

Mesostructure from Specularity Using Gradient Illumination

Yannick Francken Tom Cuypers Philippe Bekaert
Hasselt University - tUL - IBBT
Expertise Centre for Digital Media, Belgium
{firstname.lastname}@uhasselt.be

Abstract

We present a method to efficiently acquire specular mesostructure normal maps, only making use of off-the-shelf components, such as a digital still camera, an LCD screen and a linear polarizing filter. Where current methods require a specialized setup, or a considerable number of input images, we only need a cheap setup to maintain a similar level of quality. We verify the presented theory on real world examples, and provide a ground truth evaluation on photorealistic synthetic data.

Keywords: mesostructure, specularities, normals, acquisition

1 Introduction

Throughout the years computer vision, and more in particular image based modelling, has become an increasingly important tool to facilitate the design of virtual clones of real-world objects. For example, these days real world scenes or persons are often scanned instead of being modelled manually to be used in games. However, up till now the bulk of this work still focusses on acquiring mainly diffuse and rather large objects. In many cases, small surface details are ignored.

However, a number of methods to acquire small-scale surface orientation already exist, storing the results in the form of normal maps. Typically they are mapped on a geometry proxy to provide local orientation information for each surface point to enhance the shading. Normal maps can also be transformed into depth maps [Frankot and Chellappa 1988].

Currently high quality methods for accurately acquiring normal maps exist, but these methods often require a delicate and/or expensive setup [Wang and Dana 2006; Ma et al. 2007]. Alternatively, there are also easy to use methods demanding no specialized hardware, but these typically yield low(er) quality scans [Rushmeier et al. 1997] or they are time-consuming procedures when high quality is needed [Chen et al. 2006; Francken et al. 2008b]. Hence, there is still need for a method that is *simple and fast*, yet yields *pleasing* results.

In this paper we focus on such a simple and fast method allowing for obtaining visually pleasing normal maps of specular mesostructures using common hardware components. Although this method is theoretically founded, it is not our goal to improve the accuracy of state-of-the-art normal map acquisition methods. Instead we aim to

improve the acquisition simplicity for physically and visually plausible normal maps, to make scanning available to a much broader range of users. Still, as we show in our comparative study, we achieve results of a when compared to state-of-the-art techniques.

2 Related Work

The proposed work can be categorized as a form of photometric stereo. Photometric stereo [Woodham 1980] and shape from shading [Horn 1975] are techniques to obtain surface orientation from a static scene and camera by controlling the scene's illumination. The original approaches work for diffuse objects, but many extensions have been proposed that allow for specular highlights [Ikeuchi 1981; Healey and Binford 1988; Zheng and Murata 2000; Mallick et al. 2005; Kutulakos and Steger 2005; Adato et al. 2007] and more general spatially varying BRDFs [Hertzmann and Seitz 2003; Goldman et al. 2005; Hertzmann 2005].

In our work we focus on small-scale surface orientation acquisition in the form of normal maps. A considerable number of techniques exist, and we will skip through them chronologically. An easy to use and fast method for scanning diffuse mesostructures was presented by [Rushmeier et al. 1997]. The idea behind this approach is very similar to that for photometric stereo [Woodham 1980]. [Malzbender et al. 2001] presented an alternative technique to obtain so called PTMs (Polynomial Texture Maps), which can be converted to normal maps. The method employs a specialized setup which captures about 50 images under different illumination. Another interesting scanning method is presented by [Han and Perlin 2003] and uses a kaleidoscope to generate views from multiple directions to obtain a BTF (Bidirectional Texture Function). [Paterson et al. 2005] use a digital still camera and a mounted flash to scan partly specular and diffuse mesostructures by interpreting the reflected flash light. [Wang and Dana 2006] obtain relief textures from specularities using a specialized hardware setup consisting of a camera, a parabolic mirror, a beam splitter and an illumination source with corresponding lens. A more practical technique is proposed by [Chen et al. 2006], which utilizes a manually moved light source and a fixed camera. Unfortunately, many input images are needed when high angular resolution normal maps are required. A faster method is proposed by [Ma et al. 2007], where they use an extended gradient light source requiring only 4 to 8 image captures at the cost of having a more delicate and expensive setup.

Currently computer screens are more and more often employed as a controlled planar illuminant, for example for the purpose of environment matting [Zongker et al. 1999]. Recently screen-camera('s) setups have also been introduced as a device for normal map acquisition [Morris and Kutulakos 2007; Francken et al. 2008b; Francken et al. 2008a; from Specularity Consistency 2008], and our work falls into this category. However, in this paper we will considerably reduce the necessary number of image captures by carefully choosing the light patterns.

Digital Still Camera + Linear Polarizer




Figure 1: Schematic setup overview. A gradient illumination pattern is displayed on the LCD screen and reflects off the specular mesostructure (spoon), partly into the camera lens. The surface normal \vec{n} can be found by taking the halfway vector between the view vector \vec{v} and the reflection direction \vec{r} .

3 Setup

In this section we will describe our normal map acquisition setup. As we focus on simplicity and wide applicability, we employ an easy to build setup, containing the following common pieces of hardware:

LCD screen: emits linearly polarized light which is exploited to separate diffuse from specular reflections.

Polarizing filter: blocks and unblocks specular highlights by taking advantage of the polarized light properties of the LCD screen.

Still camera: although it is perfectly possible to use a video camera, we chose to apply a moderate to high resolution digital still camera because they are much more commonly used than high quality video cameras.

A schematic overview of the setup is given in Figure 1.

Because the screen is used as a controllable light source that allows for detecting per pixel reflection directions, the position of the screen with respect to the camera has to be known. Also, we need the relation between the emitted and captured light to be linear. Therefore we need to perform both a geometric and a color calibration step in advance.

3.1 Geometric Calibration

Geometric calibration consists of determining the position and orientation of the screen with respect to the camera. Because the screen is not visible to the camera, a spherical mirror is moved in front of both components, in order to make them visible to each other. As such, the needed geometric relations can be obtained using known calibration techniques [Tarini et al. 2005; Francken et al. 2007]. In order to find the internal camera parameters and the mesostructure's supporting plane, we use a standard calibration toolbox which makes use of a checkerboard pattern [Bouguet 2006].




Figure 2: Spherical projection. Left: screen as a naive window to the virtual spherical area light source (top view). Middle: scaled illumination patterns for better intensity distribution (top view). Right: schematic representation of the setup (perspective view).

3.2 Color Calibration

As already mentioned, we require a linear relation between the amount of emitted and captured illumination. Therefore we first manually adjust the brightness, contrast and gamma parameters of the screen and video card drivers to allow for a maximal intensity range, keeping in mind not to introduce unwanted non-linearities. The unavoidable minor non-linearities are finally compensated for by constructing a simple lookup-table.

Note that in the early days, LCD screens typically suffered from emitting light in a view angle dependent fashion, which is undesirable in our case. However, throughout the years the effects of this problem have been severely reduced, so we reasonably assume that every pixel functions as a diffuse light source.

4 Normals From Gradients

In this section we will explain how gradient illumination patterns displayed on an LCD screen allow us to extract surface normal information. We will start by explaining how light polarization is exploited to separate diffuse from specular reflection, as only specular information is required by the method. Then we will illustrate how normals of specular surfaces can be reconstructed from observing light reflection in general, where more details will be given for gradient illumination patterns in particular. Finally, we will describe the relevant implementation concepts.

4.1 Separating Specular and Diffuse

We scan mesostructures based on available specular information. However, typical materials are not purely specular, but rather dichromatic; a combination of specular and diffuse components [Umeyama and Godin 2004]. Therefore we want to remove the diffuse part, which is accomplished by exploiting light polarization.

The idea behind this is that the polarization direction of specularly reflected polarized light remains approximately unaltered, whereas the diffuse reflection randomizes this direction. Hence, by recording the object illuminated by the vertically polarized light of the LCD screen, and aligning the filter with the polarization, half of the diffuse component together with the specular reflections are captured. Rotating the filter 90 degrees to a perpendicular direction results in the capture of only the other half of the diffuse light without the specular component. By subtracting the latter from the former, the necessary specularities are easily be extracted. More information on this topic can be found in [Wolff 1989b; Wolff 1989a; Nayar et al. 1997; Umeyama and Godin 2004; Iizuka 2002].




Figure 3: Procedure overview. From left to right, a piece of orange skin is illuminated by patterns P_i under both blocking and non-blocking polarization directions, yielding recordings L_i^b and L_i^{nb} . Then the specular images $L_i = L_i^{nb} - L_i^b$ are computed allowing for calculating the ratio images R_x and R_y . These are finally converted into the normal map n .

4.2 Normal From Specularity

Given a detected specularity together with the corresponding light source, camera, and mesostructure position, the corresponding normal can easily be reconstructed by applying the inverse law of reflection. This law states that the normal \vec{n} equals the halfway vector between the viewing vector \vec{v} and the reflection direction \vec{r} , as illustrated in Figure 1. This is a well-known technique [Chen et al. 2006; Sanderson et al. 1988] and many extensions have been proposed to speed-up this technique using controllable extended light sources [Ma et al. 2007; Francken et al. 2008b]. Our technique is most similar to the method of [Ma et al. 2007], but requires only a simple and cheap setup as in [Francken et al. 2008b] instead. Also, as we will show in section 5, we require less image captures.

4.3 Reflection Direction from Gradients

In this section we will explain how we efficiently detect the specular reflection direction for each pixel. The main idea is that a vertical and horizontal gradient pattern are displayed on the screen, and that we find out to which pixel (region) the reflection vector points to “pick” that intensity. We will derive the resulting formulas of this modification to the work of [Ma et al. 2007], making the technique applicable to a screen-camera setup.

4.3.1 Gradient Patterns

The gradient patterns we employ are constructed by projecting a linear dark to bright pattern of a virtual surrounding spherical area light source to the screen, emulating a piece of this virtual sphere centered around the object. In order to exploit an as broad as possible range of intensity values, the linear patterns are optimally adjusted to the position of the screen which simulates a virtual view on the virtual light source, as shown in Figure 2. The exact mathematical definition is given by the following equations, where P_x

and P_y are the gradient in x and y direction respectively, and P_c is the constant white pattern:

$$P_x(\vec{\omega}) = \frac{1}{2} \left(\frac{\omega_x}{\sin(\sigma_w)} + 1 \right) \quad (1)$$

$$P_y(\vec{\omega}) = \frac{1}{2} \left(\frac{\omega_z}{\sin(\sigma_h)} + 1 \right) \quad (2)$$

$$P_c(\vec{\omega}) = 1 \quad (3)$$

$\vec{\omega} = (\omega_x, \omega_y, \omega_z)$ represents the normalized incident illumination direction, and $(2\sigma_w, 2\sigma_h)$ is the window’s width and height expressed in radians (Figure 2).

4.3.2 Reflection Direction

In order to find for each pixel the reflection direction to extract a normal map, the object to be captured is illuminated by the three patterns P_i with $i \in \{x, y, c\}$. The reflected amount of light L_i in the viewing direction \vec{v} is given by:

$$L_i(\vec{v}) = \int_{\Omega} P_i(\vec{\omega}) S(\vec{r}, \vec{\omega}) F(\vec{\omega}, \vec{n}) d\vec{\omega} \quad (4)$$

where $\vec{r} = 2(\vec{r} \cdot \vec{v})\vec{n} - \vec{v}$. S represents the specular isotropic lobe, and $F(\vec{\omega}, \vec{n}) = \max(\vec{\omega} \cdot \vec{n}, 0)$ is the foreshortening factor. Notice that we assume that there is no interreflection, nor self-shadowing and that the recording is executed in a dark room without any stray light.

First we will execute a coordinate transform T to align \vec{r} with $\vec{z} = [0, 0, 1]$ to facilitate further derivations. The transform T is

computed by rotating the integration domain Ω to Ω' . The corresponding matrix is defined as $T = [\vec{s}, \vec{t}, \vec{r}]^t$ where \vec{s} , \vec{t} and \vec{r} are orthogonal vectors with respect to each other.

The functions S and F are rotationally invariant because they only depend on the *angles* between the argument vectors. This yields $S(\vec{r}, \vec{\omega}) = S(\vec{r}', \vec{\omega}')$ and $F(\vec{\omega}, \vec{n}) = F(\vec{\omega}', \vec{n}')$ where the added apostrophe refer to the rotated versions of the vectors in the original integration domain. The pattern P_i is not invariant to rotation, so T has to be taken into account as follows: $P(\vec{\omega}) = P(T^{-1}\vec{\omega}') = P(T^t\vec{w}')$. As such, the equation is now written as:

$$L_i(\vec{v}) = \int_{\Omega'} P_i(T^t\vec{\omega}') S(\vec{z}, \vec{\omega}') F(\vec{\omega}', \vec{n}') d\vec{\omega}' \quad (5)$$

For the remainder of this section we will assume a narrow specular lobe S . Therefore $F(\vec{\omega}', \vec{n}')$ can reasonably be assumed constant (c_F) in the small solid angle of S . This assumption breaks down in the case of grazing angles, but this is not a problem in practice because of the structure of our setup.

Further simplifications will depend on the illumination pattern P_i that is used. We derive L_x for the x -gradient, which is analogous to that of P_y . After filling in P_x , isolating some constants and splitting the integral, we obtain:

$$L_x(\vec{v}) = \frac{c_F}{2 \sin(\sigma_w)} \int_{\Omega'} (\omega'_s s_x + \omega'_t t_x + \omega'_r r_x) S(\vec{z}, \vec{\omega}') d\vec{\omega}' + \frac{c_F}{2} \int_{\Omega'} S(\vec{z}, \vec{\omega}') d\vec{\omega}' \quad (6)$$

Again by taking into account the narrowness of S , we can state that $(\omega'_s s_x + \omega'_t t_x) S(\vec{z}, \vec{\omega}') \approx 0$. This can easily be seen from the fact that $\vec{\omega}$ has almost to be aligned to have some response, meaning that r_x is almost equal to 1 and so s_x and t_x are negligible. Analogously the simplified integral $\int_{\Omega'} \omega'_r S(\vec{z}, \vec{\omega}') d\vec{\omega}'$ can be approximated by $\int_{\Omega'} S(\vec{z}, \vec{\omega}') d\vec{\omega}'$. After applying these steps to both the horizontal and vertical patterns we obtain the following simple equations:

$$L_x(\vec{v}) = \frac{L_c(\vec{v})}{2} \left(\frac{r_x}{\sin(\sigma_w)} + 1 \right) \quad (7)$$

$$L_y(\vec{v}) = \frac{L_c(\vec{v})}{2} \left(\frac{r_y}{\sin(\sigma_h)} + 1 \right) \quad (8)$$

Notice that L_c is the image taken under constant (or flood lit) illumination:

$$L_c(\vec{v}) = \int_{\Omega} P_c(\vec{\omega}) S(\vec{r}, \vec{\omega}) F(\vec{\omega}, \vec{n}) d\vec{\omega} \quad (9)$$

$$= c_F \int_{\Omega'} S(\vec{z}, \vec{\omega}') d\vec{\omega}' \quad (10)$$

This factor is only necessary because of the structure of the patterns due to the impossibility of emitting negative light.

As all the required parameters for equation (7) and (8) are given except from r_x and r_y , and knowing that r is a normalized vector, the reflection direction r is found.

technique	# image captures	cheap/simple
[Morris and Kutulakos 2007]	$> 10^3$	yes
[Wang and Dana 2006]	$> 10^3$	no
[Chen et al. 2006]	35 – 200	yes
[Malzbender et al. 2001]	50	no
[Francken et al. 2008b]	10 – 30	yes
[Paterson et al. 2005]	8 – 15	yes
[Ma et al. 2007]	4 – 8	no
Our approach	3 – 6	yes

Table 1: Comparison of number of image captures necessary for reconstructing a mesostructure.

4.4 Implementation

In practice the proposed procedure allows for a straightforward and efficient implementation. For generating our results we have created a proof of concept implementation in Matlab as well as an optimized C++ version. The important concepts of the method are summarized in a five-step process:

1. Record object illuminated by P_x , P_y and P_c under blocking (L_x^b , L_y^b and L_c^b) and non-blocking (L_x^{nb} , L_y^{nb} and L_c^{nb}) polarizing filter orientations
2. Determine specular images L_x , L_y and L_c :

$$L_x = L_x^{nb} - L_x^b \quad (11)$$

$$L_y = L_y^{nb} - L_y^b \quad (12)$$

$$L_c = L_c^{nb} - L_c^b \quad (13)$$

3. Determine ratio images $R_x = \frac{L_x}{L_c}$ and $R_y = \frac{L_y}{L_c}$
4. For each pixel, find r :

$$r_x = \sin(\sigma_w)(2R_x - 1) \quad (14)$$

$$r_y = \sin(\sigma_h)(2R_y - 1) \quad (15)$$

$$r_z = \sqrt{1 - r_x^2 - r_y^2} \quad (16)$$

5. For each pixel, calculate the normal \vec{n} (halfway vector between viewing dir \vec{v} and reflection dir \vec{r})

$$\vec{n} = \frac{\vec{r} + \vec{v}}{\|\vec{r} + \vec{v}\|} \quad (17)$$

The overview shows the simplicity of the method, where only six input photographs and a few simple image operations suffice to obtain a normal map. Notice that in practice for “pure” specular materials, even three image recordings could suffice because then the diffuse images will be black. A graphical overview of our algorithm is depicted in Figure 3.

5 Results

In this section our presented method will be evaluated on real as well as on synthetically generated data sets. The real data is captured using a Canon EOS 400D camera, a cheap linear polarizing filter and a standard 19 inch LCD screen. For evaluating the real data, we compared our results with the technique of [Francken et al. 2008b], which uses Gray coded patterns instead of gradients



Figure 4: A picture of a glossy plastic button and its corresponding normal map.

and can easily be applied to the same setup. The generated photo-realistic synthetic data is rendered using the physically based renderer PBRT [Pharr and Humphreys 2004]. We apply the Cook-Torrance reflectance model with different parameter settings. Miscellaneous scans of different mesostructures are depicted in Figure 6.

5.1 Efficiency

First we will evaluate the efficiency of our method by comparing to relevant previous work. We will focus on the number of image acquisitions necessary to scan a moderate mesostructure. An overview is given in Table 1. Clearly our method has a high performance when considering the number of required input images, while maintaining a pleasing quality, as will be shown later.

We will now compare our gradient based work to the Gray code based work, where exactly the same setup is employed. First note that we are limited to sampling “only” $256 \times 256 = 65536$ different light source positions because of the 8 bit color depth of the LCD screen. As such we are *intensity*-bound, whereas [Francken et al. 2008b] are *screen resolution*-bound (typically more than one million samples are possible). However, in practice 256×256 seems to suffice more than adequately.

More important is the fact that the acquisition time of our technique is *independent* of the number of light sources we want to sample (N). We only need $O(1)$ inputs compared to $O(\log N)$, which is a considerable speedup for a large number of inputs. If for example 15000 samples are desired, we still need only 6 input images, where the Gray code approach needs about 30.

5.2 Glossiness

Theoretically our technique works only for perfect reflectors (see section 4). However, in practice it seems to behave well for glossier materials as well. An example of a normal map scan of a glossy plastic button is given in Figure 4.

In order to analyse the loss of quality in function of the increasing glossiness (or roughness), we have provide a ground truth evaluation on photo-realistic synthetic. In addition, a comparative evaluation is made for the use of Gray coded patterns. Figure 5 shows a plot of the average angular normal error with respect to the ground truth, for both the gradient and Gray coded illumination. This figure clearly shows that, except from extremely specular materials, overall the gradient patterns outperform the Gray coded ones. This can be explained from the fact that broader specular lobes cause a more extensive pattern convolution, which more or less keeps approximating the wanted center of the specular lobe (which is the reflection direction). This is not the case for Gray code patterns,

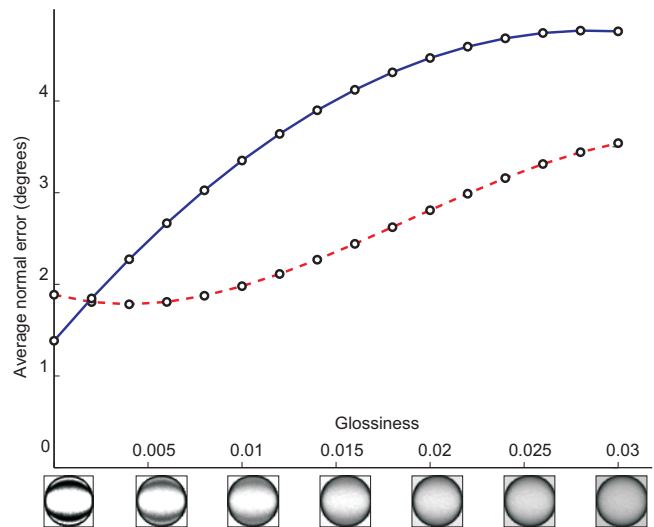


Figure 5: Error comparison for glossy materials on synthetic examples under gradient and Gray coded illumination patterns. The average angular error of the normals is plotted in function of the glossiness. The dashed curve shows the error of our gradient method, where the solid curve shows the error of Gray coded patterns.

where the convolution renders the patterns at a certain refinement level totally useless.

In conclusion we note that, where in our approach the normal maps will become blurrier for glossier materials, in the Gray code approach they will become noisier, which is typically worse. Noise in a normal map may cause visually distracting artefacts when the normal map is used for relighting or for generating a depth map. A blurrier normal map just over-smooths the surface, which is less disrupting.

5.3 Sensitivity to Noise

The main drawback of the use of gradient patterns instead of black/white binary patterns is that they are more sensitive to sensor noise. However, the rather limited amount of input images allows for longer exposure times, avoiding deviating normals due to noise. Still, it remains a more sensitive process.


5.4 Occlusion

Self-shadowing, or light that cannot reach the surface because it is occluded by the surface itself, causes dark regions where we cannot gather any information from. However, this is a general disadvantage of all shape from reflection methods, and not especially of this method. The only thing we do about it is build up a confidence map based on the per pixel darkness of the fully lit image (L_c). This tells us how usefull every pixel is and we take this into account when cleaning up normal maps (bridge minor lapses, remove noise etc.).


6 Conclusions and Future Work

We have presented a method that efficiently acquires specular mesostructure normal maps, only making use of off-the-shelf hardware components, namely a digital camera, an LCD screen and a linear polarizing filter. Whilst current methods often require spe-


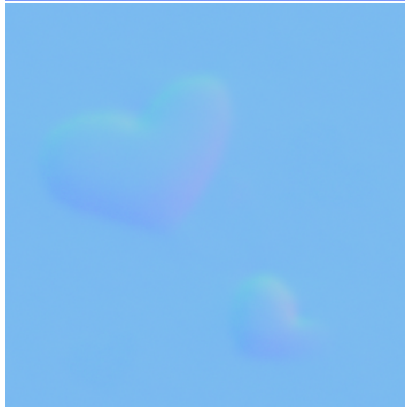
METHAL COASTER



SWORD



PAPER WITH RELIEF



TIN PLATE



Figure 6: Results. Left: input images, Middle: normal maps, Right: renderings.

cialized hardware setups or need a high number of input images, ours only needs a cheap setup, while retaining a similar quality. We have verified the presented work on real world examples, and performed a ground truth evaluation on photo-realistic synthetic data.

In the near future, we would like to look into a technique for merging the normals acquired from both specular and diffuse reflections. Both yield different normal maps for the same scene, but how can they be merged to a single, and more correct, normal map?

Acknowledgements

Part of the research at EDM is funded by the ERDF (European Regional Development Fund) and the Flemish government. Furthermore we would like to thank our colleagues for their help and inspiration. Also special thanks goes to Charles Han to provide us with tools for analyzing our normal maps.

References

- ADATO, Y., VASILYEV, Y., BEN-SHAHAR, O., AND ZICKLER, T. 2007. Toward a theory of shape from specular flow. In *Proceedings of ICCV*.
- BOUGUET, J.-Y., 2006. Camera Calibration Toolbox for MATLAB.
- CHEN, T., GOESELE, M., AND SEIDEL, H.-P. 2006. Mesostructure from specularity. In *Proceedings of CVPR*, vol. 2, 1825–1832.
- FRANCKEN, Y., HERMANS, C., AND BEKAERT, P. 2007. Screen-camera calibration using a spherical mirror. In *Proceedings of CRV*, 11–20.
- FRANCKEN, Y., CUYPERS, T., MERTENS, T., GIELIS, J., AND BEKAERT, P. 2008. High quality mesostructure acquisition using specularities. In *Conference on Computer Vision and Pattern Recognition (CVPR 2008)*, IEEE Computer Society, (to appear).
- FRANCKEN, Y., HERMANS, C., CUYPERS, T., AND BEKAERT, P. 2008. Fast normal map acquisition using an lcd screen emitting gradient patterns. In *Proceedings of CRV*, IEEE Computer Society, (to appear).
- FRANKOT, R. T., AND CHELLAPPA, R. 1988. A method for enforcing integrability in shape from shading algorithms. *PAMI* 10, 4, 439–451.
- FROM SPECULARITY CONSISTENCY, D. D. R. 2008. High quality mesostructure acquisition using specularities. In *Conference on Computer Vision and Pattern Recognition (CVPR 2008)*, IEEE Computer Society, (to appear).
- GOLDMAN, D. B., CURLESS, B., HERTZMANN, A., AND SEITZ, S. M. 2005. Shape and spatially-varying brdfs from photometric stereo. In *Proceedings of ICCV*, vol. 1, 341–348.
- HAN, J. Y., AND PERLIN, K. 2003. Measuring bidirectional texture reflectance with a kaleidoscope. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, 741–748.
- HEALEY, G., AND BINFORD, T. O. 1988. Local shape from specularity. *Computer Vision, Graphics, and Image Processing* 42, 1, 62–86.
- HERTZMANN, A., AND SEITZ, S. M. 2003. Shape and materials by example: A photometric stereo approach. In *Proceedings of CVPR*, IEEE Computer Society, Los Alamitos, CA, USA, vol. 1, 533.
- HERTZMANN, A. 2005. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *PAMI* 27, 8, 1254–1264. Member-Steven M. Seitz.
- HORN, B. 1975. Obtaining shape from shading information. In *The Psychology of Computer Vision*, 115–155.
- IZUKA, K. 2002. *Elements of Photonics*. Wiley-Interscience.
- IKEUCHI, K. 1981. Determining surface orientation of specular surfaces by using the photometric stereo method. *PAMI* 3 (Nov.).
- KUTULAKOS, K. N., AND STEGER, E. 2005. A theory of refractive and specular 3d shape by light-path triangulation. In *Proceedings of ICCV*, 1448–1455.
- MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of EGSR*.
- MALLICK, S. P., ZICKLER, T. E., KRIEGMAN, D. J., AND BELHUMEUR, P. N. 2005. Beyond lambert: Reconstructing specular surfaces using color. In *Proceedings of CVPR*, vol. 2, 619–626.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 519–528.
- MORRIS, N. J. W., AND KUTULAKOS, K. N. 2007. Reconstructing the surface of inhomogeneous transparent scenes by scatter-trace photography. In *Proceedings of ICCV*.
- NAYAR, S. K., FANG, X.-S., AND BOULT, T. 1997. Separation of reflection components using color and polarization. *IJCV* 21, 3, 163–186.
- PATERSON, J. A., CLAUS, D., AND FITZGIBBON, A. W. 2005. Brdf and geometry capture from extended inhomogeneous samples using flash photography. *Computer Graphics Forum (Special Eurographics Issue)* 24, 3 (Sept.), 383–391.
- PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- RUSHMEIER, H., TAUBIN, G., AND GUÉZIEC, A. 1997. *Applying Shape from Lighting Variation to Bump Map Capture*. IBM TJ Watson Research Center.
- SANDERSON, A., WEISS, L., AND NAYAR, S. 1988. Structured Highlight Inspection of Specular Surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10, 1 (Jan), 44–55.
- TARINI, M., LENSCH, H. P. A., GOESELE, M., AND SEIDEL, H.-P. 2005. 3d acquisition of mirroring objects. *Graphical Models* 67, 4 (jul), 233–259.
- UMEYAMA, S., AND GODIN, G. 2004. Separation of diffuse and specular components of surface reflection by use of polarization and statistical analysis of images. *PAMI* 26, 5, 639–647.
- WANG, J., AND DANA, K. J. 2006. Relief texture from specularities. *PAMI* 28, 3, 446–457.
- WOLFF, L. B. 1989. Material classification and separation of reflection components using polarization/radiometric information. In *Proceedings of a workshop on Image understanding workshop*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 232–244.

- WOLFF, L. 1989. Using polarization to separate reflection components. In *Proceedings of CVPR*, 363–369.
- WOODHAM, R. J. 1980. Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1 (jan/feb), 139–144.
- ZHENG, J. Y., AND MURATA, A. 2000. Acquiring a complete 3d model from specular motion under the illumination of circular-shaped light sources. *PAMI* 22, 8, 913–920.
- ZONGKER, D. E., WERNER, D. M., CURLESS, B., AND SALESIN, D. H. 1999. Environment matting and compositing. In *SIGGRAPH*, 205–214.