

Accuracy of Spherical Harmonic Approximations for Images of Lambertian Objects under Far and Near Lighting

Darya Frolova, Denis Simakov, and Ronen Basri

Dept. of Computer Science and Applied Math,
The Weizmann Institute of Science,
Rehovot 76100, Israel,
{darya.frolova,denis.simakov,ronen.basri}@weizmann.ac.il

Abstract. Various problems in Computer Vision become difficult due to a strong influence of lighting on the images of an object. Recent work showed analytically that the set of all images of a convex, Lambertian object can be accurately approximated by the low-dimensional linear subspace constructed using spherical harmonic functions. In this paper we present two major contributions: first, we extend previous analysis of spherical harmonic approximation to the case of *arbitrary objects*; second, we analyze its applicability for *near light*. We begin by showing that under distant lighting, with uniform distribution of light sources, the average accuracy of spherical harmonic representation can be bound from below. This bound holds for objects of arbitrary geometry and color, and for general illuminations (consisting of any number of light sources). We further examine the case when light is coming from above and provide an analytic expression for the accuracy obtained in this case. Finally, we show that low-dimensional representations using spherical harmonics provide an accurate approximation also for fairly near light. Our analysis assumes Lambertian reflectance and accounts for attached, but not for cast shadows. We support this analysis by simulations and real experiments, including an example of a 3D shape reconstruction by photometric stereo under very close, unknown lighting.

1 Introduction

Methods for solving various Computer Vision tasks such as object recognition and 3D shape reconstruction under realistic lighting often require a tractable model capable of predicting images under different illumination conditions. It has been shown in [3] that even for the simple case of Lambertian (matt) objects the set of all images of an object under varying lighting conditions occupies a volume of unbounded dimension. Nevertheless many researchers observed that in many practical cases this set lies close to a low-dimensional linear subspace [4,6,18]. Low-dimensional representations have been used for solving many Computer Vision problems (e.g., [8,10,17]).

Low dimensional representations of lighting have been recently justified analytically in [1,14]. These studies show that the set of all Lambertian reflectance

functions (the mapping from surface normals to intensities) is, to an accurate approximation, low dimensional, and that this space is spanned by the low order *spherical harmonics*. Explicit spherical harmonic bases have been used to solve a number of important problems: object recognition [1], photometric stereo [2], reconstruction of moving shapes [16], and image rendering [15].

The introduction of spherical harmonic analysis provides a useful tool for handling complex illumination, but this pioneering work [1,14] is incomplete in a practically important aspect: the analysis in [1,14] is not easily generalized for the case of arbitrary object shapes and albedos.

In this paper we consider the case of Lambertian reflectance allowing for attached, but not for cast shadows. Thus, our analysis is applicable to convex objects illuminated by arbitrary combinations of point and extended sources. We begin by showing that under distant lighting the average accuracy of spherical harmonic representations can be bound from below by a bound that is *independent of the shape* of the object. For this result we assume that lighting can be cast on an object from any direction with equal probability, and that the distribution of the intensity of lighting is independent of its direction. We further consider a second case in which lighting is illuminating the object only from above, and derive an expression that allows us to calculate the accuracy of the spherical harmonic representation in this case.

While we consider a *single* expression for the harmonic basis there are studies that seek to build an *optimal basis* for every specific object or illumination. Ramamoorthi in [13] presents analytical construction of an optimal basis for the space of images. His analysis is based on spherical harmonics, and the images are taken under point light sources (uniformly distributed). The results of [13] are generalized and extended in [11,12] for different illumination distributions and materials. While they consider specific object geometries, our goal is to bound from below the approximation accuracy for arbitrary objects.

In the second part of our paper we analyze what happens if we relax the assumption of infinitely distant illumination, and show that spherical harmonics still provide a good basis even for fairly close light. We find what distance to the light can be considered infinite, as far as a spherical harmonic approximation is concerned. Our results show that although the approximation accuracy can be very bad for extremely close light, it rapidly increases as the distance to the light grows and even at rather small distances we achieve quite a good accuracy.

The assumption of infinitely distant light greatly simplifies the analysis of illumination effects, and so it is widely utilized in Computer Vision studies. While there are studies that incorporate near light effects (as in [7]), we are unaware of previous theoretical analysis of this factor.

The paper is divided as follows. In Section 2 we briefly review the use of spherical harmonics to represent lighting and reflectance. In Section 3 we derive lower bounds on the accuracy of spherical harmonic representations for objects of arbitrary shape and albedos under infinitely distant lighting. Finally, in Section 4 we examine the case of light sources at a finite distance from an object. Proofs are omitted for lack of space and will appear in a technical report.

2 Overview: Approximation by Spherical Harmonics

Basri and Jacobs [1] and Ramamoorthi and Hanrahan [14] constructed an analytically derived representation of the images produced by a convex, Lambertian object illuminated by distant light sources. Below we provide a brief outline of their results.

According to Lambert's law [9], which states that matt materials reflect light uniformly in all directions, a surface point with normal \mathbf{n} and albedo ρ illuminated by light arriving in direction \mathbf{l} and intensity i reflects light according to the following equation:

$$E_{\text{direc}} = \max(0, \langle \rho \mathbf{n}, i \mathbf{l} \rangle) , \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the usual inner product between vectors.

If we now consider a collection of directional (point) light sources placed at infinity, we can express the intensity of lighting as a non-negative function on the unit sphere $i(\mathbf{l})$. We can then express $i(\mathbf{l})$ as a sum of spherical harmonics (similar to a Fourier basis for R^n). We denote spherical harmonics by Y_{nm} ($n = 0, 1, 2, \dots$; $m = -n, \dots, n$). Then, $i(\mathbf{l}) = \sum_{n=0}^{\infty} \sum_{m=-n}^n \ell_{nm} Y_{nm}(\mathbf{l})$. Reflectance is then obtained from lighting by a convolution on the surface of a sphere, and using the Funk-Hecke theorem (see, e.g., [5]) the intensity of a point in an image E is given by:

$$E = \sum_{n=0}^{\infty} \sum_{m=-n}^n \alpha_n \ell_{nm} (\rho Y_{nm}(\mathbf{n})) , \quad (2)$$

with $\alpha_n = \pi, 2\pi/3, \pi/4, \dots$. For the specific case of a single directional source of intensity i and direction \mathbf{l} we have $\ell_{nm} = i Y_{nm}(\mathbf{l})$, and the harmonic expansion becomes

$$E_{\text{direc}} = \sum_{n=0}^{\infty} \sum_{m=-n}^n \alpha_n (i Y_{nm}(\mathbf{l})) (\rho Y_{nm}(\mathbf{n})) , \quad (3)$$

The coefficients tend to zero as $O(n^{-2.5})$ when $n \rightarrow \infty$. For this reason we can limit ourselves to just a few low order harmonics:

$$E \approx E^N = \sum_{n=0}^N \sum_{m=-n}^n \alpha_n \ell_{nm} (\rho Y_{nm}(\mathbf{n})) , \quad (4)$$

and evaluate the quality of the approximation using the *relative squared error*, defined as:

$$\varepsilon = \frac{\|E - E^N\|^2}{\|E\|^2} , \quad (5)$$

with the norm $\|f\|^2$ of a function f is defined as the integral of f^2 over its entire domain.

The quality of this approximation depends on the frequencies present in the lighting function $i(\mathbf{l})$. Consider the reflectance of a sphere with uniform

albedo (the so called *reflectance function*). If the sphere is illuminated by a single directional source then the approximation error is given by:

$$\varepsilon_{1,sph} = \frac{\sum_{n=N}^{\infty} (2n+1)\alpha_n^2}{\sum_{n=0}^{\infty} (2n+1)\alpha_n^2}. \quad (6)$$

(This is obtained from (5) using the orthonormality of Y_{nm} .) In particular, for orders 0, 1 and 2 the relative squared error is, respectively, 65.5%, 12.5% and 0.78% (corresponding to accuracy of 37.5%, 87.5% and 99.22%). The approximation is better if lighting contains mainly low frequency components and is worse if lighting includes many high frequencies, but even with high frequencies the error is bounded [1].

Reflectance functions, however, capture only general properties of images of Lambertian objects under specific lighting and ignore properties of a particular object. In particular, objects differ in shape and color, giving rise to different distributions of normals and albedos. In addition, foreshortening distorts the distribution of normals, and, due to occlusion, only the normals facing the camera are visible in an image. The accuracy obtained by low order harmonic approximations for images of objects change as well. In particular, there exist lighting conditions under which low order harmonic approximations of the images of objects are arbitrarily bad. However, we will show in the next section that on average the low order harmonic approximations for images of objects of *arbitrary* shape and albedo are accurate, and that the accuracies derived for reflectance functions provide in fact lower bounds to the average accuracies for any convex object.

3 Infinitely Distant Light

3.1 Basic Case: Uniformly Distributed Point Light Source

Lambert's law maintains a useful duality property. The formula (1) describing the light reflected by a point due to a directional source is symmetric, so that exchanging albedo ρ by light intensity i or normal \mathbf{n} by light direction \mathbf{l} maintains the same amount of light reflected. This duality relation is maintained also if we consider a discrete set of K surface points and a discrete set of J directional sources. In addition, it is maintained by every term of the spherical harmonic representation (4) and consequently by the expression for the relative squared error (5). Below we use this duality property to prove an error bound for arbitrary object approximation.

The approximation error for some light configurations can be arbitrarily large [1], so we cannot hope to bound the error for any light. We can try instead to describe a typical case by averaging the error over different illumination conditions. We consider first the case of a single directional light source. Since we have no prior information about the direction of this source, we assume that it is

drawn with a uniform distribution over a sphere around the object. For now we assume that the intensity of the light source is fixed and relax this assumption later.

To compute the error for this case we can consider the dual problem. Consider a single point on the surface of the object under all directional lightings. If we exchange object for light we will obtain a sphere with uniform albedo illuminated by a directional light source with arbitrary direction and intensity. The approximation error then is given by (6) and is independent of both the direction of the light source or its intensity. This argument can be applied to every point on the surface of the object. For every such point we obtain by the dual formulation a sphere illuminated by a different directional source, but the accuracy of the approximation remains the same. Consequently, the average error too is given by (6). This implies that for an arbitrary object the approximation error is constant depending only on the approximation order, and is independent on the object geometry and albedos.

We can readily extend this argument also to single directional sources with intensities drawn from an arbitrary distribution. All we need to assume is that the distribution of intensity is independent of direction.

Conclusion: **if** a convex Lambertian object is illuminated by a single directional light source, uniformly distributed over the sphere, **then** the accuracy of the spherical harmonic approximation does not depend on the object geometry and albedo.

We verified this conclusion on several examples. Note that in the analysis we use harmonic decompositions in which the coefficients are determined so as to optimally fit the reflectance of a uniformly sampled sphere. These may not be the optimal coefficients for an object. Thus, these derivations only give *lower bounds* on the accuracy of low order harmonic approximations.

Table 1 shows the accuracy obtained by simulations with a 4D (first order) and 9D (second order) harmonic approximation for various objects illuminated by single directional sources. In each case we used random albedos (uniform albedos lead to similar accuracies). In the case of the face we estimated albedo by averaging 15 images of the face. Note that both the actual harmonic approximations and SVD provide very similar accuracies, indicating that spherical harmonics indeed form an optimal basis. In addition, mainly due to foreshortening, those accuracies are slightly higher than the bound (about 99.5% compared to 99.22% in the 9D case and 94-98% compared to 87.5% in the 4D case).

3.2 Multiple Light Sources

In practice objects are often illuminated with multiple light sources. Does the result we obtained for a single light source hold for more general illumination configurations? We address this question below.

We consider the case of lighting consisting of multiple sources with (possibly) different intensities. We present an expression for the approximation error

Table 1. Approximation accuracy obtained for images of various objects illuminated by single directional sources. For each case we show the bound computed numerically (SHB – spherical harmonic bound), the actual harmonic approximation with optimal coefficients obtained using least squares (LSH – least squares harmonics), and the best low dimensional approximation obtained with SVD.

4D approximation				9D approximation			
Object	SHB	LSH	SVD	Object	SHB	LSH	SVD
Sphere	87.47	87.57	87.60	Sphere	99.21	99.22	99.22
Hemisphere	87.54	95.56	95.73	Hemisphere	99.22	99.45	99.47
Random	87.53	94.00	94.31	Random	99.22	99.45	99.47
Real face	87.58	97.73	98.00	Real face	99.23	99.66	99.70

depending on the number of light sources and their intensities, and give a meaningful analysis of this dependence. As before, consider an object of arbitrary shape and albedo and assume that the direction of the light sources are drawn from a uniform distribution and that the intensities are drawn from a distribution that is independent of direction.

Let us consider illumination consisting of K point light sources with intensities i_k and directions \mathbf{l}_k : $i(\mathbf{l}) = \sum_{k=1}^K i_k \delta_{\mathbf{l}_k}$. We assume that the directions of these point sources are distributed independently and uniformly over the sphere (while their intensities are fixed). We denote the relative approximation error of order N (defined as in (5)) by $\varepsilon_{K, sph}$.

Evaluating this error we obtain the following expression:

$$\varepsilon_{K, sph} = \frac{\varepsilon_{1, sph}}{1 + V}, \tag{7}$$

where V is determined by the light intensities:

$$V = \frac{3}{8} \left(\left(\sum_{k=1}^K i_k \right)^2 / \sum_{k=1}^K i_k^2 - 1 \right) = \frac{3}{8} \left(\frac{\|\mathbf{i}\|_{l_1}^2}{\|\mathbf{i}\|_{l_2}^2} - 1 \right), \tag{8}$$

with $\|\mathbf{i}\|_{l_1}$ and $\|\mathbf{i}\|_{l_2}$ respectively denote the l_1 and l_2 norms of the vector \mathbf{i} of light intensities.

V can be interpreted as a measure of non-uniformity of light intensities. It is always non-negative, and equal to zero for a single directional source. V is largest (and the error is smallest) when all the intensities i_k are equal (for a fixed number of sources K). In this case $V = \frac{3}{8}(K - 1)$ and $V \rightarrow \infty$ when the number of light sources K tends to infinity. Here we obtain in the limit the uniform ambient light and, not surprisingly, the approximation error (7) becomes zero.

Conclusion: **if** we consider an arbitrary number of (uniformly distributed, independent) multiple directional sources, **then** the accuracy of a low order harmonic approximation for a convex Lambertian object of arbitrary shape and albedos is not less than with a single light source. In other words, a single light source is the worst case illumination for the spherical harmonic approximation.

3.3 Light From Above

Our results thus far were obtained under the assumption that directional light sources are distributed uniformly over the entire sphere (the same assumption is adopted in [11,12,13]). But in reality we often meet the situation that light is coming mainly from above. To incorporate this prior knowledge we substitute the operation of averaging over the sphere for averaging over the upper hemisphere. We derive a bound for this case that is not constant for every order, but depends on the object normals and albedos.

The formula we derive allows to compute the average approximation error for any object illuminated by a random directional light source on the hemisphere and to analyze how object geometry and albedos influence this error. Our analysis shows that, unlike in the previous case, there exist objects for which the average harmonic approximation is arbitrarily bad. However, in a typical experimental setup (horizontally oriented camera) due to foreshortening the error is typically almost the same as with light distributed over the entire sphere.

Consider an arbitrary object illuminated by a single directional source \mathbf{l} with intensity i , which is uniformly distributed over the upper hemisphere. Using the harmonic expansion (3) of the image, for every surface normal and light the approximation error for $N \geq 2$ is given by:

$$\varepsilon_{1,hsp} = \frac{\varepsilon_{1,sph}}{1 + \bar{F}}. \quad (9)$$

Here \bar{F} is a mean value of a function F (defined below), which depends on the normals and albedos of the object: $\bar{F} = \int_{object} \rho F(\theta) / \int_{object} \rho$, where θ is the angle between a surface normal and the vertical direction (θ varies from 0 to π). $F(\theta)$ is given by

$$F(\theta) = \frac{\sqrt{3}}{\pi} \sum_{n=0, n \neq 1}^{\infty} \alpha_n^2 \sqrt{2n+1} (Y_{n0}(\theta) Y_{10}(\theta) + \sqrt{2n(n+1)} Y_{n1}(\theta) Y_{11}(\theta)) \quad (10)$$

One can see the dependence of F (Figure 1) on the direction of the normals of the object.

Let us now analyze the expression (9). A positive value of \bar{F} reduces the error relative to the case of light distributed over the entire sphere. A negative value of \bar{F} increases this error (in the worst case to an arbitrarily large value). If most of the object normals (taking into account albedos) are directed upward then $\varepsilon_{1,hsp}$ is smaller than the error for the sphere $\varepsilon_{1,sph}$. And vice versa, the more normals look downward, the greater is the error.

Foreshortening also affects this error, as it affects the density of the sampled normals. A typical setting is when light comes from above (sunlight or indoor ceiling lights) and the camera's optical axis is horizontal. In this case we average over F scaled by $\sin \theta$. It appears that in this case the effect of foreshortening "helps" us: normals that spoil the approximation the most (directed toward the ground) occur less frequently due to foreshortening, while normals directed to

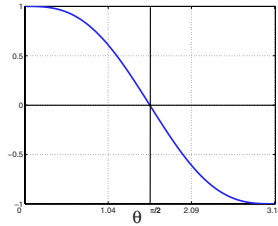


Fig. 1. The function $F(\theta)$. F is positive for $0 \leq \theta \leq \pi/2$ leading to improved accuracies for normals facing upward and negative otherwise, leading to decreased accuracies for normals facing downward.

the camera usually dominate the image. As a consequence the accuracy obtained with low order harmonic approximations is close to $1 - \varepsilon_{1, sph}$ ($F(\theta)$ is close to zero, see Figure 1).

Conclusion: **if** light distribution prior is non-uniform (from above in our analysis), **then** the accuracy of the spherical harmonic approximation depends on the object geometry. An object whose normals face the light (i.e. the dominant light direction, upwards in our analysis) is better approximated by spherical harmonics, than an object whose normals face the “dark side”.

In the typical setting that light is coming from above and the optical axis is horizontal, the approximation accuracy is usually close to the accuracy in the case of uniformly distributed light $1 - \varepsilon_{1, sph}$.

Table 2 shows the accuracy of the first five approximation orders for real and simulated objects. One can compare the results for two light distributions: when light source is distributed over the hemisphere (the first three rows) and when light source is distributed over the sphere (the bottom row).

Table 2. The first three rows show the n -th order approximation accuracy for different objects when light is coming from above. “Basic accuracy” refers to the accuracy when the light is distributed uniformly over the entire sphere.

Object	0	1	2	3	4	5	\overline{F}
Jurassic	26.99	85.40	99.09	99.09	99.77	99.77	-0.1440
Dinosaur	36.47	87.47	99.21	99.21	99.80	99.80	-0.0280
Face	58.86	91.77	99.49	99.49	99.87	99.87	0.5193
Basic accuracy	37.50	87.50	99.22	99.22	99.80	99.80	–

4 Near Light

Computer vision studies of lighting largely make the assumption of infinitely distant light. But in many images light is coming from nearby sources or is reflected from nearby objects, as for example in images taken in an indoor environment. It is then important to relax this assumption, or at least to determine what distance is sufficient to be considered infinite. In this section we attack this problem theoretically by analyzing a simplified near light model and practically by conducting experiments with near lighting.

Handling near light is complex because of two main problems. First, light originating from the same light source approaches different points on the object from different directions. Second, we can no longer assume that the intensity arriving at each point on the object is constant since it decreases as the squared distance to the light source. These problems imply in particular that we have to take into account the position of light sources and not only their direction, and thus lighting is no longer a function defined on the surface of a sphere. However, we will show an approach to still uses spherical harmonics. As before, our analysis accounts for attached shadows and will ignore the effect of cast shadows. Notice that in the case of near light the extent of attached shadows vary according to the distance to the light source.

4.1 Harmonic Approximation with Light at Finite Distance

Here we present an analytical model of the approximation of the reflectance function by spherical harmonics taking into account light at a finite distance. The model we describe assumes that all light sources are placed in some fixed distance from the object. The full treatment of near lighting requires taking into account images obtained by illuminating an object by collections of light sources at different distances from an object.

Consider a sphere of radius r with unit albedo centered at the origin (see Figure 2). Let \mathbf{p} denote a point on the sphere, and let $\mathbf{n}(\mathbf{p})$ denote the surface normal at \mathbf{p} . (Note that $\mathbf{p} = r\mathbf{n}$.) Assume that the sphere is illuminated by a point light source drawn from a uniform distribution over the sphere of radius $r + R$ also centered at the origin. Denote by \mathbf{u} a unit vector pointing toward the light source (so the light source is positioned at $(r + R)\mathbf{u}$), and let i denotes the light source intensity.

Denote by \mathbf{l} the vector from \mathbf{p} to the light source, namely, $\mathbf{l} = (r + R)\mathbf{u} - \mathbf{p}$, so that \mathbf{l} is the incident direction of the source at \mathbf{p} . The intensity that arrives at \mathbf{p} due to this source then is given by $i/\|\mathbf{l}\|^2$. Using Lambert's law the reflected intensity at \mathbf{p} due to this incident light is given by $E_{point} = \frac{i}{\|\mathbf{l}\|^2} \max(\langle \mathbf{n}, \mathbf{l} \rangle, 0)$.

Now consider a lighting function $i(\mathbf{u})$ describing the intensity of the light emitted from any number of sources placed on the sphere of radius $r + R$, the light reflected by the sphere of radius r due to these sources can be written as a convolution on the surface of a sphere of the lighting function $i(\mathbf{u})$ with the kernel

$$k = \max\left(\frac{\langle \mathbf{n}, \mathbf{l} \rangle}{\|\mathbf{l}\|^2}, 0\right). \quad (11)$$

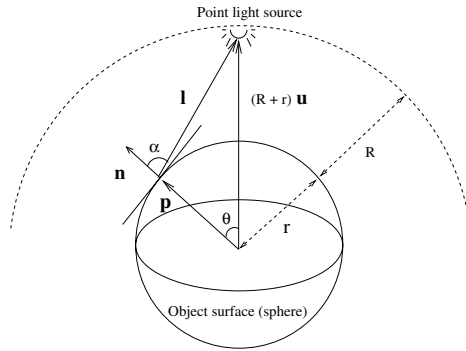


Fig. 2. A sphere of radius r is illuminated by a point light source. R is the distance between the sphere and the light. p denotes a point on the sphere and n denotes its normal vector. l is the incident lighting vector, and u is a unit vector directed toward the light source.

Expressing k as a function of θ , where θ is the angle between the light source vector h and the surface normal n we obtain:

$$k(\theta) = \max \left(\frac{(r + R) \cos(\theta) - r}{(R^2 + 2r(r + R)(1 - \cos(\theta)))^{3/2}}, 0 \right). \tag{12}$$

(Notice that k is a valid convolution kernel since it is rotationally symmetric about the north pole.) The harmonic expansion of $k(\theta)$ can be derived using the coefficients $k_n = 2\pi \int_0^\pi k(\theta) Y_{n0}(\theta) \sin(\theta) d\theta$. Integration can be limited to the positive portion of $k(\theta)$, so integration limits will now depend on the distances R and r . It is worth noting that unlike the case of infinitely distant light, harmonics of the odd orders are not eliminated by the kernel.

The relative squared energy concentrated in the first N harmonics is expressed by $\left(\frac{\sum_{n=0}^N k_n^2}{\sum_{n=0}^\infty k_n^2} \right)$. We can compute this relation for every finite N : the numerator is evaluated numerically and for the denominator we have the explicit formula:

$$\sum_{n=0}^\infty k_n^2 = \frac{\pi}{r^2 R^2} \left(\frac{2r(r + R) - 4rR + R^2 \log(1 + 2\frac{r}{R})}{4r(r + R)} \right). \tag{13}$$

4.2 Conclusions from the Simple Model

Figure 3 presents the dependence of the approximation accuracy on the distance to the light source. The distance is relative to the object size (R/r in the notation of our model). Each graph represents one approximation order (from zero to third). One can see that for extremely near light approximation accuracy is close to zero, but grows rapidly as the distance to the light increases approaching the accuracy values for infinitely distant light.

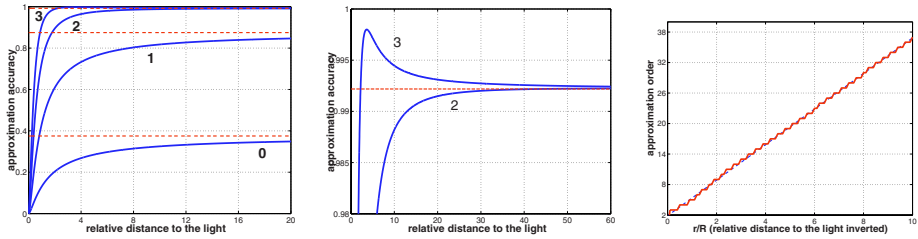


Fig. 3. Left: approximation accuracy as a function of the distance to the light source for approximation orders 0-3 (solid graphs). Dashed lines: asymptotes. Middle: a zoom-in version of the upper left portion of the left plot. Right: Order of approximation required to reach a 99.22% accuracy as a function of r/R .

The accuracy grows faster for higher approximation orders. This implies that for higher orders the light behaves as "infinitely distant" already at smaller distances. As an example consider the second and third orders (Figure 3(middle)). They both have the limit value of 99.22%. (We know from distant light analysis [1,14] that odd approximation orders higher than 1 do not contribute energy.) While the second order accuracy exceeds 99% from distance 13, the third order reaches this value already from the distance 2. Therefore for near light a third order approximation can be considerably more accurate than a second order.

As we increase the number of spherical harmonics we use, we can cope with arbitrarily close light. For example, to achieve an approximation accuracy of 99.22% we have the dependency shown in Figure 3(right). We see that the order of approximation we need is roughly inversely proportional to the distance of the light.

4.3 Experiments

We performed simulations on realistic objects to test the accuracy obtained with spherical harmonic approximations under near lighting.

We present results of simulations with a synthetic object model of a dinosaur's head ("Jurassic") obtained using the 3D Studio software. We used several light sources and moved them from very near positions to very far, rendering images for each light distance. Figure 4 shows the accuracies obtained for the first three approximation orders using the coefficients determined by our model and the coefficients obtained by least squares fitting.

We see that the approximation accuracy behaves similarly to the prediction of our model (apart from an undershoot for order zero). Starting with a small values for very near light, the accuracy for every order grows very fast and tends to its limit value for large distances. We see that even the distance scale in this simulation is very similar to the one obtained from the model.

We also performed near light experiments with two real objects: a human face and a dinosaur toy. 3D models of the objects were obtained with a laser scanner. Note that both objects are not exactly Lambertian, and some amount

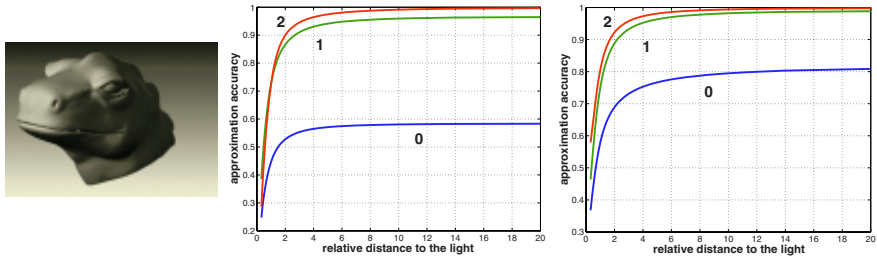


Fig. 4. Near light accuracy for Jurassic with harmonic approximations using the coefficients determined by our model (left) and the coefficients determined using least squares optimization (right).

Table 3. Least squares accuracy of harmonic approximations for a rendered model of a face (left) and a dinosaur (right). Approximate face radius is 15 cm, approximate dinosaur dimensions are: length 38 cm, width 8 cm, height 13 cm.



Distance	0	1	2	3
20 cm	77.61	91.14	92.20	92.47
50 cm	84.79	94.44	95.04	95.28
70 cm	89.56	95.00	95.40	95.63
100 cm	89.52	96.19	96.59	96.78
120 cm	92.72	97.07	97.28	97.40

Distance	0	1	2	3
20 cm	84.53	91.74	92.51	92.75
50 cm	86.49	92.82	93.49	93.68
80 cm	87.91	93.70	94.25	94.41
110 cm	88.08	93.30	93.90	94.03
140 cm	89.80	93.91	94.42	94.52

of cast shadows were present in the images. We tested the accuracies obtained for lighting at five distances using 15 pictures with varying lighting positions. The accuracies obtained are shown in Table 3. It can be seen that approximation accuracies are very high even for extremely near light.

Finally, we reconstructed the 3D shape of a dinosaur from 15 images obtained with lighting at 20 cm distance from the object using the photometric stereo method proposed in [2]. We used a first order (4D) harmonic approximation. This method uses factorization to recover lighting and shape fitting the obtained shape matrix to a spherical harmonic decomposition. The procedure allows the recovery of shape up to a 7 parameter scaled Lorentz ambiguity. To allow comparison of our results to ground truth we resolved this ambiguity manually. Figure 5 shows a subset of the images used for reconstruction, and Figure 6 shows the reconstruction obtained and for comparison the 3D shape obtained with a laser scanner. As can be seen although the object is illuminated by a very proximate lighting reconstruction is quite accurate. Similar reconstruction results were obtained with more distant lighting.



Fig. 5. Three out of 15 images with extremely near light (20 cm) used for photometric stereo reconstruction.

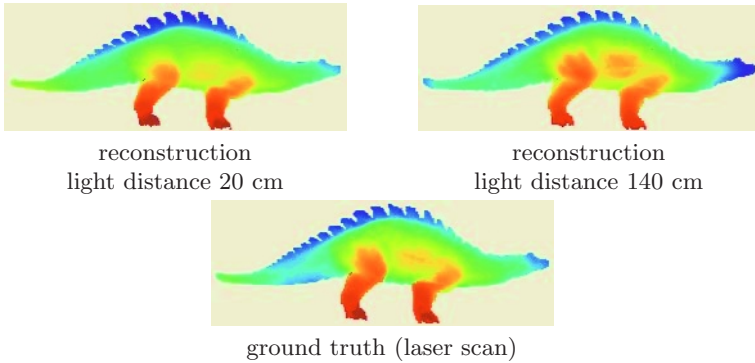


Fig. 6. Photometric stereo reconstruction results (depth maps). Colors (intensities) encode depth $z(x, y)$.

5 Summary

In this paper we have examined the use of spherical harmonic representations to model the images of Lambertian objects seen under arbitrary lighting, and extended its applicability to several cases of practical importance. We showed that under distant lighting, with reasonable assumptions on the distribution of light sources, the average accuracy of spherical harmonic representations can be bound from below independently of the shape and albedos of the object. We further examined the case of light coming from above and provided an analytic expression for the accuracy obtained in this case. Finally, we derived a model to compute the accuracy of low-dimensional harmonic representations under near lighting.

Our analysis demonstrates that spherical harmonic representations provide accurate modelling of lighting effects for a wide range of lighting conditions.

Acknowledgements. Research was supported in part by the European Community grant number IST-2000-26001 and by the Israel Science Foundation grants number 266/02. The vision group at the Weizmann Inst. is supported in part by the Moross Laboratory for Vision Research and Robotics.

References

1. R. Basri, D.W. Jacobs, "Lambertian reflectances and linear subspaces," *PAMI*, **25**(2), 2003.
2. R. Basri and D.W. Jacobs, "Photometric stereo with general, unknown lighting," *CVPR*, Vol. II: 374–381, 2001.
3. P. Bellhumeur, D. Kriegman. "What is the Set of Images of an Object Under All Possible Lighting Conditions?," *CVPR*: 270–277, 1996.
4. R. Epstein, P. Hallinan, A. Yuille. " 5 ± 2 Eigenimages Suffice: An Empirical Investigation of Low-Dimensional Lighting Models," *IEEE Workshop on Physics-Based Vision*: 108–116, 1995.
5. H. Groemer, *Geometric applications of Fourier series and spherical harmonics*, Cambridge University Press, 1996.
6. P. Hallinan. "A Low-Dimensional Representation of Human Faces for Arbitrary Lighting Conditions", *CVPR*: 995–999, 1994.
7. K. Hara, K. Nishino, K. Ikeuchi. "Determining Reflectance and Light Position from a Single Image Without Distant Illumination Assumption", *ICCV*: 560–567, 2003.
8. M. Kirby, L. Sirovich, "The application of the Karhunen-Loeve procedure for the characterization of human faces", *PAMI*, **12**(1): 103-108, 1990.
9. J. Lambert, "Photometria Sive de Mensura et Gradibus Luminis, Colorum et Umbrae," Eberhard Klett, 1760.
10. H. Murase, S. Nayar. Visual learning and recognition of 3D objects from appearance. *International Journal of Computer Vision*, **14**(1): 5–25, 1995.
11. P. Nillius, J.-O. Eklundh. "Low-Dimensional Representations of Shaded Surfaces under Varying Illumination", *CVPR*: 185–192, 2003.
12. P. Nillius, J.-O. Eklundh. "Phenomenological Eigenfunctions for Image Irradiance", *ICCV*: 568–575, 2003.
13. R. Ramamoorthi, "Analytic PCA construction for theoretical analysis of lighting in a single image of a Lambertian object," *PAMI*, **24**(10): 1322–1333, 2002.
14. R. Ramamoorthi, P. Hanrahan, "On the relationship between radiance and irradiance: determining the illumination from images of convex Lambertian object." *JOSA*, **18**(10): 2448–2459, 2001.
15. R. Ramamoorthi, P. Hanrahan, "An efficient representation for irradiance environment maps," *Siggraph*: 497–500, 2001.
16. D. Simakov, D. Frolova, R. Basri, "Dense shape reconstruction of a moving object under arbitrary, unknown lighting," *ICCV*: 1202–1209, 2003.
17. M. Turk, A. Pentland, "Eigenfaces for Recognition," *Journal of Cognitive Neuroscience*, **3**(1): 71–96, 1991.
18. A. Yuille, D. Snow, R. Epstein, P. Bellhumeur, "Determining Generative Models of Objects Under Varying Illumination: Shape and Albedo from Multiple Images Using SVD and Integrability", *IJCV*, **35**(3): 203–222, 1999.