

Mesostructure from Specularity

Tongbo Chen
MPI Informatik

Michael Goesele
University of Washington

Hans-Peter Seidel
MPI Informatik

Abstract

We describe a simple and robust method for surface mesostructure acquisition. Our method builds on the observation that specular reflection is a reliable visual cue for surface mesostructure perception. In contrast to most photometric stereo methods, which take specularities as outliers and discard them, we propose a progressive acquisition system that captures a dense specularity field as the only information for mesostructure reconstruction. Our method can efficiently recover surfaces with fine-scale geometric details from complex real-world objects with a wide variety of reflection properties, including translucent, low albedo, and highly specular objects. We show results for a variety of objects including human skin, dried apricot, orange, jelly candy, black leather and dark chocolate.

1 Introduction

The visual appearance of a real-world object is governed by reflectance properties, illumination condition, and a hierarchy of geometric components. In the geometric hierarchy, there are basically three different levels of scales, namely, the macrostructure level, the mesostructure level, and the microstructure level. The macrostructure level represents the gross surface geometry, typically expressed as a polygonal mesh or parametric spline surface. The microstructure level involves surface microfacets that are visually indistinguishable. The mesostructure level represents geometric details that are relatively small but still individually visible such as bumps or dents on a surface. Efficient mesostructure reconstruction methods can contribute greatly to high-quality graphics models in terms of fine-scale surface geometric details. An accurate and explicit mesostructure model can also benefit related mesostructure modeling techniques such as BTFs (Bidirectional Texture Functions) [4, 20].

State of the art high-resolution 3D scanning methods include [22, 5, 1, 15]. Photometric stereo methods can achieve high-resolution surface reconstruction with inexpensive setup [40, 27, 10, 9, 26, 41]. Existing techniques




Figure 1. Mesostructure reconstruction of orange skin. (a-d) Four cropped input images. (e) Recovered normal field (RGB-encoded). (f) Filtered normal field [35]. (g) Rendering of the normal field using Ward's isotropic BRDF model [14]. (h) Reconstructed 3D surface rendered at a novel viewpoint.

are, however, rarely able to capture the fine-scale details of real-world objects with translucency or highly specular reflection, such as skin, rough fruit skin, *etc.*

By drawing inspiration from photographs of real-world translucent objects and from the literature on human vision and perception [6, 24, 34, 7], we found that specular highlights are an important visual cue for surface mesostructure perception and a reliable visual information for surface detail representation. In Figure 1, the first row shows four images of a piece of orange skin under changing illumination. The small bumps on the orange skin introduce rich visual effects and can be efficiently revealed by specular highlights (see Figures 5, 6, 8, 7 for more examples). Based on this observation, we developed a simple and progressive system that uses specular highlights in order to solve the dense mesostructure reconstruction problem for a variety of real-

world complex objects, which possess a significant specular reflection component. Our method is largely independent of the underlying reflectance model, and can therefore successfully handle objects with complex reflectance that have previously been challenging.

To summarize our contributions: (1) We simplify the problem of mesostructure reconstruction from complex objects, *e.g.*, objects with translucency, or high specularity, which has up to now been expensive or even impossible to solve. (2) We use a dense specular field as the only reliable visual information for mesostructure reconstruction. (3) We developed a simple incremental and very flexible acquisition system. (4) We acquired high-quality mesostructure, for a variety of real-world objects including human skin, dried apricot, orange skin, jelly candy, black leather, and dark chocolate.

1.1 Overview

The remainder of this paper is organized as follows. Related work on surface mesostructure reconstruction and shape-from-specularity is discussed in Section 2. Section 3 describes our acquisition system and the mesostructure reconstruction method in detail. Extensive experimental results on mesostructure reconstruction for complex real-world objects are shown in Section 4. We conclude and discuss possible extensions in Section 5.

2 Related Work

Surface mesostructure is one of the key components of 3D texture [12]. It contributes strongly to the complex surface appearance of real-world objects. One method for modeling and rendering mesostructure is through BTFs (Bidirectional Texture Functions) [4], which can be regarded as a mapping from the 4D space of lighting and viewing directions to the space of 2D images. Most previous work on BTFs aims at capturing appearance data from natural materials and at efficient representation. Müller *et al.* [20] gives a comprehensive report on the state of the art of BTFs techniques. Liu *et al.* [16] used a shape-from-shading method to recover approximate 3D geometry of surface details from a BTF dataset. In [23], Neubeck *et al.* proposed a method for 3D texture reconstruction from extensive BTF data, with only a few and rather weak assumptions about reflectance and geometry. The reconstructed mesostructure can be used for the simplification of the BTF-based texture description and efficient compression of a BTF dataset. Even for the most advanced and expensive laser scanning systems, mesostructure reconstruction of highly specular or translucent objects is still a difficult problem. Most of the scanning technologies based on structured lighting based will also fail in reconstructing

fine-scale details for very low albedo, translucent, or highly-specular surfaces [22, 5]. To deal with highly-specular surface, Wang and Dana [37] presented a method that can simultaneously capture fine-scale surface shape and spatially varying BTFs by using a BTF measurement system. Similar to that work, our method will also depend on specular reflection. But we extend the idea to include not only highly-specular surface, but also very low albedo glossy or translucent glossy materials. Instead of using a complicated BTF measurement system, we developed a simple, flexible and progressive acquisition system. In [42], Yu and Chang introduced shadow graphs for 3D texture reconstruction. They show that the shadow graph alone is sufficient to solve the shape-from-shadow problem from a dense set of images. They also solved the problem of recovering height fields from a sparse set of images by integrating shadow and shading constraints. However, this method cannot work effectively for objects where shadow is no longer an accurate information, such as skin or fruit.

Photometric stereo methods [40, 27] are known to be able to capture fine-scale surface details and to provide an efficient alternative to BTF-based methods. In [10], Hertzmann and Seitz presented an example-based photometric stereo for shape reconstruction with general spatially varying BRDFs. They assumed that there are no cast shadows, no interreflections, and no subsurface scattering. Goldman *et al.* [9] proposed a photometric stereo method for iteratively recovering shape and BRDFs. They employed a local reflectance model, which cannot properly account for shadows, interreflections and subsurface scattering. In [26], Paterson *et al.* developed a simple system for BRDF and geometry capturing. Their system can handle a variety of real-world objects except highly specular or translucent materials. Wu and Tang [41] presented a simple dense photometric stereo method, using only a mirror sphere, a spotlight and a DV camera. They achieved surprisingly good results even with the presence of moderate shadows and specular highlights. To our knowledge, photometric stereo methods can rarely recover dense fine-scale surface details from translucent, highly specular, or low albedo glossy materials.

In [18], Magda and Zickler take advantage of Helmholtz reciprocity and light fields to reconstruct surfaces with arbitrary BRDFs. That method makes no assumption of the surface BRDF and works effectively for a variety of non-Lambertian surfaces (*e.g.* glossy surface), but not for highly translucent objects, where subsurface scattering dominates. In our approach, there is no explicit reflectance model assumed. We only exploit the specular reflection, which is directly related to surface geometry. Shape-from-specularity is a well-investigated field in computer vision. In contrast to most of the photometric stereo methods, where specular highlights are detected and separated as outliers, shape-from-specularity methods try to efficiently use the specular

reflectance component [25, 28, 29, 30]. Zheng and Murata [43] developed a special system to acquire gross shape from specular motion by using circular-shaped light source. Kutulakos and Steger [13] proposed an effective method for 3D shape reconstruction of refractive and specular objects by light-path triangulation. In [32] Solem *et al.* introduced variational analysis into shape-from-specularity and demonstrated the robustness of shape reconstruction from sparse specularity data. In [17], Lu *et al.* used specular reflections on surfaces to help modeling tangential hairs or grooves.

Motivated by the simplicity of existing shape-from-specularity methods, we extend them to reconstruct surface mesostructure, fine-scale geometric surface details, from objects with a significant specular component across the surface. A large set of real-world objects possess this property, such as fruit skin, human skin, plant leaves *etc.* Under the dichromatic reflectance model [31, 11], the light reflected from a surface comprises two physically different types of reflections, interface or surface reflection and body or subsurface reflection. The body part models conventional matte surfaces. Interface reflection that models highlights, is directly related to the surface (interface between object and air). We exploit only the specular reflectance, and developed a very simple system to capture dense mesostructure from complex real-world objects.

3 Mesostructure Acquisition

Most mesostructure acquisition systems are extensions of complicated and expensive BTF acquisition systems [37, 38, 23]. In contrast, we propose a simple, progressive, and flexible acquisition system for high-quality mesostructure recovery. The user can interactively adjust the number and distribution of light source positions without being constrained by the mechanical capabilities of the acquisition system.

3.1 Acquisition System

The basic setup of the mesostructure acquisition system consists of a digital camera and a point light source (see Figure 2). We used a 12-bit 1300×1030 -pixel Jenoptik ProgRes C14 digital camera for image acquisition and a 5 Watt Luxeon Star white LED as point light source. A checkerboard is used for camera calibration [3]. Four specular spheres are positioned at the four corners of the checkerboard to estimate the light source position. The sample object is placed on a small support at the center of the checkerboard. The camera faces downward to the checkerboard with optical axis perpendicular to the checkerboard plane. We keep the camera 1.5 meters away from the checkerboard. The mesostructure has ignorable magnitude, com-




Figure 2. Photograph of the acquisition system. The camera captures images while the LED is moved manually. A checkerboard is used for geometric calibration. Four specular spheres are employed for light source estimation. The sample object is placed at the center of the checkerboard. The camera is about 1.5 meters away from the sample. The light moves roughly on a virtual hemisphere with 1.5-meter radius and points always towards the sample object.

pared to such a large distance. We also assume the base geometry of the sample object has minute scale, compared to the distance between the camera and the object.

During acquisition, we keep the light source about 1.5 meters away from the object. This allows us to approximate the LED by a point light source. To keep the illumination consistent, we always point the light to the sample object. We capture one image for each position of the point light source. Using histogram thresholding, we can in real-time extract the specular reflection component of the sample object and update the specularity field, which keeps the state of how much specularity data has been captured from the sample object. Pixels for which a specular peak was detected in at least one image are marked red; otherwise they are marked black. During the acquisition, a growing por-

| sample object | number of input images |
|---------------------|------------------------|
| black leather | 35 |
| orange skin | 65 |
| jelly candy (piece) | 101 |
| jelly candy | 102 |
| human skin | 183 |
| dried apricot | 188 |
| dark chocolate | 196 |

Table 1. Numbers of input images captured for sample objects used in our experiments. Even for complex real-world objects, less than 200 images are sufficient for high quality mesostructure reconstruction.




Figure 3. Geometric constraint of specular reflection. At a specular point, the surface normal N is the bisector between the light vector L and viewing vector V . In one captured image, a large set of surface normals can be recovered simultaneously with known light source and camera positions.

tion of the specularity field will be colored red and the user can use this feedback to move the light source in a way that improves coverage of the specularity field.

This incremental refinement allows flexible control of the quality of the result mesostructure. If the final specularity field is very dense, the resulting mesostructure will be very accurate and highly-detailed. On the other hand, if only a sparse specularity field is captured, the reconstructed mesostructure will be dominated by low frequency features. In practice, less than 200 input images can already lead to a very dense specularity field and consequently highly-detailed mesostructure for complex real-world objects (see Table 1).

3.2 Light Source Estimation

Light source estimation from images is a well-researched problem in computer vision. In our experiments, we use

four specular spheres arranged in a rectangle on a common plane to estimate the light source position. Masselus et al. [19] use diffuse spheres in a similar configuration for this purpose. The law of reflection yields the geometric constraint that the specular normal N is a bisector between the light source direction L and the viewing direction V (see Figure 3), *i.e.*,

$$V = 2(N^T L)N - L. \quad (1)$$

With the known geometry of the camera and spheres, and the detected specular peaks on the spheres, we can easily compute the reflected light directions of the spheres according to Equation 1. Let the four reflected rays be $R_i = O_i + D_i t_i, i = 0, 1, 2, 3$, where O_i and D_i are the ray origin and ray direction respectively. Since there is only one point light source, we can construct an overdetermined linear system with six equations and four unknowns. Using SVD, we can compute the four unknowns $t_i, i = 0, 1, 2, 3$, and then the light source position.

3.3 Specularity Field Construction

Efficient and robust separation of diffuse and specular components of surface reflection for arbitrary materials – especially for translucent or refractive materials – is still an open problem. State of the art techniques exploit both polarization and color appearance [36, 33, 21, 39].

In experiments, we found that color appearance is no longer a reliable cue for specularity isolation, especially when translucent materials are involved. On the other hand, polarization imaging techniques pose high requirements for progressive and interactive acquisition. With these considerations, we used a simple histogram thresholding method to extract the specular reflection component.

During the progressive acquisition process, for each input image, we construct the intensity histogram as in Figure 4. Let $H(i)$ be the histogram, where i represents intensity and $H(i)$ is the number of pixels with intensity i in a specific image. We define histogram gradient as

$$G_h(i) = |H(i-1) - H(i)|. \quad (2)$$

For a specific gradient threshold T_g , we can find the associated intensity threshold T_i from the histogram $H(i)$. Then the specularity field, indicating the position of specular reflection, can be computed by thresholding the image according to T_i , *i.e.*, there is specularity if the pixel value is larger than T_i . Denote \mathcal{F}_s as the specularity field, which is defined as

$$\mathcal{F}_s(x, y) = \begin{cases} 1, & \text{specular highlight was observed at } (x, y) \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$




Figure 4. Specularity separation by histogram thresholding. (a) Input image. (b) Extracted specularity field. (c) Input image with specularity marked. (d) Histogram. With a user specified histogram gradient threshold $T_g = 10$, an intensity threshold $T_i = 491$ can be found in the histogram. The specularity field can be computed from the input image by simple thresholding. Note that different input images have the same T_g but different intensity thresholds T_i .

Each pixel (x, y) in the specularity field \mathcal{F}_s will be updated while the acquisition proceeds. The current state of the specularity field is displayed for active control of the light source. In our experiments, less than 200 images were required to recover almost complete specularity fields. For each pixel, there are usually multiple images in which it was classified as specular. We select the image where the brightest highlight occurred to calculate the normal. Images with overexposed pixels are avoided using the exposure settings of the camera.

3.4 Inferring Normal Field

With the known specularity field \mathcal{F}_s and the known light source positions for each input image, we can easily com-

pute the surface normal of each pixel (x, y) on \mathcal{F}_s (see Figure 3) yielding a normal field \mathcal{F}_n . In order to remove noise from \mathcal{F}_n , we filter it using the edge-preserving bilateral filter [35]. The resulting normal field can be used for direct rendering of mesostructure by bump mapping [2] or to reconstruct the 3D mesostructure. To simply convey the reconstructed mesostructure, we use Ward’s isotropic BRDF model [14] to render the normal field under directional illumination.

3.5 3D Mesostructure Reconstruction

To compute an accurate height field of the surface mesostructure from the estimated normal field we use the method proposed by Frankot and Chellappa [8] which uses surface smoothness as an integrability constraint.

Let $p = \partial z / \partial x$ and $q = \partial z / \partial y$ be the surface gradients in x and y directions respectively. $N(x, y)$ is the normalized surface normal at position (x, y) , which is related to the partial derivatives $\partial z / \partial x$ and $\partial z / \partial y$ through the formula:

$$N(x, y) = \frac{g}{\|g\|} \quad (4)$$

where $g = \{\partial z / \partial x, \partial z / \partial y, 1\}$. We will refer to the equality of the second order partial derivatives as the integrability constraint, *i.e.*,

$$\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x} \quad (5)$$

for all positions (x, y) on the support of the normal field. That is, they correspond to a surface with second order partial derivatives independent of the order of differentiation. Consequently, the surface height at any particular position is independent of the path of integration.

We tested this method [8] on all our experimental examples and got stable results even with a noisy input normal field. To increase performance we compute the integration through Fourier expansion.

4 Experimental Results

Previous mesostructure reconstruction techniques have difficulties with translucent, highly specular, or low albedo glossy materials. The major advantage of our method is that it can handle successfully these kinds of materials. In experiments, we choose a piece of dark chocolate and a piece of black leather, both of them have glossy reflection. To test translucent materials, we use orange skin, dried apricot, human skin, and jelly candy. The results demonstrate that our method can effectively reconstruct surface mesostructure from complex materials.

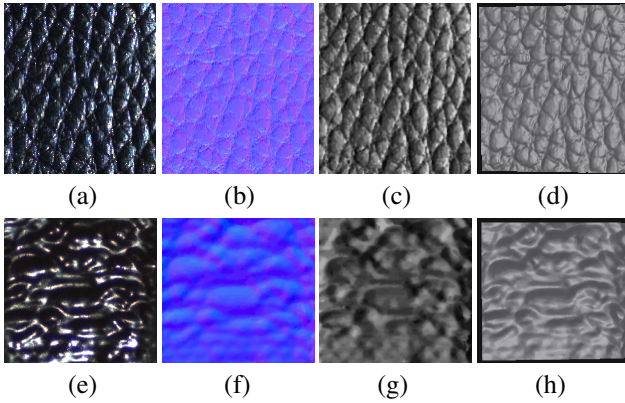


Figure 5. Mesostructure reconstruction of black leather and dark chocolate. (a, e) Input images. (b, f) Recovered normal fields. (c, g) Renderings of the normal fields using Ward’s isotropic BRDF model [14]. (d, h) Reconstructed 3D models rendered at novel viewpoints.

4.1 Very Low Albedo Glossy Materials

Figure 5 shows mesostructure reconstructions for a piece of black leather and a piece of dark chocolate. The chocolate has larger scale mesostructure than the leather. We successfully recover the large-scale shape of the chocolate and get a relatively smooth surface. With only 35 input images, we obtain accurately the fine-scale details of the leather.

4.2 Translucent Glossy Materials

Photometric stereo techniques are capable of measuring fine surface details. Effectively modeling surface details of translucent material is, however, still an open problem. We captured four real-world objects with various degrees of translucency: orange skin (Figure 1), dried apricot (Figure 6), human skin (Figure 7), and jelly candy (Figures 8 and 9).

Our system can effectively deal with the orange skin and the dried apricot whose fine-scale surface details are accurately recovered. It has, however, some difficulties with the jelly candy’s boundary, while the inner part of the jelly candy is successfully reconstructed. The reason for this is twofold. One is that the strong caustics around the jelly candy’s boundary makes accurate specular detection difficult. The other reason is that there are some surface normals on the boundary facing away from the camera, which cannot be captured with the current setup. To overcome this problem, we need additional cameras or capture several views of the object.

The skin case is notoriously known to be difficult in computer vision or graphics. The results in Figure 7 demon-

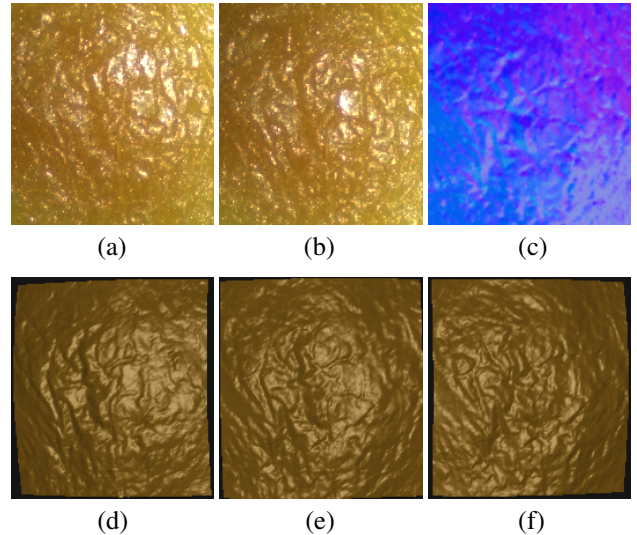


Figure 6. Mesostructure reconstruction of a piece of dried apricot. (a-b) Input images. (c) Recovered normal field. (d-f) Reconstructed surface rendered at different viewpoints.




Figure 7. Mesostructure reconstruction of human skin. (a-c) Input skin images. (d-f) Reconstructed surface rendered at novel viewpoints. Note that the pores and directional fine wrinkles are successfully reconstructed.

strate the effectiveness of our method. From the renderings at different viewpoints, we can clearly see the well-reconstructed fine wrinkles and pores. A promising extension is to measure mesostructure from general 3D skin, which can find important applications in computer graphics.

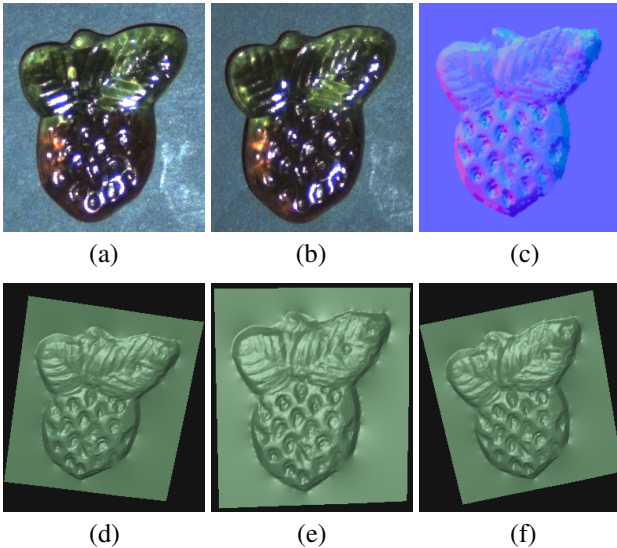


Figure 8. Mesostructure reconstruction of jelly candy. (a-b) Input images. (c) Recovered normal field. (d-f) Reconstructed surface rendered at novel viewpoints.




Figure 9. Mesostructure reconstruction of a piece of jelly candy. (a) Input image. (b) Shape by laser scanner. (c) Shape by laser scanner after covering the object with Lambertian powder, which is close to ground truth, but missed some fine details. (d) Shape by our method. It is very close to (c) but reveals more fine details.

We also captured a piece of highly translucent jelly candy with regular geometric pattern (see Figure 9). In order to compare against ground truth, we captured its surface structure also with a Minolta VI-910 laser range scanner. As expected, the laser scanner has difficulties recovering the shape because the laser light is scattered inside the object. The recovered position information is unreliable and the final geometry reconstruction shows strong artifacts. We then covered the jelly candy with a layer of very fine and homogeneous white powder in order to make it Lambertian and performed another laser range scan. Our mesostructure reconstruction compares favorably against this last scan,

which is the closest approximation of ground truth we were able to capture for this type of object.

5 Conclusion and Future Work

We presented a mesostructure reconstruction method that can successfully handle complex real-world objects. The main observation we build our approach on is that specular reflection reveals effectively fine-scale surface details and yields reliable visual information for mesostructure reconstruction. To exploit this observation, we developed a very simple system, including a camera, an LED light source, four specular spheres, and a checkerboard. The user can interactively control the light source to incrementally capture the specular field. The resulting dense specular field is used for mesostructure reconstruction. We demonstrated the effectiveness of our method by efficiently capturing the mesostructure of a variety of complex real-world objects, including human skin, dried apricot, orange skin, jelly candy, black leather, and dark chocolate.

In the future, we would like to use multiple cameras to capture a complete mesostructure model of 3D object with complex reflectance properties. Another extension of our method is a hybrid reconstruction using shape from specular when specular highlights are detected, and shape from photometric stereo for pixels below the threshold for specular highlights. Using the reconstructed mesostructure to improve the compactness of BTF representations and to realistically render highly detailed surfaces is a further research direction.

Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments. Thanks to Christian Fuchs for his help in preparing the rebuttal. We are grateful to Yan Wang for her assistance in the project.

References

- [1] F. Bernardini, H. Rushmeier, I. M. Martin, J. Mittleman, and G. Taubin. Building a digital model of Michelangelo's Florentine Pietà. *IEEE Computer Graphics and Applications*, 22(1):59–67, 2002.
- [2] J. F. Blinn. Simulation of wrinkled surfaces. *Computer Graphics*, 12(3):286–292, 1978.
- [3] J.-Y. Bouguet. MATLAB camera calibration toolbox. http://www.vision.caltech.edu/bouguetj/calib_doc/.
- [4] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, 1999.
- [5] J. Davis, D. Nehab, R. Ramamoorthi, and S. Rusinkiewicz. Spacetime stereo: a unifying framework for depth from triangulation. *PAMI*, 27(2):296–302, 2005.

- [6] R. Fleming and H. Buelthoff. Low-level images cues in the perception of translucent materials. *ACM Transactions on Applied Perception*, 2(3):346–382, 2005.
- [7] R. Fleming, A. Torralba, and E. H. Adelson. Specular reflections and the perception of shape. *Journal of Vision*, 4(9):798–820, 2004.
- [8] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *PAMI*, 10(4):439–451, 1988.
- [9] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying BRDFs from photometric stereo. In *Proceedings of ICCV*, pages 341–348, 2005.
- [10] A. Hertzmann and S. M. Seitz. Example-based photometric stereo: shape reconstruction with general, varying BRDFs. *PAMI*, 27(8):1254–1264, 2005.
- [11] G. Klinker, S. Shafer, and T. Kanade. Color image analysis with an intrinsic reflection model. In *Proceedings of ICCV*, pages 292–296, 1988.
- [12] J. J. Koenderink and A. Doorn. Illuminance texture due to surface mesostructure. *J. Opt. Soc. Am.*, 13(3):452–463, 1996.
- [13] K. N. Kutulakos and E. Steger. A theory of refractive and specular 3D shape by light-path triangulation. In *Proceedings of ICCV*, pages 1448–1455, 2005.
- [14] G. W. Larson. Measuring and modeling anisotropic reflection. *Comp. Graphics*, 26(2):265–272, 1992.
- [15] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital michelangelo project: 3D scanning of large statues. In *SIGGRAPH*, pages 131–144, 2000.
- [16] X. Liu, Y. Yu, and H.-Y. Shum. Synthesizing bidirectional texture functions for real-world surfaces. In *SIGGRAPH*, pages 97–106, 2001.
- [17] R. Lu, J. J. Koenderink, and A. M. L. Kappers. Specularities on surfaces with tangent hairs or grooves. *CVIU*, 78:320–335, 2000.
- [18] S. Magda and T. Zickler. Beyond Lambert: reconstructing surfaces with arbitrary BRDFs. In *Proceedings of ICCV*, pages 391–398, 2001.
- [19] V. Masselus, P. Dutré, and F. Anrys. The free-form light stage. In *EGRW*, pages 247–256, 2002.
- [20] G. Müller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein. Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics Forum*, 24(1):83–109, 2005.
- [21] S. K. Nayar, X. S. Fang, and T. Boulton. Separation of reflection components using color and polarization. *IJCV*, 21(3):163–186, 1997.
- [22] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi. Efficiently combining positions and normals for precise 3D geometry. *ACM Trans. Graph.*, 24(3):536–543, 2005.
- [23] A. Neubeck, A. Zalesny, and L. V. Gool. 3D texture reconstruction from extensive BTF data. In *Proceedings of International Workshop on Texture Analysis and Synthesis*, pages 13–18, 2005.
- [24] J. F. Norman. Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychol. Sci.*, 15(8):565–570, 2004.
- [25] M. Oren and S. K. Nayar. A theory of specular surface geometry. *IJCV*, 24:105–124, 1996.
- [26] J. A. Paterson, D. Claus, and A. W. Fitzgibbon. BRDF and geometry capture from extended inhomogeneous samples using flash photography. *Computer Graphics Forum*, 24(3):383–391, 2005.
- [27] H. E. Rushmeier, G. Taubin, and A. Guézic. Applying shape from lighting variation to bump map capture. In *Proceedings of EGWR*, pages 35–44, 1997.
- [28] A. C. Sanderson, L. E. Weiss, and S. K. Nayar. Structured highlight inspection of specular surfaces. *PAMI*, 10(1):44–55, 1988.
- [29] S. Savarese and P. Perona. Local analysis for 3D reconstruction of specular surfaces - part i. In *Proceedings of CVPR*, pages 738–745, 2001.
- [30] S. Savarese and P. Perona. Local analysis for 3D reconstruction of specular surfaces - part ii. In *Proceedings of ECCV*, pages 759–774, 2002.
- [31] S. A. Shafer. Using color to separate reflection components. *Color Research Application*, 10:210–218, 1985.
- [32] J. E. Solem, H. Aanaes, and A. Heyden. A variational analysis of shape from specularities using sparse data. In *Proceedings of 3DPVT*, pages 26–33, 2004.
- [33] R. T. Tan and K. Ikeuchi. Separating reflection components of textured surfaces using a single image. *PAMI*, 27(2):178–193, 2005.
- [34] J. T. Todd. The visual perception of 3D shape. *TRENDS in Cognitive Science*, 8(3):115–121, 2004.
- [35] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proceedings of ICCV*, pages 839–846, 1998.
- [36] S. Umeyama and G. Godin. Separation of diffuse and specular components of surface reflection by use of polarization and statistical analysis of images. *PAMI*, 26(5):639–647, 2004.
- [37] J. Wang and K. J. Dana. Relief texture from specularities. *PAMI*, 28:446–457, 2006.
- [38] J. Wang, X. Tong, J. Snyder, Y. Chen, B. Guo, and H.-Y. Shum. Capturing and rendering geometry details for BTF-mapped surfaces. *Visual Computer*, 21:559–568, 2005.
- [39] L. B. Wolff. Polarization-based material classification from specular reflection. *PAMI*, 12(11):1059–1071, 1990.
- [40] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [41] T. P. Wu and C. K. Tang. Dense photometric stereo using a mirror sphere and graph cut. In *Proceedings of CVPR*, pages 140–147, 2005.
- [42] Y. Yu and J. T. Chang. Shadow graphs and 3D texture reconstruction. *IJCV*, 62(1-2):35–60, 2005.
- [43] J. Y. Zheng and A. Murata. Acquiring a complete 3D model from specular motion under the illumination of circular-shaped light sources. *PAMI*, 22(8):913–920, 2000.