UBC Computer Science Technical Report TR-2001-08

Image-Based Measurement of Light Sources With Correct Filtering

Wolfgang Heidrich[1] and Michael Goesele[2]

[1)] The University of British Columbia
[2)] Max-Planck-Institut für Informatik

Abstract

In this document we explore the theory and potential experimental setups for measuring the near field of a complex luminary. This work extends on near field photometry by taking filtering issues into account. The physical measurement setups described here have not been tested at the time of writing this document, we simply describe several possibilities here. Once actual tests have been performed, the results will be published elsewhere.

1 Introduction

Real-world models for light sources are an important prerequisite for photorealistic image synthesis. Traditionally, the only measured information available for light sources has been the far field (i.e. directional information for the emitted light from a point shaped light source), for which most luminary manufacturers provide data bases. Unfortunately, the far field is only a faithful approximation of the emitted light when the light source is far from the object to be illuminated.

On the other hand, a light field [5, 8] completely represents near and far field illumination of a light source, and can thus be used to represent the light source without knowing its geometry.

In his work on *near field photometry* [1, 2, 3], Ashdown has presented methods for measuring the light field of luminaries in both its near field and its far field components. On the rendering side, Heidrich et al. [7] have described efficient algorithms based on a similar representation they called a *canned light source*. The similarity between the two representations illustrates that measurement and rendering can use the same data structures, thereby allowing for an efficient pipeline from acquisition to image synthesis.

In nearfield photometry a number of pinhole cameras are pointed at the luminary to be measured (actually, typically one camera is moved around), and then the irradiance incident to the film plane is recorded using a CCD chip. The camera positions correspond to a sampling of some virtual sampling surface $\mathcal{S}$ (see Figure 1). In the simplest case, this could be a plane, and the viewing directions at the various camera positions could all be parallel. In this case the two-plane parameterization [5, 8] is obtained, which is also used by Heidrich et al [7].

The use of pinhole cameras in this setting corresponds to taking discrete point samples on the surface $\mathcal{S}$ that surrounds the light source. Unfortunately, the light source may produce arbitrarily high spatial frequencies on $\mathcal{S}$, in which case we have just introduced aliasing by not applying a low-pass filter before the point sampling step.

It should be noted that near field photometry does not *actually* use pinhole cameras, but rather conventional cameras with lens systems. These behave *almost* like pinhole cameras, except that they integrate incident radiance over the finite aperture of the lens. The issue of aliasing has been raised by Halle [6] and was noted by Levoy and Hanrahan [8]. They both propose to use the finite aperture of the camera lens as a low-pass filter by choosing the aperture size equal to the size of a sample on the camera plane. Thus, some amount of low-pass filtering is in fact applied, but the shape of the filter kernel is more or less random (i.e. it depends on the design of the lens system rather than considerations from sampling theory and the task at hand).

In this report, we describe a measurement system that improves on nearfield photometry by projecting the light field emitted by the light source into a finite basis *before* measurement. This is done using a simple optical system. The shape and support of the basis functions are specifically designed for a particular sampling scheme. Based on the resulting measurements we can *exactly* reconstruct the least-squares approximation of the true light field in our basis. Alternatively, we can reconstruct with a more efficient, shift-invariant filter, and obtain a close approximation to the least-squares solution.

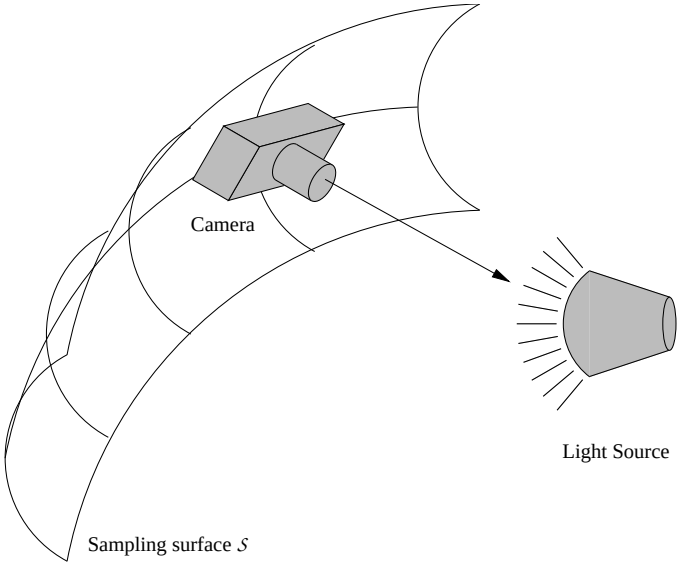

Figure 1: In near field photography the light source is measured by taking images with a pinhole camera from a number of points on a sampling surface $\mathcal{S}$.

The remainder of this report is structured as follows: first, we are going to give a rough overview of the basic idea in Section 2. We then review the theory behind our method in Section 3, before we discuss a number of physical setups in Section 4. We finally conclude with some ideas for future work.

2 Overview of the Proposed Method

The basic idea of the proposed method is to replace point sampling on $\mathcal{S}$ with an area sampling approach. Rather than distributing a number of pinhole cameras on $\mathcal{S}$, we cover $\mathcal{S}$ with optical *filters* that optically project the emitted light onto a second surface $\mathcal{M}$, the *measurement surface*. We assume that we can measure the irradiance arriving at $\mathcal{M}$ at a very high resolution. Thus, $\mathcal{M}$ corresponds to the film or CCD chip in a conventional camera.

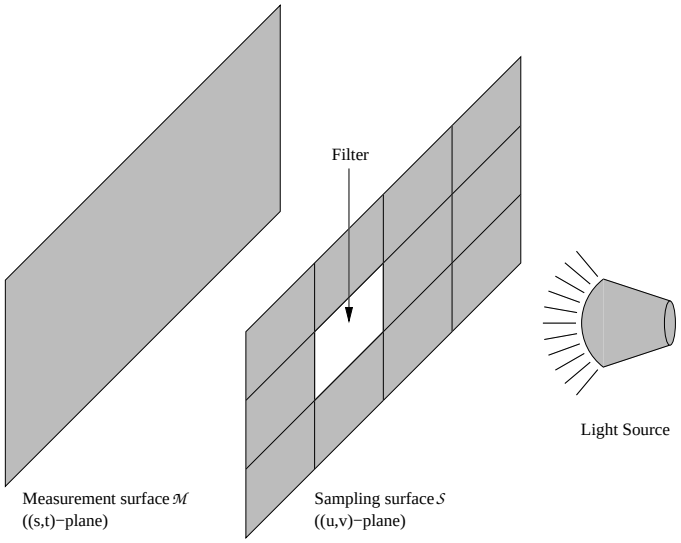

Figure 2: Conceptual setup for the proposed measurement method.

Symbol	Meaning
$\Psi_{ijkl}(u,v,s,t)$	basis function for approximating the light field
Φ_i	1D basis used for reconstruction
$\Phi_{ijkl}(u,v,s,t)$	4D tensor product basis for reconstruction
Φ'_i	biorthogonal 1D basis used for measurement
$\Phi'_{ijkl}(u,v,s,t)$	4D tensor product basis for measurement
$\mathcal{M}$	surface on which the irradiance is measured (*measurement surface*, or (s,t)-plane)
$\mathcal{S}$	surface on which the optical filters are placed (*sampling surface*, or (u,v)-plane)
$L(u,v,s,t)$	radiance passing through (u,v) on $\mathcal{S}$ and (s,t) on $\mathcal{M}$
$\tilde{L}(u,v,s,t)$	projection of $L(u,v,s,t)$ into basis $\{\Psi_{ijkl}(u,v,s,t)\}$
$L_{mn}(u,v,s,t)$	radiance $\Phi'_{mn}(u,v) \cdot L(u,v,s,t)$ projected through *one* filter $\Phi'_{mn}(u,v)$
$E_{mn}(s,t)$	irradiance caused by $L_{mn}(u,v,s,t)$ on the measurement surface $\mathcal{M}$
$E'_{mn}(s,t)$	approximation of $E_{mn}(s,t)$ with simplified geometric term

Table 1: Notation used throughout this document.

Mathematically, the optical filters project the light field emitted by the light source into some function basis. The filters are designed in such a way that match the desired reconstruction filters. I.e. if a bilinear reconstruction filter on $\mathcal{S}$ is desired, then the optical filter is designed such that it projects the light field into the space of piecewise bilinear functions on $\mathcal{S}$.

For the sake of simplicity, we will in our initial discussion of the theory assume that $\mathcal{M}$ and $\mathcal{S}$ are planes, where we parameterize $\mathcal{S}$ via the parameters (u,v), and $\mathcal{M}$ via the parameters (s,t). This corresponds to the standard two-plane parameterization of light fields. The underlying assumption of our method is that the true light field $L(u,v,s,t)$ is well represented by its projection $\tilde{L}(u,v,s,t)$ into the function space spanned by the optical filters.

Our measurement method also applies to other geometries and parameterizations, and this issue will be discussed when we describe various physical measurement setups in Section 4.

3 Theory

Before we discuss the theory behind the proposed method in detail, we first introduce the mathematical notation used throughout this document.

For the measurement, we assume that the light field emitted by the light source is well represented by a projection into a basis $\{\Psi_{ijkl}(u,v,s,t)\}_{ijkl \in \mathbf{Z}}$:

$$L(u,v,s,t) \approx \tilde{L}(u,v,s,t) := \sum_i \sum_j \sum_k \sum_l \Psi_{ijkl}(u,v,s,t) \cdot L_{ijkl}. \tag{1}$$

We assume that Ψ_{ijkl} has local support, and i, j, k, and l roughly correspond to translations in u, v, s, and t, respectively. Note, however, that the translated basis functions will not in all cases have the same shape, i.e. $\Psi_{i'j'k'l'}(u,v,s,t)$ may not be an exact copy of $\Psi_{ijkl}(u,v,s,t)$ in general.

We also define two additional sets of basis functions, one for *measuring* and one for *reconstruction*. For reconstruction we use a 1D basis $\{\Phi_i\}_{i \in \mathbf{Z}}$ with the property $\Phi_i(x) = \Phi(x+i)$. The 4D reconstruction basis is then given as the tensor product basis

$$\Phi_{ijkl}(u,v,s,t) := \Phi_{ij}(u,v) \cdot \Phi_{kl}(s,t) = \Phi_i(u) \cdot \Phi_j(v) \cdot \Phi_k(s) \cdot \Phi_l(t). \tag{2}$$

For measurement, we use the *biorthogonal* (or *dual*) $\{\Phi'_i(x)\}_{i \in \mathbf{Z}}$ of the reconstruction basis with

$$\int_{-\infty}^{\infty} \Phi'_i(x) \cdot \Phi_j(x)\, dx = \begin{cases} 1 & ;\text{ if } i = j \\ 0 & ;\text{ else} \end{cases} \tag{3}$$

and again we use a tensor-product construction for the 4D basis. The notation is summarized in Table 1.

3.1 Measured Irradiance

Our approach is based on being able to measure the irradiance $E_{mn}(s,t)$ in the (s,t)-plane that is caused by the incident radiance $L_{mn}(u,v,s,t) = \Phi'_{mn}(u,v) \cdot L(u,v,s,t)$. In Section 4 we will discuss several physical setups for performing this kind of measurement. The result of such a measurement is

$$\begin{aligned} E_{mn}(s,t) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{\cos^2\theta(u,v,s,t)}{R(u,v,s,t)^2} \cdot \Phi'_{mn}(u,v) \cdot L(u,v,s,t)\, du\, dv \\ &\approx \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{\cos^2\theta(u,v,s,t)}{R(u,v,s,t)^2} \cdot \Phi'_{mn}(u,v) \cdot \sum_i\sum_j\sum_k\sum_l \Psi_{ijkl}(u,v,s,t) \cdot L_{ijkl}\, du\, dv \quad \text{by Equation 1.} \\ &= \sum_i\sum_j\sum_k\sum_l \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{\cos^2\theta(u,v,s,t)}{R(u,v,s,t)^2} \cdot \Phi'_{mn}(u,v) \cdot \Psi_{ijkl}(u,v,s,t) \cdot L_{ijkl}\, du\, dv. \end{aligned} \tag{4}$$

Here, $\cos^2\theta(u,v,s,t)/R(u,v,s,t)^2$ is a geometric term composed of the distance R of the point on the (u,v)-plane from the point on the (s,t) plane, as well as the cosine of the angle θ between the plane normals and the vector connecting the two points. Note that this term also accounts for any differences in the parameterizations on the two planes (i.e. different grid spacings).

3.2 Exact Reconstruction

We now describe an exact reconstruction algorithm given the measurements E_{mn}. To this end, we first define what the relationship between the basis functions Ψ_{ijkl} and the reconstruction and measurement bases should be. We define

$$\Psi_{ijkl}(u,v,s,t) := \frac{R(u,v,s,t)^2}{\cos^2\theta(u,v,s,t)} \cdot \Phi_{ij}(u,v) \cdot \Phi_{kl}(s,t). \tag{5}$$

Inserting this definition into Equation 4, and using the biorthogonality relationship (Equation 3) yields

$$\begin{aligned} E_{mn}(s,t) &= \sum_i\sum_j\sum_k\sum_l \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \Phi'_{mn}(u,v) \cdot \Phi_{ij}(u,v) \cdot \Phi_{kl}(s,t) \cdot L_{ijkl}\, du\, dv \\ &= \sum_k\sum_l \Phi_{kl}(s,t) \cdot L_{mnkl}. \end{aligned} \tag{6}$$

To determine which reconstruction filter to use, we now rewrite Equation 1 using Equations 5 and 6:

$$\begin{aligned} \tilde{L}(u,v,s,t) &= \sum_m\sum_n\sum_k\sum_l \Psi_{mnkl}(u,v,s,t) \cdot L_{mnkl} \\ &= \sum_m\sum_n\sum_k\sum_l \frac{R(u,v,s,t)^2}{\cos^2\theta(u,v,s,t)} \cdot \Phi_{mn}(u,v) \cdot \Phi_{kl}(s,t) \cdot L_{mnkl} \\ &= \sum_m\sum_n \frac{R(u,v,s,t)^2}{\cos^2\theta(u,v,s,t)} \cdot \Phi_{mn}(u,v) \cdot E_{mn}(s,t) \end{aligned} \tag{7}$$

This indicates that we can *exactly* reconstruct $\tilde{L}$, the projection of L into the basis $\{\Psi_{ijkl}\}$ using the reconstruction filter $R(u,v,s,t)^2/\cos^2\theta(u,v,s,t) \cdot \Phi_{mn}(u,v)$.

Unfortunately, this reconstruction filter does contain a shift-variant component in form of the geometric term. This makes the reconstruction slightly more expensive than one would hope – ideally, the reconstruction step should only involve a convolution with a shift-invariant filter kernel. In the next section, we discuss an approximate reconstruction using a shift-invariant filter kernel.

3.3 High-Speed Approximate Reconstruction

For a high-speed reconstruction with a shift-invariant filter kernel we define Ψ_{ijkl} as a tensor-product basis: $\Psi_{ijkl}(u,v,s,t) := \Phi_{ij}(u,v) \cdot \Phi_{kl}(s,t)$. From this, we get

$$E_{mn}(s,t) = \sum_i \sum_j \sum_k \sum_l \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\cos^2 \theta(u,v,s,t)}{R(u,v,s,t)^2} \cdot \Phi'_{mn}(u,v) \cdot \Phi_{ij}(u,v) \cdot \Phi_{kl}(s,t) \cdot L_{ijkl} \, du \, dv \tag{8}$$

Unfortunately, the geometric term depends on the integration variables u and v, and can therefore not be moved outside the integral. However, if the distance d between the (u,v)- and the (s,t)-plane is large compared to the support of $\Phi_{ij}(u,v)$, and if θ is small, then this term is well approximated by one constant per point on the (s,t)-plane:

$$\frac{\cos^2 \theta(u,v,s,t)}{R(u,v,s,t)^2} \approx \frac{\cos^2 \theta_{kl}(s,t)}{R_{kl}(s,t)^2}. \tag{9}$$

This yields an approximation of the measured irradiance:

$$\begin{aligned} E_{mn}(s,t) &\approx \sum_k \sum_l \frac{\cos^2 \theta_{kl}(s,t)}{R_{kl}(s,t)^2} \cdot \Phi_{kl}(s,t) \cdot \sum_i \sum_j L_{ijkl} \cdot \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Phi'_{mn}(u,v) \cdot \Phi_{ij}(u,v) \, du \, dv \\ &= \sum_k \sum_l \frac{\cos^2 \theta_{kl}(s,t)}{R_{kl}(s,t)^2} \cdot \Phi_{kl}(s,t) \cdot L_{mnkl}. \end{aligned} \tag{10}$$

Since both the approximate geometric term and $\Phi_{kl}(s,t)$ are known, it is in principle possible to compute

$$E'_{mn}(s,t) \approx \sum_k \sum_l \Phi_{kl}(s,t) \cdot L_{mnkl} \tag{11}$$

by de-convolution. In practice, this is only feasible for basis functions $\Phi_{kl}(s,t)$ with a small support. We do not expect this to be a major problem, however, since the practical measurement setups that will be discussed in Section 4 have a very high resolution on the (s,t)-plane, so that very simple basis functions, such as bilinear interpolation, can be used. Again, we apply the definition of Ψ_{ijkl} to determine the appropriate reconstruction filter:

$$\begin{aligned} \tilde{L}(u,v,s,t) &= \sum_m \sum_n \sum_k \sum_l \Psi_{mnkl}(u,v,s,t) \cdot L_{mnkl} \\ &= \sum_m \sum_n \sum_k \sum_l \Phi_{mn}(u,v) \cdot \Phi_{kl}(s,t) \cdot L_{mnkl} \\ &\approx \sum_m \sum_n \Phi_{mn}(u,v) \cdot E'_{mn}(s,t). \end{aligned} \tag{12}$$

Thus, an *approximate* reconstruction of $\tilde{L}$, the projection of L into the basis $\{\Psi_{ijkl}\}$ is obtained using the shift-invariant reconstruction filter $\Phi_{mn}(u,v)$. For example, if Φ is a hat function, the reconstruction process is reduced to quadri-linear interpolation, which is frequently used for light field rendering algorithms. The quality of this approximation depends on the error introduced by assuming the geometric term constant over the support of basis function $\Phi_{ij}(u,v)$ in Equation 9. Some error estimates for various setups will be given in the following.

3.4 Estimates of the Approximation Error

As we move the geometric term out of the integral in Equation 10, we make the assumption that this term is close to constant over the support of the filter $\Phi'_{mn}(u,v)$.

In analyzing the geometric term for a two-plane parameterization, we get

$$\frac{\cos^2 \theta(u,v,s,t)}{R(u,v,s,t)^2} = \cos^4 \theta =: g(\theta),$$

if assume the distance between the planes to be unit-size. In order to estimate the error introduced by approximating $g(\theta)$ by a constant term, we define the absolute error

$$E_a(\theta_0) := \max_{\theta \in \mathcal{F}} g(\theta) - \min_{\theta \in \mathcal{F}} g(\theta)$$

where $\mathcal{F}$ is the support of a filter $\Phi'_{mn}(u, v)$ centered at an angle θ_0 from a measurement point on the (s, t)-plane. Similarly, we define the relative error as

$$E_r(\theta_0) := \frac{\max_{\theta \in \mathcal{F}} g(\theta) - \min_{\theta \in \mathcal{F}} g(\theta)}{g(\theta_0)}.$$

Figure 3 shows a plot of both the relative and the absolute error for several geometric setups. Firstly, we have a situation where the diameter of a filter is 10% of the distance between the two planes. Filter sizes of $5cm$ at plane spacings of $50cm$ would be a typical configuration for setups 1 and 2 as described in the next section.

We then were interested in an extreme situation, where the filters are made as small as possible compared to the plane spacing. We plotted an error curve for a filter size of $2cm$ with a plane distance of $1m$. This would certainly push the limits in terms of brightness on the measurement plane $\mathcal{M}$.

Finally, we wanted to check the quality of approximation for the other extreme, where the two planes are extremely close together (filter width equals plane spacing).

As can be seen from Figure 3, the first and second scenario produce reasonable results for angles θ up to $45°$. The maximum relative error for the first scenario is about 10%, while that for the second scenario is below 2%. This should be sufficient for a number of applications that do not need the highest precision.

In the third scenario, however, the error explodes even for relatively small angles of θ. This means the approximation cannot be used when the filter size is large compared to the plane distance.

4 Possible Experimental Setups

There are several geometric setups that might be considered. One possibility is to create a camera obscura-like setup, where light shining through the optical filter on the sampling plane $\mathcal{S}$ illuminates a diffuse surface (the measurement plane $\mathcal{M}$). A conventional digital camera is used to measure the irradiance on that plane (this requires high-dynamic range imaging along the lines of Debevec and Malik [4]).

This basic idea can be realized in two flavors: firstly, the filter itself could be moved relative to both the light source and the $\mathcal{M}$ plane, as well as the camera (Figure 4). The problem with this approach is that this setup is mechanically challenging to set up in a precise way.

An alternative possibility is to keep the observing camera, the $\mathcal{M}$ plane, and the filter in a fixed relative position (maybe even a common housing), and to move the light source relative to this setup (Figure 5). This should be mechanically much more feasible. Another advantage of this setup is that it corresponds to an $\mathcal{M}$ plane at infinity (similar to a planes at infinity described in the original work on light fields and Lumigraphs [8, 5]), which provides a nice separation of spatial and directional information. This separation can be advantageous for efficient rendering [7]. Finally, by implementing other motions than translations for the light source, one can implement other light field geometries than the two-plane parameterization. For example, a cylindrical scan could be implemented easily with a shift-invariant filter (other geometries would require different filter kernels for different filter positions which may not be feasible).

A shortcoming of both approaches discussed so far are calibration issues with respect to geometric optics and photometry. Because the camera looks at the $\mathcal{M}$ plane at an angle, perspective distortion and lens distortion have to be accounted for. Photometrically, it is difficult to avoid or compensate for interreflections in the setup (if camera, filter, and $\mathcal{M}$ plane are put into an enclosure), or residual light from the outside environment (if they are not). The material of the $\mathcal{M}$ plane is likely not 100% diffuse, which may pose a problem (although the diffuse assumption should mostly hold if the grazing angles can be avoided). Also, the material may not be as homogeneous as one would want, causing some undesirable spatial patterns in the measurements. Finally, due to geometric constraints, the distance between $\mathcal{M}$ and $\mathcal{S}$ plane may have to be relatively large so that the illumination level on the $\mathcal{M}$ plane might be quite low for darker light sources. This would mean that the camera would have to be operated at long exposure times which reduces the signal-to-noise ratio.

A possible way for avoiding these problems is a third setup, in which the CCD element of the digital camera is directly used as th e $\mathcal{M}$ plane, and the camera optical system is replaced by the filter (Figure 6). This setup is mechanically by far the easiest to realize, the CCD surface is precision manufactured, and the camera interior is optimized to avoid interreflections. The distance between filter and measurement plane is now much lower, which avoids long exposure times. However, this smaller distance also makes the measurements less suitable for the approximate reconstruction (as discussed in the previous section. Another potential problem could be a narrower field of view dictated by the relatively small size of the CCD chip.

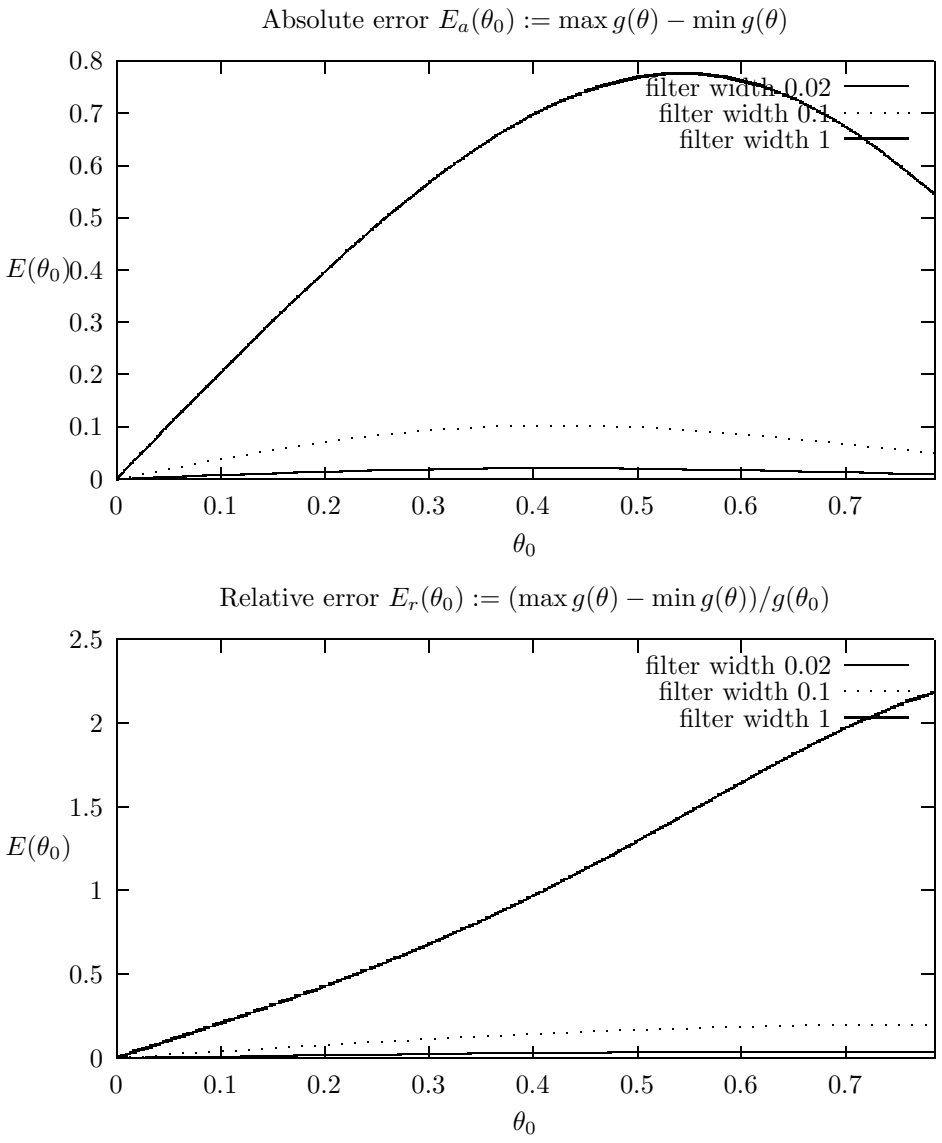

Figure 3: Absolute error (top) and relative error (bottom) of the approximation of the geometry term.

Table 4 summarizes the advantages and disadvantages of the different setups. Setup 1 clearly looks like the least promising approach. The relative merits of setups 2 and 3 will have to be determined experimentally.

5 Future Work

There is a host of possibilities for future work. These include

- Experimenting with different basis functions, trading off different degrees of smoothness for narrower or wider support.
- Exploring the possibilities for using non-planar sampling geometry. Some of these are straightforward (e.g. cylindrical sampling geometries), but in the general case this will require shift-variant filter kernels with all the disadvantages that implies (i.e. the need for different physical filters and the associated alignment problems).
- Hierarchical and adaptive sampling (Wavelets). This is strongly related to the choice of basis function and non-planar geometry.

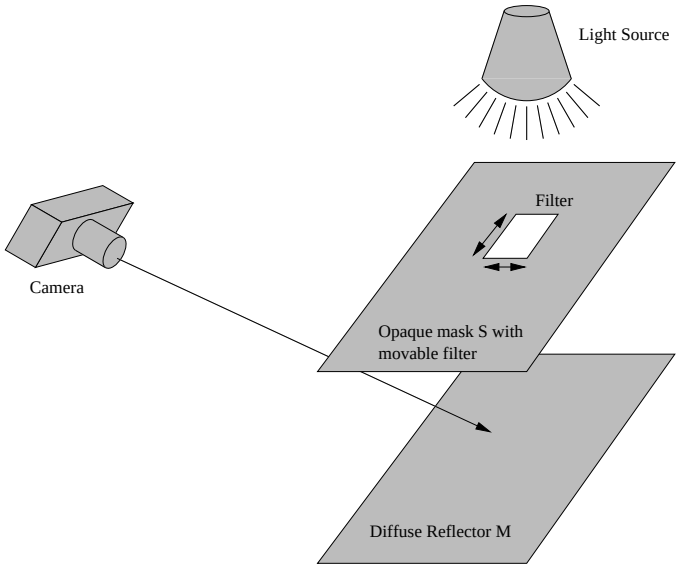

Figure 4: Setup 1

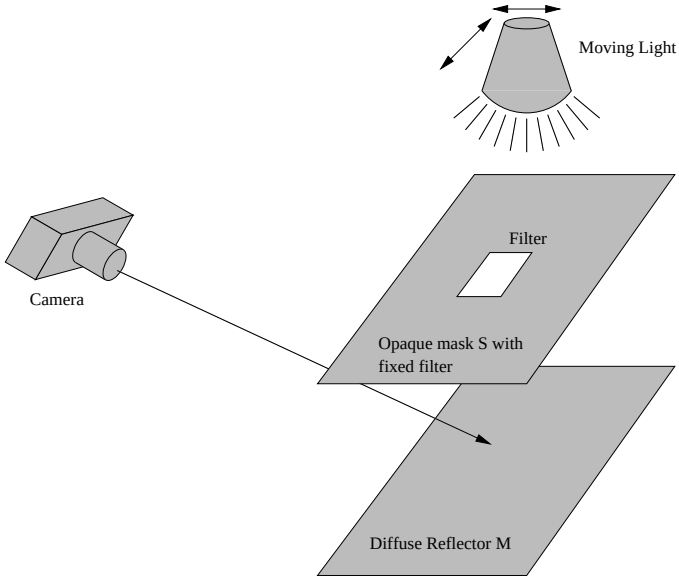

Figure 5: Setup 2

- Some of the issues arising from non-planar geometry could be solved with filters that can be changed electronically, for example by replacing the filter with an LCD panel. This approach, however, poses its own challenges, including polarization of light (meaning that already polarized light might get lost in the measurement) and limited optical density (i.e. limited contrast) of the current displays.

References

[1] Ian Ashdown. Near-Field Photometry: A New Approach. *Journal of the Illuminating Engineering Society*, 22(1):163–180, Winter 1993.

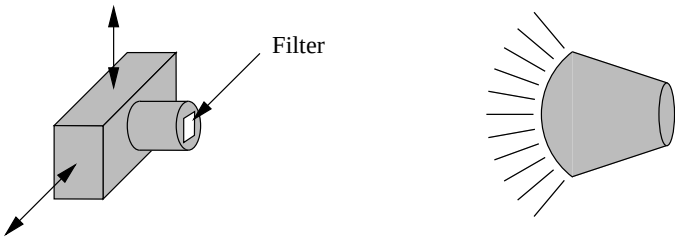

Figure 6: Setup 3

	Setup 1	Setup 2	Setup 3
Mechanical complexity	-	○	+
Mechanical calibration effort	-	○	+
Photometric calibration effort	-	-	+
Utilization of light intensity	○	○	+
Quality of approximate reconstruction	+	+	-
Separation of spatial and directional resolution	-	+	+
Wide range of directional samples	+	+	○
Extensible to non-planar sampling space	-	+	+

Table 2: A comparison of the different measurement setups.

[2] Ian Ashdown. Near-Field Photometry: Measuring and Modeling Complex 3-D Light Sources. In *ACM SIGGRAPH '95 Course Notes - Realistic Input for Realistic Images*, pages 1–15, 1995.

[3] Ian Ashdown. Making near-field photometry practical. In *Proc. of IESNA Annual Conference in Seattle*, August 1997.

[4] Paul Debevec and Jitendra Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 369–378. ACM, August 1997.

[5] Steven J. Gortler, Radek Grzeszczuk, Richard Szelinski, and Michael F. Cohen. The Lumigraph. In *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 43–54, August 1996.

[6] Michael W. Halle. Holographic stereograms as discrete imaging systems. In *Practical Holography VIII*, volume 2176, pages 73–84. SPIE, February 1994.

[7] Wolfgang Heidrich, Jan Kautz, Philipp Slusallek, and Hans-Peter Seidel. Canned lightsources. In *Rendering Techniques '98*, pages 293–300, June 1998.

[8] Marc Levoy and Pat Hanrahan. Light field rendering. In *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 31–42, August 1996.