Spherical Harmonics vs. Haar Wavelets: Basis for Recovering Illumination from Cast Shadows

Takahiro OKABE, Imari SATO, and Yoichi SATO
Institute of Industrial Science, The University of Tokyo
4-6-1 Komaba, Meguro-ku, Tokyo 153-8505, Japan
{*takahiro,imarik,ysato*}*@iis.u-tokyo.ac.jp*

Abstract

The problem of estimating an illumination distribution from images is called inverse lighting. For inverse lighting, three approaches have been developed based on specular reflection components, diffuse reflection components, and cast shadows. The present study provides theoretical insights as to why the approach based on cast shadows works in a reliable manner, and discusses what kind of basis functions are appropriate to be used for recovering illumination from cast shadows. First, we formalize the approach based on cast shadows by using spherical harmonics. Then, we analyze the approach in the frequency domain and show the advantages and the limitations of the approach. Second, motivated by the observations in the frequency domain, we propose an efficient method using Haar wavelets that provide compact supports and sparsity of coefficients. Finally, we report the results of experiments that compared the method using spherical harmonics and the method using Haar wavelets.

1. Introduction

The appearance of an object greatly depends not only upon poses of the object but also upon illumination conditions. Therefore, estimation/measurement of light sources in a scene is important for image understanding/synthesis such as face recognition [22, 5, 1, 13, 14] or augmented reality [3, 19].

The problem of estimating an illumination distribution from images is generally called *inverse lighting* [8]. To solve this problem, three approaches have mainly been proposed. The first one is based on specular reflection components observed on an object's surface [9, 3, 10]. The second one is based on shading of diffuse reflection components [8, 16, 25]. The third one makes use of shadows cast by objects in a scene [20].

With regard to the first two approaches, Ramamoorthi and Hanrahan have provided several important insights, based on the theoretical analysis in the frequency domain using spherical harmonics [16, 17]. For example, they showed that the second approach based on diffuse reflection components can estimate only low-frequency components of illumination. This explains the empirical observation reported by Marschner and Greenberg that inverse lighting using diffuse reflection components tends to become ill-conditioned [8].

On the other hand, little is known theoretically about the third approach based on cast shadows, although Sato *et al.* showed experimentally that inverse lighting using cast shadows works reliably even in the case of a Lambertian surface [20]. In the present paper, we formalize the approach based on cast shadows by using spherical harmonics, and analyze the approach in the frequency domain. In particular, we provide theoretical insights as to why the approach based on cast shadows works in a reliable manner. In addition, based on sampling theorem [4, 18], we derive the maximum frequency of band-limited illumination that can be estimated by using our method.

Our proposed method using spherical harmonics is appropriate for estimating illumination localized in the frequency domain. However, our method has two limitations. First, we usually observe only a part of the shadows cast by an object since we cannot observe regions behind the object. This prevents our method from estimating high-frequency components of illumination. Second, the assumption of band-limited illumination does not hold, because natural illumination usually contains high-frequency components. Consequently, aliasing is often caused in the estimated illumination distribution.

To cope with these limitations, we propose an efficient method using the simplest multi-resolution basis known as Haar wavelets [24]. There are three advantages in using Haar wavelets. The first one is that we can represent illumination of each incident direction at different resolutions, because the basis functions have compact supports with var-

0-7695-2158-4/04 $20.00 (C) 2004 IEEE

ious sizes. The second one is sparsity of the wavelet coefficients. Ng *et al.* showed, in the context of image synthesis, that an illumination distribution is efficiently approximated by using a small number of the basis functions [11]. The third one is the orthonormality of the basis functions. Then, the contribution of a basis function in terms of L_2 norm is simply obtained by the square of the corresponding coefficient.

Utilizing these advantages, our proposed method represents an illumination distribution by using a small number of the basis functions. However, in contrast to image synthesis under a given illumination distribution, we cannot know important basis functions *a priori* which make a large contribution to the illumination distribution. Accordingly, our method adaptively estimates illumination in a way similar to the procedure called the *oracle* [6]. Namely, we improve the resolution of illumination by adding new basis functions and discard those with smallest squared coefficients at the same time; we reliably estimate a small number of coefficients representing high-resolution illumination.

To compare the performance of our proposed methods, we conducted a number of experiments using both synthetic and real images. We confirmed that the method using Haar wavelets works successfully for recovering illumination from cast shadows.

The rest of the paper is organized as follows. In Section 2, we summarize our assumptions. In Section 3, we propose a method using spherical harmonics and analyze the method in the frequency domain. In Section 4, we propose an efficient method using Haar wavelets. We report the results of the experiments that compared the performance of our methods in Section 5. Finally, in Section 6, we present concluding remarks.

2. Assumptions

Before we study inverse lighting based on cast shadows in detail, we summarize the assumptions of our proposed methods. First, as is usual with inverse lighting, the geometry of a scene is assumed to be known. Second, we assume that the scene is illuminated by distant light sources. Therefore, an illumination distribution is represented by a function independent of locations within the scene. Finally, we do not take interreflection into account.

For the sake of simplicity, we consider a scene where an object is placed on a Lambertian plane with uniform albedo. Let (θ, ϕ) be the polar coordinates whose north pole $(\theta = 0)$ corresponds to the direction of the surface normal. Then, the brightness E at a point $\boldsymbol{x}$ on the plane is represented by

$$E(\boldsymbol{x}) = \int_0^{2\pi} \int_0^{\pi/2} V(\boldsymbol{x}, \theta, \phi) L(\theta, \phi) \cos\theta \sin\theta d\theta d\phi, \tag{1}$$

where $L(\theta, \phi)$ is the illumination radiance coming from the direction (θ, ϕ) and $V(\boldsymbol{x}, \theta, \phi)$ is the visibility function. At each point $\boldsymbol{x}$, $V = 0$ if the direction (θ, ϕ) is occluded by the object and $V = 1$ otherwise. Inverse lighting based on cast shadows is a problem of estimating the illumination radiance $L(\theta, \phi)$ from the brightness $E(\boldsymbol{x})$.

3. Inverse Lighting Using Spherical Harmonics

In order to investigate the advantages and the limitations of using cast shadows for inverse lighting, we formalize the approach based on cast shadows by using spherical harmonics and analyze that approach in the frequency domain[1].

3.1. Formalization

Spherical harmonics is an orthonormal basis defined over a unit sphere. In this paper, $Y_{lm}(\theta, \phi)$ $(l \geq 0, -l \leq m \leq l)$ denote spherical harmonics defined by

$$Y_{lm}(\theta, \phi) = N_{lm} P_l^m(\cos\theta) e^{im\phi}, \tag{2}$$

where $P_l^m(\cdot)$ and N_{lm} are the associated Legendre functions and the normalization constants respectively.

The assumption of distant illumination enables us to represent an illumination distribution by a function defined on the sphere. Then, we can expand the illumination distribution $L(\theta, \phi)$ by using spherical harmonics as

$$L(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} Y_{lm}(\theta, \phi). \tag{3}$$

Therefore, inverse lighting results in the problem of estimating the spherical harmonics coefficients L_{lm}.

Substituting equation (3) into equation (1), we can derive

$$E(\boldsymbol{x}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_{lm} T_{lm}(\boldsymbol{x}). \tag{4}$$

The coefficients $T_{lm}(\boldsymbol{x})$ are computed from the geometry of the scene as

$$\begin{aligned} T_{lm}(\boldsymbol{x}) \quad &\equiv \quad \int_0^{2\pi} \int_0^{\pi} V(\boldsymbol{x}, \theta, \phi) \max(\cos\theta, 0) \\ &\times \quad Y_{lm}(\theta, \phi) \sin\theta d\theta d\phi, \end{aligned} \tag{5}$$

where we replace $\cos\theta$ with $\max(\cos\theta, 0)$ for further discussion.

When we approximate the infinite sum in equation (4) by the sum up to $l_{\max}$, we obtain a linear equation with $(l_{\max} + 1)^2$ unknowns. Hence, we can solve a set of linear equations with respect to L_{lm} $(0 \leq l \leq l_{\max}, -l \leq m \leq l)$

[1] A part of this section is based on our technical report [12].

by observing the brightness at a sufficient number of points on the Lambertian plane. This is a basic idea for recovering illumination from cast shadows by using spherical harmonics.

3.2. Analysis in Frequency Domain

As described in Section 1, while empirical results show that inverse lighting based on cast shadows works reliably even in the case of a Lambertian surface, little is known theoretically. Accordingly, we consider inverse lighting based on cast shadows in the frequency domain and provide two theoretical insights: why the approach based on cast shadows works well, and to which frequency the approach can estimate illumination.

First, we discuss the former issue based on the discreetness of the visibility function. We define the transfer function[2] $T(\boldsymbol{x}, \theta, \phi)$ as

$$T(\boldsymbol{x}, \theta, \phi) \equiv V(\boldsymbol{x}, \theta, \phi) \times \max(\cos\theta, 0). \quad (6)$$

Thus, $T_{lm}(\boldsymbol{x})$ in equation (5) are spherical harmonics coefficients of the transfer function. Therefore, equation (4) means that the coefficients of the transfer function contribute to the brightness of the Lambertian surface.

Let us consider the case where $V(\boldsymbol{x}, \theta, \phi)$ always equals to 1, that is, the case of a convex Lambertian object. In this case, it has been shown that the coefficients of the transfer function decay rapidly with increasing l, and that more than 99% of its energy is captured by $l \leq 2$ [16, 17, 1]. This means that the transfer function acts as a low-pass filter, and high-frequency components of illumination have little or no contribution to the brightness. This is the main reason why inverse lighting based on diffuse reflection components can estimate only low-frequency components $(L_{lm}, l \leq 2)$.

On the other hand, in the case where the visibility function takes discrete values $\{0, 1\}$, some portions of the cosine curve in the transfer function collapse as shown in Figure 1. Thus, the transfer function generally has nonzero high-frequency components, depending on the visibility function determined by the geometry of a scene. This enables high-frequency components of illumination to contribute to the brightness of the Lambertian surface through those of the transfer function. Consequently, inverse lighting based on cast shadows has a potential for estimating high-frequency components.

As an obvious example, let us consider a pinhole camera: a surface covered by a box with a pinhole. In this case, a point on the surface is bright only when it is illuminated through the pinhole. Namely, the transfer function becomes a delta function whose spherical harmonics coefficients do not decay with increasing l. Thus, inverse lighting based on cast shadows works thoroughly in this case. However, depending on the geometry of the scenes, the frequency characteristics of transfer functions vary greatly. Hence, in practice, it is essential to know the maximum frequency of illumination to which the approach based on cast shadows can estimate from a given input image.

[2]The same transfer function is used in the context of forward rendering of cast shadows under low-frequency illumination [23].

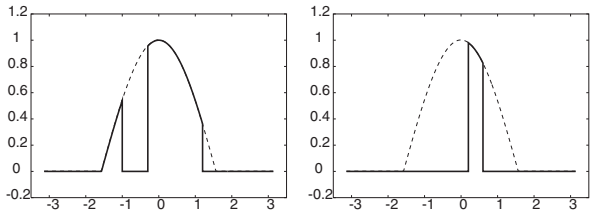

Figure 1. Slices of the transfer function: Some portions of the cosine curve (dashed line) collapse due to the visibility function.

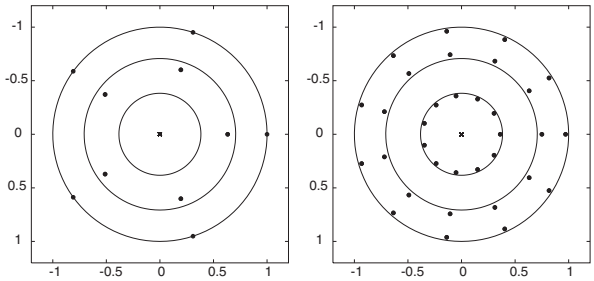

Figure 2. Sampling directions: $B = 3$ (left) and $B = 6$ (right). Circles in the figures correspond to $\theta = \pi/8$, $\pi/4$, and $\pi/2$ respectively.

Second, we consider this issue based on sampling theorem on a sphere[3] [4]. Let us assume that an illumination distribution is represented by a band-limited function with bandwidth B: $L_{lm} = 0$ $(l \geq B)$. Then, the illumination distribution can be reconstructed from the illumination radiance sampled at a finite number of directions. In particular, we can reconstruct the illumination distribution from only $B \times (2B - 1)$ sampling directions as

$$L_{lm} = \sum_{i=0}^{B-1} \sum_{j=0}^{2B-2} \frac{2\pi w_i}{2B-1} L(\theta_i, \phi_j) Y_{lm}(\theta_i, \phi_j), \quad (7)$$

based on Gaussian quadratures [15] (see Appendix). Figure 2 shows sampling directions necessary when $B = 3$ and $B = 6$. Here, they are represented by points on a unit sphere and then projected onto a plane corresponding to $\theta = \pi/2$. The distribution of sampling directions is symmetric with respect to the plane.

[3]The sampling theorem has been used for acquiring basis images of an object under varying illumination conditions [18].

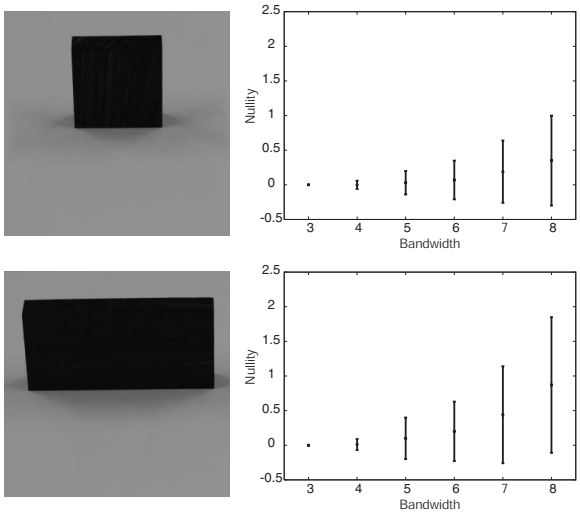

Figure 3. Nullity against bandwidth for real scenes shown in the left panels.

The sampling theorem provides a set of sampling directions necessary for reconstructing an illumination distribution uniquely. Therefore, by verifying whether we can estimate the illumination radiance at the set of sampling directions, we determine the maximum frequency of illumination to which the approach based on cast shadows can estimate correctly. Namely, substituting equation (7) into equation (4), we represent the brightness of the Lambertian surface as

$$E(\boldsymbol{x}) = \sum_{i=0}^{B-1} \sum_{j=0}^{2B-2} L(\theta_i, \phi_j) a_{ij}, \tag{8}$$

and obtain a set of linear equations with respect to $L(\theta_i, \phi_j)$. Here, $a_{ij} = 2\pi w_i/(2B-1) \sum_{l,m} Y_{lm}(\theta_i, \phi_j) T_{lm}(\boldsymbol{x})$. Then, we confirm the number of linearly independent equations. For our current implementation, we used a standard MATLAB function *rank*.

In Figure 3, the right panels show the nullity $(= (\text{the number of unknowns}) - rank)$ against the bandwidth B for real scenes shown in the left panels. Since the distribution of the sampling directions is not uniform over a unit sphere as shown in Figure 2, we randomly rotated a set of sampling directions 300 times and then computed the mean and the standard deviation of the nullity. Here, we set the illumination radiance coming from the back of the plane to 0 by definition. One can find, focusing on the case where the nullity is considerably less than 1, that inverse lighting based on cast shadows can estimate illumination up to $B = 5$ ($B = 4$), that is, $l = 4$ ($l = 3$) for the upper (the lower) example.

3.3. Limitations

Our proposed method using spherical harmonics is effective for recovering illumination that is well localized in the frequency domain and consists of relatively low-frequency components such as a cloudy sky. However, our method has two limitations.

First, we usually observe only a part of the shadows cast by an object in an input image. It often happens that some portions of the view of a camera are occluded by the object and the regions behind the object cannot be seen from the camera. This problem of unobservable regions prevents our method from estimating high-frequency components of illumination as shown in Figure 3. The fewer regions we observe, the larger nullity we obtain.

Second, natural illumination usually contains high-frequency components such as a point light source or an area light source with a sharp edge. It is certain that spherical harmonics can efficiently represent illumination distribution localized in the frequency domain. On the other hand, a large number of basis functions are required for representing illumination localized in the angular domain. Moreover, the assumption in Section 3.2 that an illumination distribution is represented by a band-limited function does not hold. This problem of high-frequency components often causes aliasing in the estimated illumination distribution.

4. Inverse Lighting Using Haar Wavelets

In order to cope with the limitations described in the previous section, we propose an efficient method using Haar wavelets.

4.1. Motivation

The use of Haar basis functions is motivated by the following three aspects.

- **Compact Supports**
  The basis functions have compact supports with various sizes. Here, the support of a function is the region over which the function is nonzero. This enables us to represent illumination of each incident direction at different resolutions and alleviate the problem of unobservable regions.

- **Sparsity**
  In general, many of wavelet coefficients of a function are either zero or negligible. In order to efficiently synthesize cast shadows, Ng *et al.* represented illumination distributions by using Haar wavelets [11]. They showed that fewer than 1% of the basis functions are

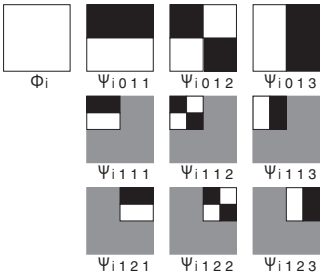

Figure 4. Haar basis functions: The basis functions take 1, 0, and -1 in white, gray, and black regions. Normalization constants are ignored for display.

sufficient to represent natural illumination with high accuracy. This property alleviates the problem of high-frequency components appeared in the method using spherical harmonics.

- **Orthonormality**
  The basis functions are orthonormal. Then, the contribution of a basis function to an illumination distribution in terms of L_2 norm is simply obtained by the square of the corresponding coefficients. This makes our proposed method tractable.

4.2. Proposed Method

By taking advantage of Haar wavelets, we propose an adaptive method that is similar to the oracle [6]. In our method, an illumination distribution is represented by using a cube map. Since we consider the scene where an object is placed on a plane, we omit the bottom of the cube map. Then, we construct two-dimensional nonstandard Haar basis functions [24] on each facet of the cube map. We denote the l-th basis function $(l = 1, 2, 3)$ on the i-th facet $(i = 1, 2, \cdots, 5)$ at the j-th resolution $(j = 0, 1, \cdots, j_{\max})$ with the k-th support $(k = 1, 2, \cdots, 2^j)$ as Ψ_{ijkl}. We also denote the basis function which is constant over the i-th facet as Φ_i. Figure 4 shows the first 10 basis functions on the i-th facet, where we ignore the normalization constants for display.

Thus, the illumination distribution is represented by

$$L(\theta, \phi) = \sum_i \left(c_i \Phi_i(\theta, \phi) + \sum_{j,k,l} d_{ijkl} \Psi_{ijkl}(\theta, \phi) \right), \tag{9}$$

where c_i and d_{ijkl} are coefficients of the corresponding basis functions. By substituting equation (9) into equation (1), we obtain a set of linear equations with respect to these coefficients.

Our proposed method represents an illumination distribution efficiently and accurately by using a small number of basis functions, and estimates these coefficients reliably. In the context of image synthesis under a given illumination distribution, one can compute wavelet coefficients of the illumination distribution and compress it based on L_2 norm in advance [11]. However, in the context of inverse lighting, we cannot know which basis functions are negligible *a priori* by definition. Moreover, we have to take into account the problem of unobservable regions in order to reliably estimate illumination.

Accordingly, we propose the following adaptive method. Our method begins with the constant basis functions Φ_i, and iteratively adds new basis functions in ascending order of the resolution j. Basically, we add three basis functions Ψ_{ijkl} $(l = 1, 2, 3)$ with the same support simultaneously, and compute their coefficients. Then, we discard basis functions if the squares of their coefficients are smaller than the predefined threshold.

In addition, our method takes account the following two aspects. First, we compute the coefficients under the constraint that the estimated illumination distribution is non-negative everywhere in a similar way to the previously proposed method [20]. For our current implementation, we used a MATLAB function *lsqlin*. Because the average of Ψ_{ijkl} over its support is zero, it is obvious that we cannot discard the constant basis functions Φ_i. Second, in the same way as in Section 3.2, we check the number of linearly independent equations before computing the coefficients. If the problem of computing the coefficients becomes ill-conditioned by the addition of new basis functions Ψ_{ijkl} $(l = 1, 2, 3)$, we discard them and do not divide the region corresponding to the k-th support anymore[4]. Namely, we represent illumination of each incident direction up to the resolution to which our method can reliably estimate its radiance.

The sparsity of wavelet coefficients of a function physically results from regions where the function is close to being flat. In this sense, our proposed method is similar to the previous method which adaptively increases sampling directions for representing illumination [20]. The advantage of our method is that we adaptively refine estimation based on mathematical background instead of on heuristic cost function.

5. Experimental Comparisons

In this section, we report the results of the experiments that compared the method based on spherical harmonics

[4] Ng *et al.* proposed a selection scheme of the basis functions based on the energy they produce [11]. What is important for reliably recovering illumination is not the energy but the rank of the matrix they provide.

Figure 5. Light probes: Kitchen (left), campus (center), and galileo (right). We omit the bottom of the cube maps and show low-dynamic range light probes for display.

with the method based on Haar wavelets by using both synthetic and real images.

5.1. Experiments Using Synthetic Images

We compared two proposed methods by using synthetic images so that we could carefully evaluate their performance on the conditions that their assumptions would perfectly hold: known geometry, distant illumination, and no interreflection. For this evaluation, we used synthetic images rendered under the high-dynamic range light probes provided by Debevec [3] and recovered illumination distributions from them. Figure 5 shows the cube maps of the light probes used in our experiments. We considered a scene where a rectangular solid was placed on a plane. The geometry of the scene was the same as that of the real scene described in Section 5.2.

Figure 6 shows ground truth images (left column), images obtained by using the method based on spherical harmonics (center column), and those obtained by using the method based on Haar wavelets (right column). The images in the second and the third rows are synthesized under the illumination distributions in the first row. We recovered illumination distributions from an input image (d) synthesized by using the cube map of kitchen (a). We set the maximum frequency of illumination, to which the method using spherical harmonics estimated, to 5 based on the analysis described in Section 3.2. The method using Haar wavelets estimated illumination up to the resolution of $16 \times 16 \times 5$ pixels by using 102 basis functions. Here, we normalize the illumination distributions (a and c) so that the pixel value at the maximum radiance takes 255. The estimated illumination by using spherical harmonics (b) is normalized so that the pixel value at the maximum radiance on the top facet takes 255 and the pixel with no radiance becomes 128.

One can confirm that the estimated illumination distribution by using spherical harmonics (b) is affected by aliasing due to the truncation of its high-frequency components, and the cast shadows synthesized under the estimated illumination are blurred (e). In contrast, the estimated illumination distribution by using Haar wavelets (c) resembles that

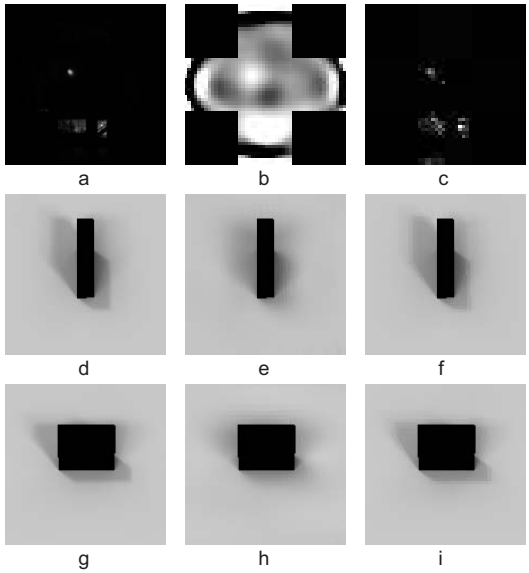

Figure 6. Comparison by using synthetic images: Each column shows ground truth images (left column), images obtained by using the method based on spherical harmonics (center column), and those obtained by using the method based on Haar wavelets (right column) respectively.

Table 1. Root-mean-square errors.

Illumination	Spherical Harmonics	Haar Wavelets
Kitchen	3.7	1.6
Campus	2.7	1.3
Galileo	5.9	2.1
Real	4.8	3.2

shown in the ground truth (a) and explains the cast shadows (f) in the input image (d).

To quantitatively evaluate the performance of our proposed methods, we synthesized images of a different scene under the estimated illumination distributions (h and i), and compared them with the ground truth image (g). In Table 1, we show the Root-Mean-Square (RMS) errors of pixel values. One can find that the RMS errors in the method using Haar wavelets are smaller than those in the method using spherical harmonics. Therefore, it is certain that the method using Haar wavelets works better than the method using spherical harmonics for recovering illumination from cast shadows.

5.2. Experiments Using Real Images

Finally, we compared two proposed methods by using real images. In this experiment, two images of an ebony

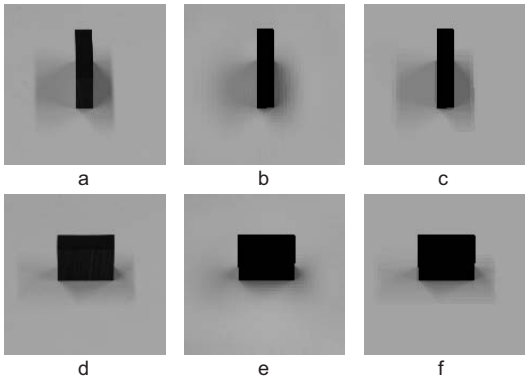

Figure 7. Comparison by using real images: Each column shows real images (left column), images obtained by using the method based on spherical harmonics (center column), and those obtained by using the method based on Haar wavelets (right column) respectively.

block with different poses under natural indoor illumination were taken by a Nikon D1x camera: one image was for the estimation purpose and the other was for the evaluation purpose. Camera calibration of Nikon D1x was conducted by using the camera calibration toolbox for MATLAB developed by Bouguet [2].

Figure 7 shows the real images for estimation (a) and evaluation (d), images obtained by using the method based on spherical harmonics (center column), and images obtained by using the method based on Haar wavelets (right column). As the results of the experiments using synthetic images, the method based on spherical harmonics causes some blurring of cast shadows (b and e). On the other hand, the method based on Haar wavelets successfully reproduces cast shadows (c and f) in the real images (a and d). In addition, Table 1 shows that the RMS error in the method using spherical harmonics becomes larger than that in the method using Haar wavelets. Hence, it can be concluded that the method using Haar wavelets works reliably also for real images.

6. Conclusions and Future Work

In this paper, we provided theoretical insights into why inverse lighting based on cast shadows works in a reliable manner, and discussed what kind of basis functions are appropriate to be used for recovering illumination from cast shadows. In summary, the main contributions of the present study consist of (i) a method using spherical harmonics that is effective for recovering illumination localized in the frequency domain, (ii) theoretical insights as to why inverse lighting based on cast shadows works reliably, (iii) an efficient method using Haar basis functions that alleviates the limitations of the method using spherical harmonics, and (iv) experimental comparisons between our proposed methods.

With regard to wavelet basis functions used in our proposed method, Haar wavelets might not be the optimal choice for representing illumination distibutions. The use of more sophisticated wavelets such as spherical wavelets [21] is one of the future directions of our work. The integration of the approach based on cast shadows with those based on specular reflection components and diffuse reflection components [7] is another future direction. Finally, we note that cast shadows play an important role for recovering Bidirectional Reflectance Distribution Function (BRDF). In a way similar to inverse lighting, the transfer function defined by the product of the illumination radiance and the visibility function generally provides high-frequency components. Namely, cast shadows have the potential for making inverse BRDF stable even under diffuse illumination.

Acknowledgements

A part of this work was supported by Grants-in-Aid for Scientific Research from the Ministry of Education, Culture, Sports, Science and Technology of Japan (No. 13224051).

References

[1] R. Basri and D. Jacobs, "Lambertian reflectance and linear subspaces", *IEEE Trans. PAMI*, 25(2), pp.218–233, 2003.

[2] J.-Y. Bouguet, Camera calibration toolbox for MATLAB, www.vision.caltech.edu/bouguetj.

[3] P. Debevec, "Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography", In *Proc. ACM SIGGRAPH '98*, pp.189–198, 1998.

[4] J. Driscoll and D. Healy, "Computing Fourier transforms and convolutions on the 2-sphere", *Advances in Applied Mathematics*, 15, pp.202–250, 1994.

[5] A. Georghiades, P. Belhumeur, and D. Kriegman, "From few to many: illumination cone models for face recognition under variable lighting and pose", *IEEE Trans. PAMI*, 23(6), pp.643–660, 2001.

[6] S. Gortler and M. Cohen, "Hierarchical and variational geometric modeling with wavelets", In *Proc. ACM Symposium on Interactive 3D Graphics*, pp.35–42, 1995.

[7] Y. Li, S. Lin, H. Lu, and H.-Y. Shum, "Multiple-cue illumination estimation in textured scenes", In *Proc. IEEE ICCV 2003*, pp.1366–1373, 2003.

[8] A. Marschner and D. Greenberg, "Inverse lighting for photography", In *Fifth Color Imaging Conference*, pp.262–265, 1997.

[9] G. Miller and C. Hoffman, "Illumination and reflection maps: simulated objects in simulated and real environments", In *ACM SIGGRAPH '84 course notes*, 1984.

[10] K. Nishino, Z. Zhang, and K. Ikeuchi, "Determining reflectance parameters and illumination distribution from sparse set of images for view-dependent image synthesis" In *Proc. IEEE ICCV 2001*, pp.599–606, 2001.

[11] R. Ng, R. Ramamoorthi, and P. Hanrahan, "All-frequency shadows using non-linear wavelet lighting approximation" In *Proc. ACM SIGGRAPH 2003*, pp.376–381, 2003.

[12] T. Okabe, I. Sato, Y. Sato, and K. Ikeuchi, "Illumination estimation from cast shadows: analysis based on spherical harmonics expansions", *IPSJ SIG Notes*, CVIM 2002-133-27, pp.201–208, 2002 (in Japanese).

[13] T. Okabe and Y. Sato, "Object recognition based on photometric alignment using RANSAC", In *Proc. IEEE CVPR 2003*, pp.I-221–228, 2003.

[14] T. Okabe and Y. Sato, "Support vector machines for object recognition under varying illumination conditions", In *Proc. ACCV 2004*, pp.724–729, 2004.

[15] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C*, Cambridge University Press, 1992.

[16] R. Ramamoorthi and P. Hanrahan, "On the relationship between radiance and irradiance: determining the illumination from images of a convex Lambertian object", *J. Opt. Soc. Am. A*, 18(10), pp.2448–2459, 2001.

[17] R. Ramamoorthi and P. Hanrahan, "A signal-processing framework for inverse rendering", In *Proc. ACM SIGGRAPH 2001*, pp.117–128, 2001.

[18] I. Sato, T. Okabe, Y. Sato, and K. Ikeuchi, "Appearance sampling for obtaining a set of basis images for variable illumination", In *Proc. IEEE ICCV 2003*, pp.800–807, 2003.

[19] I. Sato, Y. Sato, and K. Ikeuchi, "Acquiring a radiance distribution to superimpose virtual objects onto a real scene", *IEEE Trans. VCG*, 5(1), pp.1–12, 1999.

[20] I. Sato, Y. Sato, and K. Ikeuchi, "Illumination from shadows", *IEEE Trans. PAMI*, 25(3), pp.290–300, 2003.

[21] P. Schroder and W. Sweldens, "Spherical wavelets: efficiently representing functions on the sphere", In *Proc. ACM SIGGRAPH '95*, pp.161–172, 1995.

[22] A. Shashua, "On photometric issues in 3D visual recognition from a single 2D image", *Int'l. J. Computer Vision*, 21(1/2), pp.99–122, 1997.

[23] P. Sloan, J. Kautz, and J. Snyder, "Precomputed radiance transfer for real-time rendering in dynamic, low frequency lighting environments", In *Proc. ACM SIGGRAPH 2002*, pp.527–536, 2002.

[24] E. Stollnitz, T. Derose, and D. Salesin, *Wavelets for computer graphics*, Morgan Kaufmann Publishers, 1996.

[25] Y. Wang and D. Samaras, "Estimation of multiple illuminants from a single image of arbitrary known geometry", In *Proc. ECCV 2002*, LNCS 2352, pp.272–288, 2002.

Appendix

The spherical harmonics coefficients of an illumination distribution $L(\theta, \phi)$ are given by

$$L_{lm} = \int_0^{2\pi} \int_0^{\pi} L(\theta, \phi) Y_{lm}^*(\theta, \phi) \sin\theta d\theta d\phi. \quad (10)$$

Assuming that $L(\theta, \phi)$ is a band-limited function with bandwidth B, it is represented as

$$L(\theta, \phi) = \sum_{l'=0}^{B-1} \sum_{m'=-l'}^{l'} L_{l'm'} Y_{l'm'}(\theta, \phi). \quad (11)$$

Substituting equation (11) into equation (10), we can derive

$$\begin{aligned} L_{lm} &= N_{lm} \sum_{l'=0}^{B-1} \sum_{m'=-l'}^{l'} N_{l'm'} L_{l'm'} \\ &\times \left(\int_0^{2\pi} e^{im'\phi} e^{-im\phi} d\phi \right) \\ &\times \left(\int_0^{\pi} P_{l'}^{m'}(\cos\theta) P_l^m(\cos\theta) \sin\theta d\theta \right), \end{aligned} \quad (12)$$

where we use equation (2). In the integral with respect to θ, we change the variable to $x = \cos\theta$. Thus, the integrand is of degree $(2B-2)$ at most with respect to x. Therefore, $P_{l'}^{m'}(x) P_l^m(x) = P_B(x) q_{B-2}(x) + r_{B-1}(x)$, where P_B is the Legendre polynomial and q_{B-2} and r_{B-1} are functions of degrees $(B-2)$ and $(B-1)$ respectively. Because P_B is orthogonal to any function of degree less than B, we can derive

$$\int \cdots d\theta = \int_{-1}^{1} P_{l'}^{m'}(x) P_l^m(x) dx = \sum_{i=0}^{B-1} w_i P_{l'}^{m'}(x_i) P_l^m(x_i), \quad (13)$$

where $P_B(x_i) = P_B(\cos\theta_i) = 0$ and w_i are weight coefficients. We can easily transform the integral with respect to ϕ as

$$\int \cdots d\phi = \frac{2\pi}{2B-1} \sum_{j=0}^{2B-2} e^{im'\phi_j} e^{-im\phi_j}, \quad (14)$$

where $\phi_j = 2\pi j/(2B-1)$. Hence, we can uniquely reconstruct the illumination distribution $L(\theta, \phi)$ from the illumination radiance sampled at the set of directions (θ_i, ϕ_j) as

$$L_{lm} = \sum_{i=0}^{B-1} \sum_{j=0}^{2B-2} \frac{2\pi w_i}{2B-1} L(\theta_i, \phi_j) Y_{lm}(\theta_i, \phi_j). \quad (15)$$