

A Planar Light Probe

Neil G. Alldrin

David J. Kriegman

nalldrin@cs.ucsd.edu kriegman@cs.ucsd.edu

University of California, San Diego

9500 Gilman Dr. Dept. #0404, La Jolla, CA 92093

Abstract

We develop a novel technique for measuring lighting that exploits the interaction of light with a set of custom BRDFs. This enables the construction of a planar light probe with certain advantages over existing methods for measuring lighting. To facilitate the construction of our light probe, we derive a new class of bi-directional reflectance functions based on the interaction of light through two planar surfaces separated by a transparent medium. Under certain assumptions and proper selection of the two surfaces, we show how to recover Fourier series coefficients of the incident lighting parameterized over the plane. The results are experimentally validated by imaging a sheet of glass with spatially varying patterns printed on either side.

1. Introduction

Images of a scene depend not only on the viewpoint and scene contact, but also upon the way it is illuminated. Knowledge of the lighting is directly used in computer vision techniques such as shape from shading, shape from shadows, shape from specularities, and photometric stereo and can be indirectly exploited in tasks such as recognition. In 3-D graphics, lighting is specified during rendering, and for photo-realistic augmented or mixed reality, the lighting of rendered objects must match the lighting in the real scene. In general, lighting is a function on the 4-D space of rays or directed line segments within a scene. However, it is common and often effective to treat sources as being far from the scene elements in which case lighting can be viewed as a positive function on a sphere. Unless otherwise stated, we will consider the later situation here. The dynamic range of the lighting in a single scene can be very large between direct illumination from say the sun vs. indirect illumination from say a dark shadowed region. Yet the total energy in the darker regions may be significant because only a very small fraction of the light source directions will correspond to a bright, direct illuminant.

The output of an illumination estimation process is some representation of the lighting. Most generally, lighting estimation is viewed as a sampling of the 4-D (or 2-D) lighting space, and a large number of direct measurements are returned as a radiance or environment map [3]. work assumes that there are just a small number of point light sources located nearby or at a distance [17, 27, 8, 12, 26, 24], and so these approaches return the coordinates and strengths of the sources. In general, these methods can be viewed as assuming a parameterized generative model of lighting and attempt to estimate the parameters. Others take a non-parametric approach by representing lighting as a linear superposition of some set of basis functions, and lighting estimation amounts to estimating the coefficients for each basis function [11]. Lighting and reflectance (BRDF) have been well characterized by a spherical harmonic basis [1, 18], and theoretical results [1] and empirical evidence [4, 10] support the idea that a low order expansion of lighting (3rd order) is sufficient for many Lambertian scenes. Haar wavelets have also been used to estimate lighting from cast shadows [15].

A wide variety of techniques have emerged for measuring lighting, and to understand their relative advantage, it is helpful to first consider the different attributes of illumination capture. (1) Does the technique require inserting a physical probe in the scene, require knowledge of the scene geometry, or can it passively infer lighting directly from images of an unknown scene? (2) Does it provide lighting as a 2-D or 4-D function? (3) Does it produce a low or high dynamic range (LDR or HDR) illumination map? (4) Does it require a single image (implying applicability to video with dynamic lighting) or does it use multiple images, perhaps to construct an HDR image? (5) What is the size, bulk, and cost of the probe? (6) What is the resolution (spatial frequency response) of the output? An ideal technique would passively infer lighting from a single image, would provide a high resolution 4-D light field, would produce HDR output, would be applicable to video, and would be low cost.

No technique meets this ideal. Without a probe, the problem is ill-posed and requires some sort of prior to arrive at

a solution [17]; so here we will consider techniques that require insertion of probe into the scene and that treat lighting as a function on a sphere. The most straight forward technique for distant lighting is simply to use a camera to directly measure the light, perhaps using a fish eye lens [5, 6], or as part of a catadioptric omni-cam [14]. Alternatively, a camera can observe a mirrored sphere which has been placed in the scene, that reflects light from all directions (though at a lower resolution toward the occluding contour). While these techniques provide high spatial resolution, they require HDR imaging [3, 13] to fully characterize dynamic range of lighting; this is accomplished by capturing multiple LDR images, and is therefore unsuitable for video. In addition cameras or spherical probes in the scene can be relatively expensive and/or bulky.

A second approach is to introduce a matte light probe into the scene which can be a sphere [27, 26, 24] or an object with known shape [25]. When considered in terms of the spherical harmonic expansion of the BRDF [1, 18], mirrored and Lambertian balls couldn't be more different. The mirrored ball, whose impulse response is akin to a delta function, passes all spatial frequencies whereas the Lambertian ball acts as a low pass filter and severely attenuates high frequencies. It is argued that only an expansion to 3rd order is feasible. However, the advantage of a Lambertian probe is that the dynamic range of the images of a sphere under any lighting condition are low, and so lighting can be estimated from a single image, making the technique suitable for video. [2] image multiple spheres with varying reflectance (matte, glossy, and specular) to recover lighting.

A third way of constructing a probe is to take advantage of non-convex geometry and the resulting shadows [20, 21, 22, 15]. Consider a sundial. The irradiance arriving at a particular point on the underlying surface is the product of the incident lighting with the visibility function induced by the geometry of the sundial. In [15] the relation between lighting and cast shadows are analyzed in the frequency domain in terms of spherical harmonics and Haar wavelet bases. [19] analyzes cast shadows in Fourier space.

In the paper presented here, we introduce a fourth mechanism for measuring lighting by using a multi-layered, transparent medium which differentially absorbs and reflects light. In particular, the homogenous middle layer of the planar probe is a transparent medium (e.g., air or glass), the top layer is patterned to partially absorb and transmit light differently at different locations (e.g., a transparency). The bottom layer is also patterned to reflect light differently at different locations (e.g., a piece of printed paper). The key idea is that the design of the two patterns leads to a spatially varying effective BRDF. That is when image plane irradiance is averaged over an area (within a single pixel or over multiple pixels), the manufactured probe can




Figure 1. The imaging setup for our light probe.

be treated as a meso-scopic geometric structure akin to the micro-facet or pit models used in constructing models of BRDFs [23, 16, 9]. Given a set of these BRDFs distributed over the plane, we can recover lighting over the upper hemisphere by treating each BRDF as a basis function. A special case of the analysis is when the upper and lower patterns are binary (completely opaque or transparent) in which case the effective BRDF is solely the result of shadowing and masking [16].

The advantage of such a light probe is that its capabilities lie between a mirrored and Lambertian probe in that it can measure higher order frequencies than a Lambertian ball, yet high dynamic range imaging is not needed. As consequence, it is suitable for capturing lighting in video. Furthermore, a particular application of this probe is augmented reality where it is common to include a planar geometric probe with fiducial markers (See for example AR Toolkit [7]) for determining relative orientation, and our illumination probe could be readily integrated for lighting estimation and photo-consistent rendering.

In the rest of this paper, we first introduce a design for a probe, and characterize the effective BRDF of the probe as a function f of the patterned upper and lower layer. We then show how the probe is designed, and how lighting can be estimated from the probe. Finally, we report on results of experiments that validate the potential of the probe.

2. Designing a BRDF for Lighting Recovery

Suppose we have a material consisting of two parallel layers, referred to as the top and bottom layers, separated by a transparent medium (see figure 1). Further suppose both layers are spatially varying so that the top layer reflects and transmits light as a function of position and the bottom layer reflects light as a function of position. We assume that the wavelength of light is much smaller than both the thickness of the transparent layer and the spatially varying patterns so that no diffractive effects occur. When light hits the top layer, some is directly reflected and some is trans-




Figure 2. Mapping of the incident lighting onto the plane via refraction.

mitted through the transparent medium where it is reflected by the bottom layer, travels back through the transparent medium, and finally passes out through the top layer. We seek to analyze the reflectance properties of such a material with the ultimate goal of recovering the lighting from a set of distinct BRDFs constructed in this way.

We denote the position on the plane by \vec{x} , the incident radiance from direction $\vec{\omega}_i$ arriving at position \vec{x} as $l_i(\vec{x}, \vec{\omega}_i)$, and the reflected radiance in direction $\vec{\omega}_r$ exiting at position \vec{x} as $l_r(\vec{x}, \vec{\omega}_r)$. We relate the incident radiance from differential solid angle $d\vec{\omega}_i$ at position \vec{x} to the reflected radiance at position \vec{x}' in direction $\vec{\omega}_r$ through the BSSRDF $S(\vec{x}, d\vec{\omega}_i \rightarrow \vec{x}', \vec{\omega}_r)$ so that the reflected radiance at position \vec{x} in direction $\vec{\omega}_r$ is

$$l_r(\vec{x}, \vec{\omega}_r) = \int_{\vec{x}' \in A} \int_{d\vec{\omega}_i \in \Omega} l_i(\vec{x}', d\vec{\omega}_i) S(\vec{x}', d\vec{\omega}_i; \vec{x}, \vec{\omega}_r) d\vec{\omega}_i^N dA \quad (1)$$

where $\vec{\omega}_i^N$ is the projected solid angle onto the plane with surface normal N . This expression is stating that all light arriving in some area A contributes to the reflected radiance at position \vec{x} according to BSSRDF S .

We now begin to specify S . First, we split it into a reflective term and a scattering term,

$$S(\vec{x}, d\vec{\omega}_i \rightarrow \vec{x}', \vec{\omega}_r) = f_r(\vec{x}, d\vec{\omega}_i \rightarrow \vec{\omega}_r) + f_s(\vec{x}, d\vec{\omega}_i \rightarrow \vec{x}', \vec{\omega}_r). \quad (2)$$

The reflective term, $f_r(\vec{x}, d\vec{\omega}_i \rightarrow \vec{\omega}_r)$, is just a standard spatially varying BRDF and represents the light directly reflected off the top surface. The scattering term, $f_s(\vec{x}, d\vec{\omega}_i \rightarrow \vec{x}', \vec{\omega}_r)$ represents light passing through the top layer at position \vec{x}' with solid angle $d\vec{\omega}_i$ and re-emerging at position \vec{x} in direction $\vec{\omega}_r$.

Assuming no absorption as light travels through the transparent medium, specular transmission through the top layer, and that scattering beyond the initial reflection at the




Figure 3. Relationships between various angles and distances.

bottom layer is negligible, we can split the scattering term f_s into the product of three parts,

$$f_s(\vec{x}, d\vec{\omega}_i \rightarrow \vec{x}', \vec{\omega}_r) = f_t(\vec{x}, d\vec{\omega}_i) F_i(\vec{\omega}_i) \dots \cdot g_r(\vec{x} + \vec{\Delta}_i, d\vec{\omega}_{i'}, \vec{\omega}_{r'}) d\vec{\omega}_{i'}^N \dots \cdot F_r(\vec{\omega}_r) f_t(\vec{x}', d\vec{\omega}_r) d\vec{\omega}_r^N \quad (3)$$

where $f_t(\vec{x}, d\vec{\omega}_i)$ represents the initial transmission through the top layer, $g_r(\vec{x} + \vec{\Delta}_i, d\vec{\omega}_{i'}, \vec{\omega}_{r'})$ is the reflection from the bottom layer, $f_t(\vec{x}', d\vec{\omega}_r)$ is the transmission out through the top layer, and $F_i(\vec{\omega}_i)$ and $F_r(\vec{\omega}_r)$ are Fresnel transmission terms for entering and leaving the transparent medium. $\vec{x} + \vec{\Delta}(\vec{\omega}_i)$ is the position where a given ray of light hits the bottom layer.

Now observe that if we fix the exitant angle, then only a single incident angle contributes to the scattering term. To understand this, note that only a single position \vec{x}' contributes to the exitant radiance, which is completely specified by $\vec{\omega}_i$, $\vec{\omega}_r$, the thickness of the transparent medium h , and Snell's law (specified with the indices of refraction η_1 and η_2 for the outside medium and transparent medium respectively). We can take this idea one step further by parameterizing the incident and exitant hemisphere in terms of $\vec{\Delta}_i$ and $\vec{\Delta}_r$ respectively, which denote the distance along the plane between where a ray intersects the top layer and bottom layer. We can now specify the scattering term as a function of \vec{x} , $\vec{\Delta}_i$, and $\vec{\Delta}_r$,

$$f_s(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) = f_t(\vec{x}, \vec{\Delta}_i) F_i(\vec{\Delta}_i) \dots \cdot g_r(\vec{x} + \vec{\Delta}_i, \vec{\Delta}_i, \vec{\Delta}_r) d\vec{\omega}_{i'}^N \dots \cdot F_r(\vec{\Delta}_r) f_t(\vec{x} + \vec{\Delta}_r + \vec{\Delta}_i, \vec{\Delta}_r) d\vec{\omega}_r^N. \quad (4)$$

If the incident lighting is distant, then it remains constant across the plane and we can parameterize it entirely in terms

of $\vec{\Delta}_i$. Furthermore, if the refractive index of the transparent layer is higher than that of the outside environment, then because of Snell's law, $|\vec{\Delta}_i|$ will be bounded by some finite value dictated by the critical angle and the thickness of the medium. Without loss of generality, we scale the coordinate system along the plane so that $|\vec{\Delta}_i|_{\max} = \frac{1}{2}$. To simplify subsequent integrals, we define the incident lighting so that $l(\vec{\Delta}) = 0$ for all $|\vec{\Delta}| > \frac{1}{2}$ (where it was previously undefined).

Putting everything together, we can rewrite equation 1 as,

$$l_r(\vec{x}, \vec{\Delta}_r) = \int_{\vec{\Delta}_i \in A} l_i(\vec{\Delta}_i) S(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) \frac{d\vec{\omega}_i^N}{dA} dA \quad (5)$$

where $A = [-\frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$.

2.1. Fourier Analysis

In the previous section we formulated the exitant radiance as the integral of the incident lighting $l_i(\vec{\Delta}_i)$ and simplified BSSRDF $S(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r)$. We now turn our attention to the *average* behavior of the system. To simplify the math, we fold the projected solid angle terms and the Fresnel transmission terms into modified versions of f_r , f_t , and g_r so that,

$$\tilde{f}_r(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) = f_r(\vec{x}, \vec{\Delta}, \vec{\omega}_r) \frac{d\vec{\omega}_i^N}{dA} \quad (6)$$

$$\tilde{f}_t(\vec{x}, \vec{\Delta}) = f_t(\vec{x}, \vec{\Delta}) \quad (7)$$

$$\tilde{g}_r(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) = g_r(\dots) F_i(\vec{\Delta}_i) F_r(\vec{\Delta}_r) \frac{d\vec{\omega}_i^N}{dA}. \quad (8)$$

We can then write equation 5 as

$$l_r(\vec{x}, \vec{\Delta}_r) = \int_{\vec{\Delta}_i \in A} l_i(\vec{\Delta}_i) \tilde{S}(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) dA \quad (9)$$

where

$$\begin{aligned} \tilde{S}(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) &= \tilde{f}_r(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) + \dots \\ &\tilde{f}_t(\vec{x}, \vec{\Delta}_i) \cdot \tilde{g}_r(\vec{x} + \vec{\Delta}_i, \vec{\Delta}_i, \vec{\Delta}_r) \dots \\ &\cdot \tilde{f}_t(\vec{x} + \vec{\Delta}_i + \vec{\Delta}_r, \vec{\Delta}_r). \end{aligned} \quad (10)$$

If \tilde{S} varies spatially with period 1 in the x and y directions, then the average exitant radiance is

$$\bar{l}_r(\vec{\Delta}_r) = \int_{\vec{x} \in A} \int_{\vec{\Delta}_i \in A} l_i(\vec{\Delta}_i) \tilde{S}(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) d\vec{\Delta}_i d\vec{x} \quad (11)$$

where $A = [-\frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$.

2.1.1 An Ideal Case

Suppose we are free to choose any form for \tilde{f}_r , \tilde{f}_t , and \tilde{g}_r so long as it varies in x and y with period 1. To simplify things,

let $\tilde{f}_r = 0$, $\vec{\Delta}_r = \vec{0}$ and suppose $\tilde{g}_r(\vec{x}, \vec{\Delta}_i, \vec{\Delta}_r) = \tilde{g}_r(\vec{x}, \vec{\Delta}_r)$ and $\tilde{f}_t(\vec{x}, \vec{\Delta}_i) = \tilde{f}_t(\vec{x})$ do not depend on $\vec{\Delta}_i$. Then the average exitant radiance is

$$\begin{aligned} l_r(\vec{0}) &= \int_{\vec{x} \in A} \tilde{f}_t(\vec{x}) \tilde{g}_r(\vec{x}, \vec{0}) \\ &\cdot \int_{\vec{\Delta}_i \in A} l_i(\vec{\Delta}_i) \tilde{f}_t(\vec{x} - \vec{\Delta}_i) d\vec{\Delta}_i d\vec{x} \end{aligned} \quad (12)$$

(13)

Recalling our goal of recovering l_r , if we choose $\tilde{f}_t(\vec{x}) = \delta(\vec{x}) + \delta(\vec{x} - \vec{x}')$ and $\tilde{g}_r(\vec{x}) = \delta(\vec{x} - \vec{x}')$, where δ is the Kronecker delta function, then we get

$$l_r(\vec{0}) = l_i(\vec{x}') + l_i(\vec{0}). \quad (14)$$

Thus, setting \tilde{f}_t and \tilde{g}_r to delta functions enables point sampling of the incident radiance. Assuming we could measure the effects of a delta function, this would allow full recovery of the lighting (or a sampled version if \vec{x}' is restricted to a finite set of values). However, this is an unrealistic solution in practice because delta-like functions would result in very subtle changes in image intensity which would be hard to recover with an image sensor.

Another possible choice for \tilde{f}_t is

$$\begin{aligned} \tilde{f}_t(\vec{x}, \vec{\Delta}_i) &= e^{-i2\pi\vec{n}\cdot\vec{x}} \\ &= e^{-i2\pi ux - i2\pi vy} \end{aligned} \quad (15)$$

where $\vec{n} = \{u, v\}$ and u, v are integers. In this case the average exitant radiance is

$$\begin{aligned} l_r(\vec{0}) &= \int_{\vec{x} \in A} e^{-i2\pi\vec{n}\cdot\vec{x}} \tilde{g}_r(\vec{x}, \vec{0}) \\ &\cdot \int_{\vec{\Delta}_i \in A} l_i(\vec{\Delta}_i) e^{-i2\pi\vec{n}\cdot(\vec{x} - \vec{\Delta}_i)} d\vec{\Delta}_i d\vec{x} \end{aligned} \quad (16)$$

$$\begin{aligned} l_r(\vec{0}) &= \int_{\vec{x} \in A} e^{-i2\pi\vec{n}\cdot\vec{x}} \tilde{g}_r(\vec{x}, \vec{0}) \\ &\cdot \int_{\vec{\Delta}_i \in A} l_i(\vec{\Delta}_i) e^{i2\pi\vec{n}\cdot\vec{\Delta}_i} d\vec{\Delta}_i d\vec{x} \end{aligned} \quad (17)$$

$$l_r(\vec{0}) = \tilde{G}_{2\vec{n}} L_{\vec{n}}^* \quad (18)$$

where $L_{\vec{n}} = L(u, v)$ is the u, v th 2D Fourier series coefficient of the lighting, $\tilde{G}_{2\vec{n}} = \tilde{G}(2u, 2v)$ is the $(2u, 2v)$ th Fourier series coefficient of \tilde{g}_r , and $*$ denotes the conjugate operator. To recover $L_{\vec{n}}$ we simply need to choose \tilde{g}_r so that it contains frequencies of $2\vec{n}$. The most logical choice is to set $\tilde{g}_r(\vec{x}, \vec{0}) = e^{i4\pi\vec{n}\cdot\vec{x}}$, yielding $G_{2\vec{n}} = 1$ and

$$l_r(\vec{0}) = L_{\vec{n}}^*. \quad (19)$$

From this equation we directly obtain $L_{\vec{n}}$. While the assumptions used to reach this result are unrealistic¹, it does

¹Not only is positivity violated, but imaginary numbers are used!

provide hope that one can construct a BRDF that directly outputs frequency components of the lighting. Since low frequency lighting is often sufficient for rendering purposes, we should be able to obtain a useful lighting representation using only a small set of such BRDFs.

2.1.2 A More Realistic Case

To satisfy the laws of physics, we must modify the above formulation in a number of ways:

1. Positivity and conservation of energy must be enforced. A valid BRDF or BTDF is greater than or equal to zero for all inputs and the integral of a BRDF or BTDF over all incident directions must sum to 1 or less.
2. Fresnel reflectance varies with incident angle, so our assumption that the top and bottom layers are constant across incident angles is violated.
3. We must depend exclusively on non-imaginary numbers.

Assumption 1 can be met simply by adding a constant term to \tilde{f}_t or \tilde{g}_r and then scaling the signal with a multiplicative factor. Thus, we will get a new signal $k_m(k_a + f(\vec{x}))$ that satisfies positivity and conservation of energy. Assumption 2 implies that we can no longer factor \tilde{f}_t and \tilde{g}_r out of the inner integral. However, in many cases we can factor these terms into a spatially varying component that doesn't vary with incident angle and a non-spatially varying component that remains inside the integral. If the spatially varying components have appropriate signals we can recover the product of the lighting with the non-spatially varying components of \tilde{f}_t and \tilde{g}_r . Once this is recovered we can divide out the undesired terms and recover the original lighting. Assumption 3 is easily overcome by using sinusoids instead of complex exponentials (ie, the real-valued form of the Fourier series).

3. Experimental Validation

3.1. Setup

To validate our theory, we printed a set of sinusoidal patterns on a transparency sheet and on a sheet of matte paper and separated the two patterns with a sheet of glass (see figure 4). The thickness of the glass was measured to be 0.096 inches and the refractive index assumed to be 1.52. To flatten the transparency sheet we also placed a sheet of glass above the transparency. Thus, our light probe consists of two sheets of glass, a transparency sheet and a piece of matte paper. For each frequency² (u, v) we devote

²Minus redundant frequencies caused by conjugate symmetry.




Figure 4. Our experimental setup. We have our planar light probe next to a mirrored ball, which is used to capture the baseline lighting.




Figure 5. Closeup of the patterns on our planar light probe.

two regions where the top layer is a sinusoid of the form $\frac{1}{2}(1 + \sin(-2\pi(ux + vy)))$ and the bottom is a sinusoid of the form $\frac{1}{2}(1 + \sin(4\pi(ux + vy) + \tau))$. If we assume the bottom layer is Lambertian then it can be shown that the reflected radiance is of the form $aL_0 + bL_{a\vec{n}} + cL_{b\vec{n}} + s(x)$, where L_0 is the average incident radiance, $L_{a\vec{n}}$ and $L_{b\vec{n}}$ are the even and odd portions of the Fourier coefficients, and $s(x)$ represents the specularities that occur at the surface of the glass. We add a spatial dependence on the surface reflection because while we assume the light and camera are distant, in practice this assumption is violated and surface reflections vary spatially across the surface (albeit slowly). To counteract the effect of the spatially varying specular term, we sample the specular reflection by placing unpatterned regions at uniform intervals across the light probe. There are four types of unpatterned region: (top clear, bottom white), (top clear, bottom dark), (top dark, bottom white), (top dark, bottom dark). Because Fresnel reflection occurs at the top surface of the glass, by subtracting the average intensity in a constant region from some other type of region, we effectively cancel the specular reflection term.

We layout the planar light probe in terms of blocks,




Figure 6. Low order frequencies recovered from our light probe vs. the baseline provided by a mirrored ball. (Top) Real and imaginary estimated frequency components compared to the baseline frequencies. (Bottom,Right) Actual lighting, (Bottom,Middle) Actual lighting resulting from an order 1 Fourier series approximation, (Bottom,Left) Estimated lighting.




Figure 7. Low order frequencies for a different environment recovered from our light probe vs. the baseline provided by a mirrored ball. (Top) Real and imaginary estimated frequency components compared to the baseline frequencies. (Bottom,Right) Actual lighting, (Bottom,Middle) Actual lighting resulting from an order 1 Fourier series approximation, (Bottom,Left) Estimated lighting.

where each block consists of a pattern designed to extract a single frequency component of the incident lighting. To calibrate the light probe, suppose we have K blocks and N images of the probe under different known lighting conditions. Then we can form a feature vector $V \in [K \times N]$ of our observations and a corresponding linear system $L = MV$ where L contains the known lighting coefficients. After solving for the matrix M we can then obtain unknown lighting coefficients from images of our light probe using MV .

3.2. Results

Figures 6,7, and 8 show some results from our light probe. For very low frequencies the results seem relatively stable; however, the stability quickly declines for higher frequencies³. There are many potential causes for this. One likely possibility is that high frequency terms are present in many of the regions we observe as a result of discretization in the printing process. Since we are trying to fit only low order terms, if higher frequencies are present it will interfere with the fitting process. A potential solution would be to use a better printing process such as dye-sublimation⁴ to get more continuous half-tones. Another possible workaround would be to directly measure the BRDFs present on the light




Figure 8. Higher frequency approximations. Current estimation techniques are only working well for very low frequencies. Notice all the noise in the estimation; we hope to fix this in future work.

³Note: In these figures, redundant conjugate symmetric coefficients have been omitted.

⁴A 1200dpi laser printer was used in our experiments.

probe and use these directly as basis functions for the lighting. This is the direction we are currently working towards

as it should yield the best results; however it requires tedious and error prone BRDF measurements that can be difficult to get right.

4. Conclusions

We have shown theory that suggests BRDFs can be manufactured that are sensitive to specific frequencies of the lighting. Based on this theory, we have constructed a planar light probe capable of estimating low frequencies of the lighting. Such a probe would be useful for many applications, particularly applications requiring lighting estimation from low-dynamic range images. For most materials, a fifth order frequency approximation of the lighting is enough to photo-realistically render it, so if we can push our probe a little further it will be a truly useful device.

Acknowledgements

This work was supported in part by National Science Foundation grants IIS-0308185 and EIA-0224431.

References

- [1] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(2):218–233, February 2003.
- [2] P. Debevec et al. Estimating surface reflectance properties of a complex scene under captured natural illumination. Technical Report ICT-TR-06.2004, University of Southern California ICT, December 2004.
- [3] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 369–378, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [4] R. Epstein, P. Hallinan, and A. Yuille. 5+/-2 eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *PBMCV*, 1995.
- [5] N. Greene. Environment mapping and other applications of world projections. *IEEE Comput. Graph. Appl.*, 6(11):21–29, 1986.
- [6] P. Haeberli and M. Segal. Texture mapping as a fundamental drawing primitive. In *Fourth Eurographics Workshop on Rendering*, pages 259–266. Eurographics, June 1993.
- [7] H. Kato. Artoolkit.
- [8] C. Kim, A. P. Petrov, H. Choh, Y. Seo, and I. Kweon. Illuminant direction and shape of a bump. *Optical Society of America Journal A*, 15:2341–2350, Sep 1998.
- [9] J. J. Koenderink, A. J. V. Doorn, K. J. Dana, and S. Nayar. Bidirectional reflection distribution function of thoroughly pitted surfaces. *Int. J. Comput. Vision*, 31(2-3):129–144, 1999.
- [10] K. Lee, J. Ho, and D. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Analysis & Machine Intelligence*, pages 684–698, May 2005.
- [11] S. R. Marschner and D. P. Greenberg. Inverse lighting for photography. In *Proceedings of the Fifth Color Imaging Conference, Society for Imaging Science and Technology*, 1997.
- [12] D. Miyazaki, R. T. Tan, K. Hara, and K. Ikeuchi. Polarization-based inverse rendering from a single view. In *ICCV*, pages 982–987, 2003.
- [13] S. Nayar and T. Mitsunaga. High dynamic range imaging: Spatially varying pixel exposures. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, pages 472–479, 2000.
- [14] S. K. Nayar. Catadioptric omnidirectional camera. In *CVPR '97: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97)*, page 482, Washington, DC, USA, 1997. IEEE Computer Society.
- [15] T. Okabe, I. Sato, and Y. Sato. Spherical harmonics vs. haar wavelets: Basis for recovering illumination from cast shadows. In *CVPR (1)*, pages 50–57, 2004.
- [16] M. Oren and S. K. Nayar. Generalization of lambert's reflectance model. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246, New York, NY, USA, 1994. ACM Press.
- [17] A. Pentland. Finding the illuminant direction. *J. Optical Soc. Am.*, 72:448–455, 1982.
- [18] R. Ramamoorthi. *A signal-processing framework for forward and inverse rendering*. PhD thesis, Stanford, 2002. Adviser-Pat Hanrahan.
- [19] R. Ramamoorthi, M. Koudelka, and P. Belhumeur. A fourier theory for cast shadows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(2):288–295, 2005.
- [20] I. Sato, Y. Sato, and K. Ikeuchi. Illumination distribution from brightness in shadows: adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. In *Proceedings of IEEE ICCV'99*, volume 2, pages 875 – 882, September 1999.
- [21] I. Sato, Y. Sato, and K. Ikeuchi. *Illumination distribution from shadows*. Kluwer Academic Publishers, 2001.
- [22] I. Sato, Y. Sato, and K. Ikeuchi. Illumination from shadows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(3):290–300, 2003.
- [23] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from rough surfaces. *Journal of the Optical Society of America*, 57(9):1105–1114, 1967.
- [24] Y. Wang and D. Samaras. Estimation of multiple illuminants from a single image of arbitrary known geometry. In *ECCV '02: Proceedings of the 7th European Conference on Computer Vision-Part III*, pages 272–288. Springer-Verlag, 2002.
- [25] M. Weber and R. Cipolla. A practical method for estimation of point light-sources. In *BMVC*, 2001.
- [26] Y. Zhang and Y.-H. Yang. Multiple illuminant direction detection with application to image synthesis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(8):915–920, 2001.
- [27] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(7):680–702, 1991.