

Inverse Rendering of Polished Materials under Constant Complex Uncontrolled Illumination

Michael Bang Nielsen and Anders Brodersen

University of Aarhus

Aabogade 34

DK-8200 Aarhus N

Denmark

{bang, rip}@daimi.au.dk

ABSTRACT

Inverse rendering infers realistic descriptions of illumination and material properties from photographs. However, the applicability of state of the art inverse rendering algorithms in real world scenarios is currently limited by the many assumptions made, eg. controlled or simple illumination. In this paper we concentrate on estimating polished materials and describe an inverse rendering system that works under constant complex uncontrolled distant illumination. In particular, we develop an analytical dual angular frequency space shading model for polished materials. We use this shading model in algorithms for estimating homogeneous materials, textures, illumination and both homogeneous materials and illumination simultaneously to the correct scale. We verify the algorithms experimentally in two indoor spectrally different real world scenarios. Furthermore, some of the practical challenges and problems inherent to inverse rendering under complex uncontrolled illumination are discussed.

Keywords

Inverse Rendering, Spherical Harmonics, Environment Maps

1. INTRODUCTION

Synthesizing images indistinguishable from photographs and realistically mixing computer generated and real world imagery has found application in many areas such as movies, entertainment and architectural design. However, it requires highly detailed and realistic input descriptions of materials and illumination.

Inverse rendering employs (semi-) automatic algorithms to measure real world material properties and illumination directly from photographs, thereby enabling these parameters to be changed independently. Using real world measurements of such data as input to physically based rendering algorithms adds further realism to computer generated images.

There has been a lot of research in the field of inverse rendering and many have reported impressive results. However, the applicability of the approaches in gen-

eral real world situations are limited, as they most often require point source illumination [1][2][7][10][19][24][26] or that the illumination can be parameterized in some way [8][17][25]. Thus experiments are most frequently conducted in laboratory or outdoor clear sky environments.

Recently Ramamoorthi and Hanrahan[13] developed a general signal processing framework for inverse rendering. They reported very promising algorithms towards solving the inverse rendering problem under arbitrary complex illumination for materials that can be described by the Torrance-Sparrow microfacet BRDF (Bidirectional Reflectance Distribution Function) model [23]. However, their algorithms are only designed for and tested under carefully constructed semi-complex illumination, see figure 1. As we will argue later, applying their algorithms under constant complex uncontrolled illumination requires additional work.

In this paper, we address inverse rendering under constant complex uncontrolled distant illumination, thus eliminating the assumption of simple or controlled illumination explicitly or implicitly made by other algorithms. We do not consider extending the algorithms suggested by Ramamoorthi and Hanrahan for the Torrance-Sparrow microfacet BRDF to work under more general illumination conditions.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Journal of WSCG, Vol.12, No.1-3, ISSN 1213-6972
WSCG'2004, February 2-6, 2003, Plzen, Czech Republic.
Copyright UNION Agency – Science Press




Figure 1 The image to the left shows the carefully constructed semi-complex illumination used by Ramamoorthi and Hanrahan[12]. The two other images show the uncontrolled complex illumination used by us.

Instead we focus on polished materials that consist of a diffuse and a perfectly reflecting specular contribution. Many materials in the real world adhere to this type, e.g. plastics, polished marble or wood, building facades etc.

We use the parametric coupled BRDF model introduced by Shirley et al.[21] that consists of a non-lambertian diffuse term and a perfectly reflecting specular term modulated by Schlick's approximation to the Fresnel reflectance[20]. The model couples the diffuse and the specular terms in order to ensure energy conservation. Using the spherical harmonics orthonormal system [5][6][9] we derive an analytical dual angular frequency space shading model for polished materials, analogous to the derivations carried out by Ramamoorthi and Hanrahan[13][15][16] for other types of materials, and use it to quickly predict the reflected light field during material estimation. Our inverse rendering algorithms are similar to those of Ramamoorthi and Hanrahan tailored for the shading model for polished materials. Furthermore, we test the algorithms under constant complex uncontrolled illumination in two indoor spectrally different real world scenarios exhibiting complex illumination patterns. The illumination in these scenarios can be seen in figure 1. In particular we demonstrate recovering homogeneous BRDFs, textures, illumination and both homogeneous BRDFs and illumination simultaneously to the correct scale. All test scenes consist of a single sphere thus allowing us to ignore shadowing and inter-reflections.

The main contributions of this paper are

- Derivation of an analytical dual angular frequency space shading model for polished materials.
- Application of inverse rendering algorithms to objects with polished materials under constant complex uncontrolled illumination conditions. To our knowledge this has not previously been successfully demonstrated.

In addition we verify that multispectral inverse rendering must be employed in order to successfully render recovered materials under spectrally different

illumination, and that recovery under white light produces the best results. We also describe our test setup in detail and identify and discuss a number of problems inherent to inverse rendering under general illumination conditions.

2. PREVIOUS WORK

Many algorithms are designed for point source illumination in a laboratory environment as this simplifies BRDF measurements. Ward[24] and Marschner et al.[10] both use a single point light source and measure BRDFs of homogeneous opaque surfaces in a laboratory. Sato et al.[19] also estimate textures. Lensch et al.[7] measure spatially varying BRDFs of complex objects and Debevec et al.[2] recover the reflectance field of a human face. Typically the above methods produce very convincing results for a wide range of materials, but require a fairly large number of input images and only work in laboratory settings. Thus they are not applicable in real world scenarios.

The method of Yu et al. [26] demonstrates estimation of parametric BRDFs and diffuse textures. Contrary to the methods above it can handle inter-reflections, but still it requires that the only direct illumination stems from a number of point light sources. Boivin et al.[1] propose a method to estimate parametric BRDFs and texture from a single image and known point source illumination. It can account for global illumination effects but is not capable of accurately recovering textured specular surfaces.

Some existing methods work outside of a laboratory which means that they must account for an increase in the complexity of the illumination, however, restrictions are still imposed.

Yu and Malik[25] recover albedo maps and BRDFs of a building. However, their algorithm is specifically tailored for a single isolated building, cannot handle highly specular surfaces and depends on the accuracy of the parametric sky radiance model used. Similarly Sato et al.[17] and Love[8] estimate BRDFs under clear sky illumination by assuming a parametric model of sky and sun radiance. Gibson et al.[4] try to overcome constraints of previous algorithms by al-

lowing incomplete descriptions of the illumination. However, ambiguities are introduced. Sato et al.[18] demonstrate recovering both a coarse discretization of the incoming complex illumination and the BRDF of planar surfaces, but still require shadows to be present. Nishino et al.[12] estimate spatially varying reflectance parameters and the illumination distribution from a small shiny object. However, the complexity of the illumination used for testing is limited.

Ramamoorthi and Hanrahan[13] deal with complex illumination, and their practical algorithms give very impressive results. However, the algorithms are only tested in a highly controlled environment under carefully constructed semi-complex illumination.

3. SHADING MODEL FOR POLISHED MATERIALS

When rendering a scene, the reflection equation must be evaluated at least at every visible surface point in order to determine the color of each pixel in the image. Inverse rendering takes as input one or several images and attempts to invert the reflection equation to recover the (possibly spatially varying) BRDF and/or the illumination. In this paper we perform this inversion numerically which requires us to evaluate the reflection equation a large number of times to predict the reflected light field based on the evolving estimates of the BRDF and the illumination. In order to make inverse rendering feasible for polished materials, we derive an analytical dual angular frequency space shading model from the coupled BRDF model[21] and the reflection equation. By representing part of the integral in frequency space wrt. the spherical harmonics orthonormal system[5][6][9], this shading model allows for quick evaluation and in fact reduces the evaluation time by several orders of magnitude compared to brute force integration in the angular domain, see eg.[14]. The derivations are not trivial, but are far too long to be presented in detail here. The full derivations can be found in [11].

We allow complex uncontrolled illumination, but require it to be distant and remain constant during data acquisition. Furthermore we assume that the scene consists of a single, stationary and convex object which allows us to ignore shadowing and inter-reflections. The above implies that the incoming illumination can be represented in 2D, eg. as an environment map, and that the reflected light field at a surface point depends only on the material, surface normal and the viewing direction.

The coupled model[21] for polished materials consists of a non-lambertian diffuse term coupled with a perfectly reflecting specular term multiplied by Schlick's approximation[20] to the Fresnel reflectance. A detailed analysis[11] shows that the specular term contains arbitrarily high frequencies and is best

represented in angular space. It can also be shown analytically[11] that the spherical harmonics coefficients of the coupled diffuse term tend to zero as the order increases and that 99.73% of total energy of the diffuse term is in fact concentrated in the 25 first terms of its spherical harmonic expansion, corresponding to truncation to order $l=4$.

In the formula below we make use of spherical coordinates. Coordinates in the local tangent frame of a surface point are primed and global coordinates are unprimed.

The shading model for polished materials based on the coupled BRDF model becomes

$$L_o(\theta'_o, \phi'_o) = [R_0 + (1 - \cos \theta'_o)^5 (1 - R_0)] L_i(\theta_r, \phi_r) + k R_d [1 - (1 - \cos \theta'_o)^5] \sum_{l=0}^4 \sum_{m=-l}^l h_l L_{lm} \sqrt{4\pi / (2l+1)} Y_{lm}(\alpha, \beta)$$

Equation 1

Where L_o is the outgoing illumination, L_i is the incoming illumination, (θ_o, ϕ_o) is the outgoing/viewing direction, (θ_r, ϕ_r) is the viewing direction reflected about the surface normal, $R_0 \in [0;1]$ is the specular reflectance coefficient at normal incidence, $R_d \in [0;1]$ is the diffuse reflectance coefficient and $k = 21 / (20\pi(1 - R_0))$ is a constant. In our practical experiments we use a different R_d but the same R_0 for each of the red, green and blue color components, corresponding to a dielectric. L_{lm} is the (l,m) 'th spherical harmonic coefficient[9] of the 2D incoming illumination and is easily precomputed if the illumination is known[11]. $Y_{lm}(\alpha, \beta)$ is the (l,m) 'th spherical harmonic[9] and (α, β) is the surface normal.

Finally, h_l is given by an analytical expression which can be found in [11]. The numerical values for orders of l from 0 to 4 are given in table 1.

l	h_l
0	0.8440
1	1.0051
2	0.5308
3	0.0326
4	-0.1271

Table 1 Numerical values of h_l .

In our practical inverse rendering algorithms described later, we use the shading model presented here and derived in detail in [11].

4. EXPERIMENTAL SETUPS

We have tested our algorithms in two spectrally different scenarios. In both scenarios, the scene consists of a single sphere. We have estimated the materials of three different spheres in each of the two scenarios. These spheres, one textured and two homogeneous, are shown in Figure 2.



Figure 2 The three test spheres

We have chosen two indoor scenarios, as this allows us to take the pictures under truly complex and uncontrolled yet constant illumination. In the first scenario we have placed the sphere in the middle of a large office. The complexity of the lighting however remains high as we have direct light coming from more than 20 lights in the ceiling and indirect light reflected off of both specular and diffuse surfaces. This scenario, which we will refer to as the *office scenario*, is depicted in Figure 3 left.



Figure 3 Test setup for the two scenarios

In the second scenario, we have moved the setup into a blue-painted room. Switching on several projectors, including one with a red filter attached, we again have a complex illumination setup including direct white and red light from the projectors and indirect, mostly blue, light from the walls. This scenario, referred to as the *blue room scenario*, is depicted in Figure 3 right.

All imagery is acquired with an off-the-shelf Nikon Coolpix 990 digital camera mounted on a tripod. The camera is calibrated beforehand such that the camera intrinsics are known. To estimate the material of a given sphere we acquire photographs from a number of viewpoints. For the textured sphere, we took pictures from four different viewpoints spaced 90 degrees apart, whereas for the two homogeneous spheres, we only took pictures from three different viewpoints spaced 120 degrees apart. From each viewpoint we take 10 pictures. The first picture is always taken with a checker board pattern visible which is used for camera registration, see Figure 4, using techniques similar to those of Simon et al.[22]. In the remaining nine pictures the checker board pat-

tern is hidden by two pieces of black cardboard to minimize the amount of non-distant illumination. The nine pictures are acquired with different shutter speeds and assembled into an HDR image.



Figure 4 Test setup with and without calibration pattern visible

For the algorithms where the incoming illumination must be known, we acquire photographs of a mirrored sphere from three viewpoints and assemble these into a light probe image[3]. Contrary to the approach by Debevec, we cannot in this setup make the assumption that all imagery is acquired with a near-orthographic camera. Therefore we have developed a new algorithm that does not make this approximation[11]. The omni directional light probe images recovered in the two scenarios are shown in Figure 1. Finally the size and position of the sphere is found by intersecting the cones defined by the camera positions and the conic section that the sphere maps to in each of the images.

5. ESTIMATING MATERIALS

We will now show how the presented theory can be used to construct practical inverse rendering algorithms. The algorithms presented here and in sections 6 and 7 are all similar to the algorithms of Ramamoorthi and Hanrahan[13], except that they are adapted to the shading model for polished materials as described.

The inputs to our algorithms are a geometrical model, i.e. the position and radius of the sphere, a number of HDR images for which the intrinsic and extrinsic camera parameters are known and, in the algorithms where the incoming illumination is assumed known, an omni directional environment map is also given.

In order to estimate materials and illumination, we need to generate a number of samples for which the surface normal, viewing direction and observed RGB radiance is known. These samples are subsequently used in the estimation process. For homogeneous polished materials the samples are either points distributed uniformly on the sphere and then projected into each HDR image, or taken as each pixel in a number of HDR images that back-project to the sphere. For textured materials samples are generated over the sphere geometry according to the texture parameterization and projected into each HDR image. We use a parameterization by spherical coordinates.

Estimating Homogeneous Materials

Recall the shading model for polished materials from Equation 1. For notational convenience in the algorithm to be explained we will rewrite it as

$$L_0(\theta_o, \phi_o) = S(\theta_r, R_0)L_i(\theta_r, \phi_r) + R_d k(R_0)D(\alpha, \beta, \theta_o, L_{im})$$

For each sample extracted, we can set up an equation of this form where the unknowns are the specular reflectance at normal incidence, R_0 , and the diffuse reflectance parameter, R_d .

The algorithm is based on the observation that the equation above is nonlinear only in the R_0 parameter. Given R_0 , the equation simplifies to the linear problem of estimating R_d . In practice our measurements are influenced by noise and the parametric BRDF model does not describe the surface material perfectly hence we use a larger number of samples and solve the resulting system of nonlinear equations in a least squares manner (For our tests, we used the Matlab function *lsqlin* and *lsqnonlin* for solving linear and nonlinear systems of equations).

The algorithm is based upon nested procedures. In the outer procedure a nonlinear least squares method is used to estimate the parameter R_0 , which is then passed on to the inner procedure. Given R_0 , the inner procedure then finds the estimate for R_d , constrained to lie in the interval $[0;1]$, as described above. A residual, computed from the difference between the observed radiance and the radiance predicted using the estimated parameters, is returned to guide the search for the optimal R_0 .

Estimating Textured Materials

Extending the algorithm for homogeneous polished materials described above to estimate a diffuse texture is done by allowing R_d to vary with surface position. Only the inner procedure that solves the linear problem of determining the diffuse texture requires a few additions.

The new inner procedure iterates over all samples with the same texture coordinate, for each possible texture coordinate. For each texel, the diffuse reflectance parameter is found as the clamped weighted average of all the diffuse reflectance parameters predicted by the samples with that particular texture coordinate, clamped to lie in the interval $[0;1]$. The weights of the samples are computed as the ratio of the diffuse component to the specular component. Furthermore, larger weights are given to samples that observe the surface close to the normal direction, as the foreshortening is minimal in this configuration and the sample thus corresponds to a smaller surface area and is less influenced by noise. This weighting scheme favors well lit (large diffuse component) samples with small specular components observed

from angles close to the normal direction. For poorly lit samples and samples with a large specular component, the relative noise in the measurement grows and completely obliterates the diffuse component. It is therefore not possible to estimate the diffuse reflectance parameter reliably from such samples, which is the motivation for using this weighting scheme.

Results

Table 2 shows the BRDF parameters recovered using the algorithms described above. The diffuse textures recovered from the textured sphere in the two scenarios are shown in Figure 5. The spheres were not oriented identically in the two test scenarios and the texture recovered in the office scenario has been warped to obtain the same orientation as the other texture. Note that parts of the texture could not be estimated either because they were not visible at all (different parts for the two scenarios) or they were covered by a specular highlight, making the diffuse contribution too small to be estimated correctly.




Figure 5 Diffuse textures recovered from the textured sphere. a) Blue room scenario b) Office scenario.

To evaluate the results, we have rendered a set of spheres with the estimated materials, lit by the illumination recovered using a light probe. These rendered images along with the photographs of the real spheres from the same viewpoints are shown in Figure 6. The real images are shown in the top row and were not used as part of the estimation process. This shows that the estimated materials can be used to successfully predict the appearance of the spheres from novel viewpoints.

Scenario	$R_{d,red}$	$R_{d,green}$	$R_{d,blue}$	R_0
Red sphere, office	0.9765	0.0186	0.0000	0.0452
Red sphere, blue room	1.0000	0.1593	0.0121	0.0463
White sphere, office	0.6242	0.6210	0.4083	0.0385
Textured sphere, office	-	-	-	0.0504
Textured sphere, blue room	-	-	-	0.0480

Table 2 Recovered parameters of the BRDF for polished materials

One of the important benefits of inverse rendering is the ability to use recovered materials in novel illumi-

nation. To see how well the recovered materials transfer to a spectrally different illumination setup, we have rendered the spheres using the illumination from the scenario different from the one in which the materials were estimated. These images are shown in Figure 7.

As pointed out by Yu and Malik[25], limitations of the RGB color model only enable the recovery of a wavelength dependent pseudo BRDF (contrary to the true BRDF) from RGB photographs. This results in the color aberrations in the first (red sphere) and fourth (textured sphere) image from the left in figure 7. Parameters recovered in the office under white light more successfully predict the appearance of the spheres under spectrally different novel illumination.




Figure 6 The top row shows original photographs not used in material estimation. The bottom row shows renderings using recovered materials from novel viewpoints in original illumination

One obvious problem with the images in Figure 6 and Figure 7 is the “halo” at the objects’ edges. This is partly due to the coupled BDRF model’s approximation to the Fresnel reflectance. However, note that the black background enhances this, as the color at the edges is (and should be) equal to the color of the background. This is because the object reflects light perfectly at the edges due to the Fresnel factor. The slightly blurred highlights near the edges of the objects are probably due to the low resolution of the environment map used in the rendering (we do not use spherical harmonics for rendering but instead a Monte-Carlo raytracer). Furthermore the material of the textured sphere cannot be described perfectly by the coupled BRDF model which eg. results in the dark diffuse component appearing too bright.




Figure 7 Top row shows original photographs. Bottom row shows renderings of estimated materials in novel illumination

6. ESTIMATING ILLUMINATION

We now consider recovering the illumination from an object of an arbitrary, known polished material. This is a problem of great importance for future practical applications. Explicitly capturing the illumination is unpractical and hard in most real world situations. Recovering it from observations of real world objects already present is much more convenient.

Algorithm

The algorithm first determines a low frequency representation of the illumination in frequency space, then use this to determine the diffuse contribution of the reflected light field and finally subtract the diffuse contribution from the observed reflected light field to determine a high resolution angular space version of the illumination. With this approach, estimating the spherical harmonic coefficients of the low frequency part of the illumination from the reflected light field amounts to solving a system of 25 linear equations in 25 unknowns for each color channel. The derivation of this system of equations is similar to that of Ramamoorthi and Hanrahan [13][16] and we will not present them here.

Results

We have used the above algorithm to estimate the illumination from pictures of the red and the white spheres. However, as the known material properties used as input to this algorithm are estimated as described in the previous section, any errors in those estimates propagate to these tests. We therefore also tested the algorithm on images of the red sphere rendered using a standard Monte-Carlo ray tracer. Figure 8 shows the recovered illumination. The HDR image generation process introduced visible banding artifacts, especially pronounced in the images of the white sphere (see Figure 6), not present in the original photographs. These artifacts are the main reason for the visible bands in the recovered illumination. The recovered illumination for the red rendered sphere is closer to the one recovered using the mirrored sphere. The noise present stems from the Monte-Carlo raytracer used to render the input pictures.




Figure 8 Recovered illumination. Top left: Light probe. Top right White sphere. Bottom left: Red Sphere. Bottom right: Red rendered sphere

7. ESTIMATING ILLUMINATION AND BRDF SIMULTANEOUSLY

Ramamoorthi and Hanrahan[13] demonstrated recovering the BRDF and illumination up to a scale from a simplified Torrance-Sparrow micro-facet BRDF. A consequence of estimating polished materials using the coupled BRDF model is that the BRDF and illumination can be recovered simultaneously up to the correct scale as the BRDF model does not allow a linear scale.

Algorithm

The algorithm consists of nested procedures. In the outer procedure a nonlinear least squares method optimizes the parameters of the BRDF. In the inner procedure the illumination is first estimated based on the current BRDF parameters using the algorithm described in the previous section. Next, the reflected light field predicted by the BRDF parameters and the illumination is computed. The residual between the predicted and observed reflected light field is then computed and returned to the nonlinear least squares solver to guide the search for the optimal BRDF parameters.

Results

We have evaluated the algorithm on the red sphere. The original photographs and the renderings are shown in Figure 9 and the numerical values of the BRDF parameters are displayed in Table 3.

Scenario	$R_{d,red}$	$R_{d,green}$	$R_{d,blue}$	R_0
Red sphere, known illumination	0.9765	0.0186	0.0000	0.0452
Red sphere, unknown illumination	0.9763	0.0414	0.0244	0.0452

Table 3 The recovered BRDF parameters for the red sphere with known and unknown illumination




Figure 9 Images in the top row are original photographs. Images in the bottom are rendered using materials and illumination recovered simultaneously

The numerical values are quite close to the parameters found under known illumination, apart from slightly higher green and blue diffuse reflectance parameters. The synthetic images in the bottom row in Figure 9 are rendered using the BRDF parameters

and illumination recovered simultaneously. The recovered illumination suffers from the same artifacts as the illumination recovered with known BRDF parameters.

8. DISCUSSION AND FUTURE WORK

Ramamoorthi and Hanrahan[13] introduced a general theoretical framework for analyzing inverse rendering problems in frequency space under arbitrary illumination. However, their practical algorithms are specifically designed for and tested only under carefully constructed semi-complex illumination which can easily be separated into low and high frequencies as required by their algorithms. As can be seen from Figure 1, a separation is easy as the illumination consists of area light sources that contain only low frequencies and high frequency point light sources. Furthermore, none of the area and point light sources overlap. It is not considered how to separate eg. an illumination environment consisting of an area light source with high frequency edges or more general illumination conditions such as the ones used by us and depicted in figure 1. An alternative to separating the illumination would be to treat every pixel in the environment map as a point light source, but this would require efficient integration schemes to be investigated in order to make the algorithms practically feasible. For these reasons the algorithms of Ramamoorthi and Hanrahan are not applicable under complex uncontrolled illumination. Thus comparison of our results to an extended version of their algorithms is left as future work.

In inverse rendering it is crucial that the chosen BRDF model can successfully describe the materials being estimated. Our experiments show that although the BRDF model we have chosen fits the material of the three spheres quite well, it is not good enough. This is partially due to the fact that the spheres have local anisotropic features and do not exhibit perfectly specular reflection. It is an interesting question whether parametric BRDF models are at all sufficient to capture the realistic behaviour of real world materials convincingly enough.

As we have seen in this paper, applying inverse rendering under complex uncontrolled illumination means that parts of the surface can be both under- and over-exposed in all input photographs, and for such parts it is difficult to reliably estimate local characteristics such as textures.

Future work will need to consider non-constant illumination such as that occurring in realistic outdoor scenarios and this issue can in fact severely limit the practical applicability of inverse rendering. Furthermore it is not obvious how to acquire sufficient input data in order to reduce ambiguities and still keeping

the amount of input data required within a practically feasible limit.

As we saw in section 5, limitations in using inverse rendering with RGB input photographs only allow us to recover the pseudo-BRDF[25]. To overcome this problem, a multispectral BRDF model needs to be used. The algorithms presented here need to be made more robust to errors in the input data such as misregistration with respect to the geometric models and the illumination. Extending the algorithms to be able to handle non-distant illumination, inter-reflections, shadowing, non-convex objects and a broader class of materials will also improve on this aspect.

9. CONCLUSIONS

In this paper we have applied inverse rendering under constant complex uncontrolled illumination which to our knowledge has not previously been successfully demonstrated. In particular, we have developed an analytical dual angular frequency space shading model for polished materials. We have presented algorithms based on this shading model for estimating polished materials and evaluated them in two spectrally different real world scenarios exhibiting complex illumination patterns. Furthermore we have identified and discussed some of the problems inherent to inverse rendering, particularly under complex uncontrolled illumination.

10. ACKNOWLEDGMENTS

This work has been partially funded by the EU IST project WorkSPACE [www.daimi.au.dk/workspace/] We would like to thank professor Kaj Grønbak and Peter Ørbæk for a lot of indispensable help throughout this entire project. In addition we thank Lars Bo Kristensen, Jørgen Lindskov Knudsen, Ole Østerby, Michael Christensen, Niels Olof Bouvin and Christina Nielsen for helping out where ever needed.

11. REFERENCES

- [1] Boivin, S. and Galalowicz, A. Image-based rendering of diffuse, specular and glossy surfaces from a single image. SIGGRAPH 2001.
- [2] Debevec, P. et al. Acquiring the reflectance field of a human face. SIGGRAPH 2000.
- [3] Debevec, P. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. SIGGRAPH 98.
- [4] Gibson, S. et al. Flexible image-based photometric reconstruction using virtual light sources. Computer Graphics Forum, (3):203-214, 2001. ISSN 1067-7055. 3
- [5] Hobson, E.W. The theory of spherical and ellipsoidal harmonics. Cambridge University Press, 1955.
- [6] Inui, T., Tanabe, Y., and Onodera, Y.. Group theory and its applications in physics. Springer Verlag, 1990.
- [7] Lensch, H. et al. Image-Based reconstruction of spatially varying materials. Rendering Techniques 2001.
- [8] Love, R.C.. Surface Reflection Model Estimation from Naturally Illuminated Image Sequences. PhD thesis, Leeds, 1997.
- [9] Mathworld, www.mathworld.com
- [10] Marschner, S.R., Westin, S.H., Lafortune, E.P.F., Torrance, K.E., Greenberg, D.P. Image-based brdf measurement including human skin. Eurographics Rendering Workshop 1999.
- [11] Nielsen, M.B. and Brodersen, A. Inverse rendering under uncontrolled illumination. Master's thesis, University of Aarhus, Denmark, 2002. http://www.daimi.au.dk/~bang/masters_thesis.
- [12] Nishino, K., Zhang, Z. and Ikeuchi K. Determining Reflectance Parameters and Illumination Distribution from a Sparse Set of Images for View-dependent Image Synthesis. IEEE ICCV '01.
- [13] Ramamoorthi, R. and Hanrahan, P. A signalprocessing framework for inverse rendering. SIGGRAPH 2001.
- [14] Ramamoorthi, R. and Hanrahan, P. An Efficient Representation for Irradiance Environment Maps. SIGGRAPH 2001.
- [15] Ramamoorthi, R. and Hanrahan, P. On the relationship between radiance and irradiance: determining the illumination from images of a convex lambertian object. Journal of the Optical Society of America, 2001.
- [16] Ramamoorthi, R. Practical algorithms for inverse rendering under complex illumination. SIGGRAPH 2002 Course Notes. Course "Acquiring Material Models using Inverse Rendering".
- [17] Sato Y. and Ikeuchi K. Reflectance analysis under solar illumination. Technical Report CMU-CS-94-221, CMU, 1994.
- [18] Sato, I., Sato, Y. and Ikeuchi, K. Illumination distribution from brightness in shadows: adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. IEEE ICCV'99.
- [19] Sato, Y., Wheeler, M.D. and Ikeuchi, K. Object shape and reflectance modeling from observation. SIGGRAPH 97.
- [20] Schlick, C. A customizable reflectance model for everyday rendering. Eurographics1993.
- [21] Shirley, P., Hu, H., Smits, B., and Lafortune, E. A practitioners' assessment of light reflection models, 1997. Pacific Graphics '97.
- [22] Simon, G., Fitzgibbon, A.W. and Zisserman, A. Markerless Tracking using Planar Structures in the Scene. In Proc. International Symposium on Augmented Reality, October 2000.
- [23] Torrance, K.E. and Sparrow, E.M.. Theory for off-specular reflection from roughened surfaces. In JOSA, volume 57(9), pages 1105-1114, 1967.
- [24] Ward Larson, Gregory J. Measuring and modeling anisotropic reflection. SIGGRAPH 92.
- [25] Yu, Y. and Malik, J. Recovering photometric properties of architectural scenes from photographs. SIGGRAPH 98.
- [26] Yu, Y., Debevec, P., Malik, J. and Hawkins, T. Inverse global illumination: Recovering reflectance models of real scenes from photographs. SIGGRAPH 99