include few images (often a single image) of each sample (for instance, the widely used the Brodatz database). Our texture database covers a diverse collection of rough surfaces and captures the variation of image texture with changing illumination and viewing directions. This database is also available at www.cs.columbia.edu/CAVE/curet.

The measurements and model fitting results are pertinent to a variety of areas including remote-sensing, photogrammetry, image understanding and scene rendering. Important implications of this work for computer vision are discussed.




Figure 1. Taxonomy of surface appearance. When viewing and illumination directions are fixed, surface appearance can be described by either *reflectance* (at coarse-scale observation) or *texture* (at fine-scale observation). When viewing and illumination directions vary, the equivalent descriptions are the *bidirectional reflectance distribution function* (BRDF) and the *bidirectional texture function* (BTF). Analogous to the BRDF, the BTF is a function of four independent angles (two each for viewing and illumination directions).

2 Measurement Methods

Our measurement device is comprised of a robot¹, lamp², personal computer³, spectrometer⁴ and video camera⁵. Measuring the BRDF requires radiance measurements for a range of viewing/illumination directions. For each sample and each combination of illumination and viewing directions, an image from the video camera is captured by the frame grabber. These images have 640x480 pixels with 24 bits per pixel (8 bits per R/G/B channel). The pixel values are converted to radiance values using a post-processing calibration and segmentation scheme described in [4]. The calibrated, segmented images serve as the BTF measurements and these images are averaged to obtain the BRDF measurements.

The need to vary the viewing and source directions over the entire hemisphere of possible directions presents a practical obstacle in the measurements. This difficulty is reduced considerably by orienting the sample to generate the varied conditions. As illustrated in Figure 2, the light source remains fixed throughout the measurements. The light rays incident on the sample are approximately parallel and uniformly illuminate the sample. The camera is mounted on a tripod and its optical axis is parallel to the floor of the lab. During measurements for a given sample, the camera is moved to seven different locations, each separated by 22.5 degrees in the ground plane at a distance of 200 cm from the sample. For each camera position, the sample is oriented so that its normal is directed toward the vertices of the facets which tessellate the




Figure 2. Illustration of the discrete sample orientations, light source and camera positions used in the measurements. For each of the 7 camera positions illustrated, the robot orients the sample's global normal to the directions indicated by the vertices on the quarter-sphere. The illumination direction remains fixed.

fixed quarter-sphere illustrated in Figure 2. With this arrangement, a considerable number of measurements are made in the plane of incidence (i.e. source direction, viewing direction and sample normal lie in the same plane). Also, for each camera position, a specular point is included where the sample normal bisects the angle between the viewing and source direction. Sample orientations with corresponding viewing angles or illumination angles greater than 85 degrees are excluded from the measurements to avoid self-occlusion and self-shadowing. This exclusion results in the collection of 205 images for each sample. For anisotropic samples, the 205 measurements are repeated after rotating the sample about the global normal by either 90 degrees or 45 degrees, depending on the structure of the anisotropy.

3 Samples For Measurements

The collection of real-world surfaces used in the measurements are illustrated in Figure 3. Samples of these surfaces were mounted on 10x12 cm bases which were constructed to fit onto the robot gripper. Each sample, though globally planar, exhibits considerable depth variation or macroscopic surface roughness. The samples were chosen to span a wide range of geometric and photometric properties. The categories include specular surfaces (aluminum foil, artificial grass), diffuse surfaces (plaster, concrete), isotropic surfaces (cork, leather, styrofoam), anisotropic surfaces (straw, corduroy, corn husk), surfaces with large height variations (crumpled paper, terrycloth, pebbles), surfaces with small height variations (sandpaper, quarry tile, brick), pastel surfaces (paper, cotton), colored surfaces

¹SCORBOT-ER V by ESHED Robotec (Tel Aviv, Israel).

²Halogen bulb with a Fresnel lens.

³PC with Videomaker frame grabber by VITEC Multimedia.

⁴SpectraScan PR-704 by Photoresearch (Chatsworth, CA).

⁵Sony DXC-930 3-CCD color video camera.




Figure 3. The collection of 61 real-world surfaces used in the measurements. The name and number of each sample is indicated above its image. The samples were chosen to span a wide range of geometric and photometric properties. The categories include specular surfaces (aluminum foil, artificial grass), diffuse surfaces (plaster, concrete), isotropic surfaces (cork, leather, styrofoam), anisotropic surfaces (straw, corduroy, corn husk), surfaces with large height variations (crumpled paper, terrycloth, pebbles), surfaces with small height variations (sandpaper, quarry tile, brick), pastel surfaces (paper, cotton), colored surfaces (velvet, rug), natural surfaces (moss, lettuce, fur) and man-made surfaces (sponge, terrycloth, velvet). Different samples of the same type of surfaces are denoted by letters, e.g. Brick_a and Brick_b. Samples 29, 30, 31 and 32 are close-up views of samples 2, 11, 12 and 14, respectively.




Figure 4. Spheres rendered using the BRDF measurements obtained from camera position 1 (illumination at 22.5° to the right). Interpolation was used to obtain radiance values between the measured points.

(velvet, rug), natural surfaces (moss, lettuce, fur) and man-made surfaces (sponge, terrycloth, velvet).

4 BRDF Database

The BRDF measurements form a database with over 12,000 reflectance measurements (61 samples, 205 measurements per sample, 205 additional measurements for anisotropic samples). The measured BRDFs are quite diverse and reveal the complex appearance of many ordinary surfaces.

Figure 4 illustrates examples of spheres rendered with the measured BRDF as seen from camera position 1, i.e. with illumination from 22.5° to the right. Interpolation is used to obtain a continuous radiance pattern over each sphere. The rendered sphere corresponding to velvet (Sample 7) shows a particularly interesting BRDF that has bright regions when the global surface normal is close to 90 degrees from the illumination direction. This effect can be accounted for by considering the individual strands comprising the velvet structure which reflect light strongly as the illumination becomes oblique. This effect is consistent with the observed brightness in the interiors of folds of a velvet sheet. Indeed, the rendered velvet sphere gives a convincing impression of velvet.

The rendered spheres of plaster (Sample 30) and roofing shingle (Sample 32) show a fairly flat appearance which is quite different from the Lambertian prediction for such matte objects, but is consistent with [12] and [13]. Concrete (Sample 49) and salt crystals

(Sample 43) also show a somewhat flat appearance, while rough paper (Sample 31) is more Lambertian. The plush rug (Sample 19) and moss (Sample 61), have similar reflectance patterns as one would expect from the similarities of their geometry. Rendered spheres from two anisotropic samples of wood (Sample 54 and Sample 56) are also illustrated in Figure 4. The structure of the anisotropy of sample 54 consists of horizontally oriented ridges. This ridge structure causes a vertical bright stripe instead of a specular lobe in the rendered sphere. Sample 56 shows a similar effect, but the anisotropic structure for this sample consists of near vertical ridges. Consequently the corresponding rendered sphere shows a horizontal bright region due to the surface geometry.

5 Fitting to BRDF Models

A concise description is required for functional utility of the measurements. We employ the Oren-Nayar model [13] and the Koenderink et al. representation [9] to obtain parametric descriptions of the BRDF measurement database. The resulting database of parameters can be used directly in a variety of algorithms where accurate, concise and analytical reflectance descriptions are needed. In vision, these applications include shape-from-shading and photometric stereo. In computer graphics, the reflectance parameters are useful for realistic rendering of natural surfaces. As with the measurement database, the complete database of reflectance parameters is also available electronically.

5.1 BRDF Fitting Results

The 3 parameter Oren-Nayar model was fit to the 205 radiance measurements for each sample using the Levenberg-Marquardt method. The series used in the Koenderink et al. representation was truncated at order 2 and order 8, resulting in a 5 parameter and a 55 parameter representation, respectively. Linear estimation was used to fit the truncated series to the BRDF measurements.

Plots for the entire set of fitting results are too numerous to include here and are provided in [4]. The samples included here were chosen to be representative of the overall results. These samples are 11-plaster, 12-rough paper, 49-concrete, 25-quarry tile, 43-salt crystals, and 36-limestone. The measurements and modeling results are depicted in Figure 5. Row A of Figure 5 shows the raw measurement data plotted as a function of decreasing source angle, i.e. polar illumination angle, to emphasize the non-monotonic and non-Lambertian behavior of the reflectance. Rows B,C and D show the scatter-plots of the Oren-Nayar model fit, the order 2 Koenderink et al. fit and the order 8 Koenderink et al. fit, respectively. These scatter-plots show the 205 measurement values plotted as a function of the corresponding estimated value; a straight line indicates a perfect fit. Rows E,F and G show the residuals from




Figure 6. Cylinder rendered with 2D texture mapping (left) and 3D texture mapping (right). The 2D texture mapping was done by warping a frontal view image of the texture (with illumination at 22.5 degrees to the right). The 3D texture mapping was done by merging 13 images from the BTF of Sample 45 (concrete).

the Oren-Nayar model fit, the order 2 Koenderink et al. fit and the order 8 Koenderink et al. fit, respectively. The residuals are plotted as a function of increasing viewing angle to show the concentration of errors at oblique views for some of the samples. All measurements are shown in radiance with units of watts per steradian per square-meter.

For the Oren-Nayar model fits, the best performance is with diffuse samples, like plaster and concrete. Samples like quarry tile, which show sharp peaks in the radiance plots of Row A, have more fitting errors. For many of the fits, the average error is less than 10% of the peak reflectance value. For many applications this amount of error may be acceptable given the conciseness of this 3 parameter model.

It is useful to compare the performance of the order 2 (5 parameters) Koenderink et al. decomposition and the Oren-Nayar model since these representations have approximately the same number of parameters. The performance is better for most cases using the Oren-Nayar model. The low orders of the Koenderink et al. model are best used for Lambertian-type reflectance (with order 0 identical to Lambertian reflectance). As a result, samples like quarry tile and limestone are not well-represented.

The 55 parameter Koenderink et al. decomposition accurately represents each of the reflectance plots as shown by the near straight line scatter-plots for each sample in Row D. In fact, most of the samples in the database are represented accurately with this decomposition. Note that the behavior of the fit between measured points is not indicated by these plots. The possible presence of oscillations associated with high order fits needs further investigation.

6 Texture Database

The appearance of a rough surface, whether manifested as a single radiance value or as image texture, depends on viewing and source direction. Just




Figure 7. Changes in the spectrum due to changes in imaging conditions. (Top row) Two images of sample 11 with different source and viewing directions. (Bottom row) Fourier spectrum of the images in the top row, with zero frequency at the center and brighter regions corresponding to higher magnitudes. The orientation change in the spectrum is due to the change of source direction which causes a change in the shadow direction.

as the BRDF describes the coarse-scale appearance of a rough surface, the BTF (bidirectional texture function) is useful for describing the fine-scale appearance of a rough surface. Our measurements of image texture comprise the first BTF database for real-world surfaces. The database has over 12,000 images (61 samples, 205 measurements per sample, 205 additional measurements for anisotropic samples).

Important observations on the BTF can be made from the database. Consider texture mapping using Sample 45 (concrete) as shown in Figure 6. The differences in the 2D texture-mapped cylinder and the 3D texture-mapped cylinder (using the database images) are readily apparent. Because of the varying surface normals across the sample, foreshortening effects are quite complicated and cannot be accounted for by common texture-mapping techniques. A detailed discussion of the pitfalls of current texture rendering schemes is given in [9].

Consider the same sample shown under two different sets of illumination and viewing directions in Figure 7. The corresponding Fourier spectra are also shown in Figure 7. Notice that the spectra are quite different. Most of the difference is due to the change in azimuthal angle of the source direction which causes a change in the shadowing direction and hence a change in the dominant orientation of the spectrum. If the image texture was due to a planar albedo or color variation, changes in the source direction would not have this type of effect on the spectrum. Source direction changes would only cause a uniform scaling of the intensity over the entire image.

7 Implications for Vision

Our BRDF measurement database provides a thorough investigation of the reflectance properties of real-world rough surfaces. This database fills a long-standing need for a benchmark to test and compare

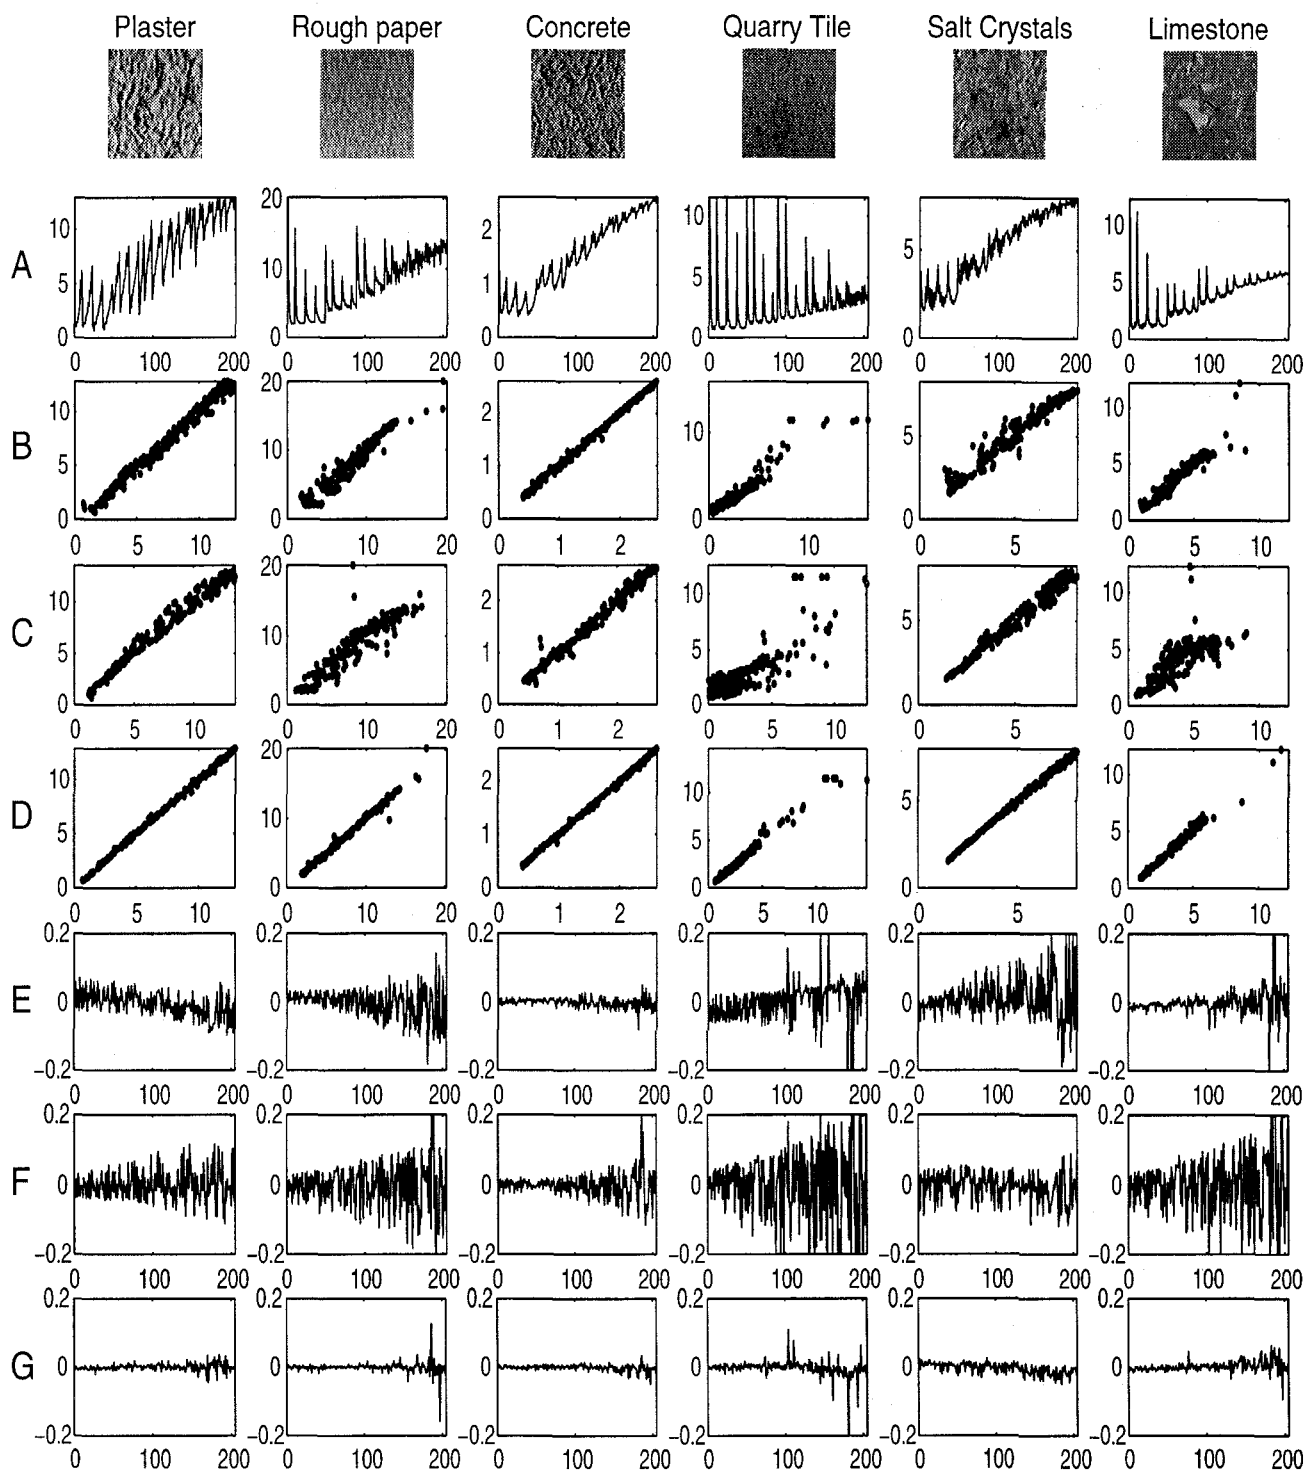


Figure 5. BRDF measurements and model-fitting results. Row A shows the raw measurement data plotted as a function of decreasing source angle to emphasize the non-Lambertian behavior of each sample. Rows B,C and D show the scatter-plots of the Oren-Nayar model fit (3 parameters), the order 2 Koenderink et al. fit (5 parameters) and the order 8 Koenderink et al. fit (55 parameters), respectively. Rows E,F and G show the residuals from the Oren-Nayar model fit, the order 2 Koenderink et al. fit and the order 8 Koenderink et al. fit, respectively. The residuals are plotted as a function of increasing viewing angle to show the concentration of errors at oblique views for some of the samples. The residuals shown are normalized by maximum measured radiance for that sample.

BRDF models as we have done here for the Oren-Nayar model and the Koenderink et al. decomposition.

Our BRDF parameter database, obtained by fitting the measurements to the Oren-Nayar model and the Koenderink et al. decomposition, can be used in place of the popular Lambertian reflectance model in such algorithms as shape-from-shading [6] and photometric stereo [21]. Since these algorithms rely on a reflectance model to ascertain shape, inaccuracies of the Lambertian model can significantly affect their performance. The model parameters can also be used instead of popular shading models [5],[2] for photorealistic rendering of real-world surfaces.

Since the parameter database covers two BRDF representations, a choice can be made to balance accuracy and conciseness. For isotropic surfaces the 3-parameter Oren-Nayar model can be employed. For isotropic and anisotropic surfaces, when a richer description can be afforded, the 55 parameter Koenderink et al. model can be used. For the 61 surfaces we have investigated, the parameters for both models are readily available.

Our BTF database is the first comprehensive investigation of texture appearance as a function of viewing and illumination direction. As illustrated in Figure 6 and Figure 7, surface roughness causes notable effects on the BTF which are not considered by current texture algorithms. Present algorithms for shape-from-texture [14],[16],[10], texture segmentation [22],[10] and texture recognition [15] are only suitable for 2D textures, i.e. planar texture due to albedo variation. Texture rendering also typically assumes a 2D planar texture that is mapped to a 3D surface. When the surface is rough, the rendering tends to be too flat and unrealistic. Texture analysis and synthesis of real-world rough surfaces remains an important unsolved problem. The database illustrates the need for 3D texture algorithms and serves as a starting point for their exploration.

Our BRDF measurement database, BRDF model parameter database and BTF measurement database together represent an extensive investigation of the appearance of real-world surfaces. Each of these databases has important implications for computer vision.

Acknowledgements

The research is sponsored in part by the National Science Foundation, DARPA/ONR under the MURI Grant No. N00014-95-1-0601 and by REALISE of the European Commission.

References

- [1] M. Born and E. Wolf, *Principles of Optics*, Pergamon Press, New York, 1959.
- [2] P. Bui-Tuong, "Illumination for computer generated pictures," *Communications of the ACM*, Vol. 18, pp. 311-317, 1975.
- [3] S. Chatterjee, "Classification of natural textures using Gaussian Markov random fields," *Markov Random Fields: Theory and Applications*, pp. 159-177, Academic Press, Boston, 1993.
- [4] K. J. Dana, B. Van Ginneken, S. K. Nayar and J.J. Koenderink, *Columbia University Technical Report CU-CS-048-96*, December 1996.
- [5] H. Gouraud, "Continuous shading of curves surfaces," *IEEE Trans. on Computers*, pp. 623-629, June 1971.
- [6] B.K.P. Horn and M.J. Brooks, *Shape from Shading*, MIT Press, Cambridge, Mass, 1989.
- [7] R.L. Kashyap, "Characterization and estimation of two-dimensional ARMA models," *IEEE Transactions on Information Theory*, Vol. IT-30, No. 5, September 1984.
- [8] J.J. Koenderink, A.J. van Doorn and M. Stavridi, "Bidirectional reflection distribution function expressed in terms of surface scattering modes," *European Conference on Computer Vision*, pp. 28-39, 1996.
- [9] J.J. Koenderink and A.J. van Doorn, "Illuminance texture due to surface mesostructure," *Journal of the Optical Society of America A*, Vol. 13 pp. 452-463, 1996.
- [10] J. Krumm and S.A. Shafer, "Texture segmentation and shape in the same image," *IEEE Conference on Computer Vision*, pp. 121-127, 1995.
- [11] S.K. Nayar, K. Ikeuchi and T. Kanade, "Surface reflection: physical and geometrical perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 7, pp. 611-634, July 1991.
- [12] S.K. Nayar and M. Oren, "Visual appearance of matte surfaces," *Science*, Vol. 267, pp. 1153-1156, Feb. 1995.
- [13] M. Oren and S.K. Nayar, "Generalization of the Lambertian model and implications for machine vision," *International Journal of Computer Vision*, Vol. 14, pp. 227-251, 1995.
- [14] M.A.S. Patel and F.S. Cohen, "Shape from texture using Markov random field models and stereo-windows," *IEEE Conference on CVPR*, pp. 290-305, 1992.
- [15] R.W. Picard, T. Kabir and F. Liu, "Real-time recognition with the entire Brodatz texture database," *IEEE Conference on CVPR*, pp. 638-9, 1993.
- [16] B.J. Super and A.C. Bovik, "Shape from texture using local spectral moments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 17, pp. 333-343, 1995.
- [17] H.D. Tagare and R.J.P. DeFigueiredo, "A framework for the construction of reflectance maps for machine vision," *CVGIP: Image Understanding*, Vol. 57, No. 3, pp. 265-282, May 1993.
- [18] K.E. Torrance and E.M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *Journal of the Optical Society of America*, Vol. 57, No. 9, pp. 1105-1114, 1967.
- [19] L. Wang and G. Healey, "Illumination and geometry invariant recognition of texture in color images," *IEEE Conference on CVPR*, pp. 419-424, 1996.
- [20] L.B. Wolff, "A diffuse reflectance model for smooth dielectrics," *Journal of the Optical Society of America A - Special Issue on Physics Based Machine Vision*, Vol. 11, pp. 2956-2968, November 1994.
- [21] R.J. Woodham, "Photometric methods for determining surface orientation from multiple images," *Optics Engineering*, Vol. 19, No. 1, pp. 139-144, 1980.
- [22] Z. Xie and M. Brady, "Texture segmentation using local energy in wavelet scale space," *ECCV*, Vol. 1, pp. 304-313, 1996.