

# Noise-Resistant Fitting for Spherical Harmonics

Ping-Man Lam, Chi-Sing Leung, and Tien-Tsin Wong

**Abstract**—Spherical harmonic (SH) basis functions have been widely used for representing spherical functions in modeling various illumination properties. They can compactly represent low-frequency spherical functions. However, when the unconstrained least square method is used for estimating the SH coefficients of a hemispherical function, the magnitude of these SH coefficients could be very large. Hence, the rendering result is very sensitive to quantization noise (introduced by modern texture compression like S3TC, IEEE half float data type on GPU, or other lossy compression methods) in these SH coefficients. Our experiments show that as the precision of SH coefficients is reduced, the rendered images may exhibit annoying visual artifacts. To reduce the noise sensitivity of the SH coefficients, this paper first discusses how the magnitude of SH coefficients affects the rendering result when there is quantization noise. Then, two fast fitting methods for estimating the noise-resistant SH coefficients are proposed. They can effectively control the magnitude of the estimated SH coefficients, and hence suppress the rendering artifacts. Both statistical and visual results confirm our theory.

**Index Terms:** Spherical harmonics, BRDF, image-based relighting, pre-computed radiance transfer, constrained least square, texture compression, noise-resistant fitting.

## I. INTRODUCTION

Spherical harmonics (SH) [1], [2] [3] are a set of orthonormal basis functions defined over a unit sphere. They have been widely used in representing various spherical illumination properties including bidirectional reflectance distribution function (BRDF) [4], [5], [6], [7], apparent BRDF (ABRDF) [8], [9], distant environment [4], and pre-computed radiance transfer (PRT) [10], [11] [12]. In these applications, we may have to store many SH coefficients.

Further compression on these SH coefficients must be considered in practical applications due to the enormous

Manuscript received xxxx; revised xxxx. This work is supported by a research grant from City University of Hong Kong under Project 7001703.

C.-S. Leung (e-mail: eeleungc@cityu.edu.hk) and P.-M. Lam are with the Department of Electronic Engineering, City University of Hong Kong, Kowloon, Hong Kong

T.-T. Wong (e-mail: ttwong@cse.cuhk.edu.hk) is with the Department of Computer Science & Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong

data volume [11] [13] [14]. These compression methods (e.g. texture compression like S3TC and IEEE half float data type supported by GPU) are usually lossy and introduce noise to the SH coefficients. If care is not taken when estimating the SH coefficients, a small percentage error introduced by these compression processes may induce serious artifacts to the ultimate rendering (Fig. 1). With the increasing interest in the SH basis for practical applications<sup>1</sup>, it is crucial to study the estimation methods of the noise-resistant SH coefficients.

As mentioned in [10] [11] [12], the reflectance properties in many graphics applications are usually defined or collected over the upper hemisphere only. In other words, they are hemispherical functions. To approximate a hemispherical function with the (full-sphere) SH basis functions, the lower hemisphere can be firstly filled with zeros [10] or mirrored values [16], and then the conventional SH projection (integration over the whole sphere) can be applied to estimate the SH coefficients. However, the approximation error is usually very large as pointed out in [11][12].

Instead of using the full-sphere SH basis to approximate a hemispherical function, there are some works to search for other orthonormal bases defined over a hemisphere. In [17], Gautron *et al.* proposed a novel hemispherical orthonormal basis based on shifted polynomials. In [18], the Gram-Schmidt procedure is applied to the SH basis on the upper hemisphere to create a new set of orthonormal basis functions. The basis is designed to construct the principal components for a convex Lambertian object under a varying distant illumination. In [19], an orthonormal basis is used to analyze the COBE DMR mission data in astronomy. This orthonormal construction can be applied to handle data scattered over part of a sphere. However, it would be more desirable to maintain the computation in the SH basis due to the wide availability of existing SH-based literatures and tools.

In [11], Sloan *et al.* proposed to approximate the hemispherical function by the unconstrained least square (ULS) fitting with the SH basis but values on the lower hemisphere are ignored. Mathematically, this ap-

<sup>1</sup>For example, the SH-based PRT routines are available in the DirectX development toolkit [15].




Fig. 1. Visual artifacts introduced by further compression on SH coefficients. SH coefficients in (a) & (b) are obtained by unconstrained least square method while SH coefficients in (c) are obtained by the proposed noise-resistant method. (a) Results rendered from the original 32-bit floating point SH coefficients. (b) Results after compressing these SH coefficients to 4 bits per coefficient with wavelet-based coding. (c) The same wavelet-based coding is applied to the noise-resistant SH coefficients.

proach defines the hemispherical SH basis with upper hemisphere having the same shape as the SH basis and the lower hemisphere being zeroed. Although this hemispherical SH basis is no longer orthonormal, it does not harm the efficiency of computing illumination by performing dot product in the SH domain. This is simply resolved by pre-multiplying an autocorrelation matrix [11].

Unfortunately, the magnitude of SH coefficients estimated from hemispherical samples, using the conventional ULS, is usually very large. The implication of these large-magnitude coefficients is that the rendering is very *noise sensitive*. If no further compression is applied to the SH coefficients, the rendering (reconstruction) can be achieved perfectly. However, further compression is unavoidable due to the enormous size. Since most effective compression methods, such as transform coding, vector quantization, and principal component analysis (PCA) [20], are lossy, they unavoidably introduce noise to the SH coefficients and hence to the rendering results. In this paper, we prove that if the magnitude of SH coefficients is large, a small percentage error in SH coefficients will induce annoying and unexpected visual artifacts to the rendered images. Fig. 1 shows the rendering results from a SH-projected test scene. In this example, the first 49 SH coefficients are used for representing the reflection at each vertex. The SH coefficients in Fig. 1(a) and (b) are estimated by ULS. A pleasant result is obtained (Fig. 1(a)) if these SH coefficients are not compressed and stored as 32-bit floating point per coefficient. When they are compressed to 4 bits per coefficient by wavelet-based coding [21] [22], serious visual artifact is apparent (Fig. 1(b)).

To our knowledge, such a relationship between the noise sensitivity and the magnitude of the SH coeffi-

cients has not been investigated before. Our analysis suggests that the rendering error is proportional to the magnitude of the SH coefficients. Using the conventional ULS, the magnitude of the SH coefficients could be very large. We shall explain the underlying reason. To control the magnitude, we first define a constraint on the magnitude of the SH coefficients. We then propose two fast *constrained least square methods for estimating the noise-resistant SH coefficients* and verify them in various graphics applications. Experiments support our theoretical finding (Fig. 1(c)).

The rest of the paper is organized as follows. Section II reviews the spherical harmonics and its graphics applications. Section III describes the configuration of experiments used throughout the paper. Section IV analyzes the effect of quantization noise in SH coefficients on the rendering. Section V proposes two noise-resistant fitting methods. Statistical and visual results are presented in Section VI to evaluate our methods. Finally, conclusions are drawn and future directions are discussed in Section VII.

## II. SPHERICAL HARMONICS IN GRAPHICS

### A. Mathematics

Spherical harmonics [2] are a set of orthonormal basis functions defined over a unit sphere  $\Omega$ . A spherical function  $f(\vec{s})$  can be approximated by linearly combining a finite number of SH basis functions, given by

$$f(\vec{s}) \approx \sum_{i=1}^n c_i y_i(\vec{s}), \quad (1)$$

where  $\vec{s}$  is a unit vector in 3D,  $y_i(\vec{s})$ 's are the SH basis functions,  $c_i$ 's are the SH coefficients and  $n$  is the number of SH basis functions used. The SH coefficients  $\{c_1, \dots, c_n\}$  form a SH vector to be stored.

## B. Spherical Harmonics in Rendering

The SH expansion is commonly used for representing BRDF, ABRDF, PRT and other spherical reflectance properties due to its compactness [4] [5] [8] [23] [9] [13] [7] [10] [11] [12] [6]. With the SH representation, we can render a scene/object under complex lighting configurations directly from the SH coefficients.

We denote  $L_{p,\text{out}}(\vec{v})$  as the reflected radiance from a non-emissive surface element  $p$ , where  $\vec{v}$  is the viewing direction. In this paper, we use the local coordinate frame [7] [10] [11] [12] to describe the lighting property. That means the surface normal of a surface element is equal to  $(n_{p,x}, n_{p,y}, n_{p,z}) = (0, 0, 1)$ . If the inter-reflection is not considered, the reflected radiance under a distant environment light source is given by [24],

$$L_{p,\text{out}}(\vec{v}) = \int_{\vec{s} \in \Omega_h} L_{p,\text{in}}(\vec{s}) \rho(\vec{s}, \vec{v}) \max(0, s_z) V_p(\vec{s}) d\vec{s} \quad (2)$$

where the unit vector  $\vec{s}$  denotes the light vector defined over the upper hemisphere  $\Omega_h$ ;  $s_z$  denotes the  $z$  coordinate of the light direction  $\vec{s}$ ;  $L_{p,\text{in}}(\vec{s})$  is the incident radiance from the surrounding environment;  $\rho(\vec{s}, \vec{v})$  is the BRDF,  $\rho(\vec{s}, \vec{v})$  and  $\max(0, s_z)$  can be collapsed into the BRDF product function  $\rho^*(\vec{s}, \vec{v}) = \rho(\vec{s}, \vec{v}) \max(0, s_z)$ ; and  $V_p(\vec{s})$  is the visibility. Note that each element has its own visibility function. Since the local frame is employed, the illumination-contributing hemispherical environment changes from point to point. Hence, the computational bottleneck is usually the rotation of  $L_{p,\text{in}}(\vec{s})$ .

As our objective is not on the sophistication of rendering, but on the noise-resistant fitting of SH, let us first consider a simple scenario when the surface reflectance is diffuse. Our example becomes “shadowed diffuse transfer without inter-reflection effect” [11]. Note that our following discussion on the noise-resistant fitting remains valid for other scenarios, such as glossy surfaces<sup>2</sup>. Now, the reflected radiance is given by

$$L_{p,\text{out}} = \int_{\vec{s} \in \Omega_h} L_{p,\text{in}}(\vec{s}) \rho \max(0, s_z) V_p(\vec{s}) d\vec{s}. \quad (3)$$

Clearly,  $\rho \max(0, s_z) V_p(\vec{s})$  is a hemispherical function.

For the illumination with a distant environment, we can express the environment lighting  $L_{p,\text{in}}(\vec{s})$  by the SH basis over the *full sphere* and  $\rho \max(0, s_z) V_p(\vec{s})$  by the SH basis over the *hemisphere*. The reflected radiance becomes

$$L_{p,\text{out}} = \vec{c}^T \tilde{A} \vec{l} \quad (4)$$

<sup>2</sup>For a glossy surface, we need to doubly SH-project the lighting property both on the lighting domain and the viewing domain. In this situation, the input domain of  $\rho(\vec{s}, \vec{v}) \max(0, s_z) V_p(\vec{s})$  is a cross-product of two hemispheres.




Fig. 2. Various sampling patterns over the hemisphere. (a) Hemispherical uniformly distributed; (b) hemispherical regular grid; and (c) scattered sampling.

where  $\vec{c}$  contains the SH coefficients of the reflectance property and self-shadowing of the element;  $\vec{l}$  contains the SH coefficients of the environment; and  $\tilde{A}$  is the autocorrelation matrix of the SH basis over the hemisphere. By considering the pre-computation or the concept of the dual basis, (4) becomes

$$L_{p,\text{out}} = \vec{c}^T \vec{l} \quad (5)$$

where  $\vec{l} = \tilde{A} \vec{l}$ . Hence, the reflected radiance can be efficiently obtained by calculating the dot product of these two vectors<sup>3</sup> [4].

For the local illumination due to a directional or point light source with intensity of  $L_{p,\text{in}}$ , we need to determine its incident light vector  $\vec{s}$ . The reflected radiance is then computed as

$$L_{p,\text{out}} = L_{p,\text{in}} \rho \max(0, s_z) V_p(\vec{s}). \quad (6)$$

As each surface element may have its own visibility function, it is not desirable to store all visibility functions plainly. A compact solution is to fit  $\rho \max(0, s_z) V_p(\vec{s})$  with the SH basis over the hemisphere.

For glossy and opaque surfaces illuminated under a fixed lighting condition, the reflected radiance is a function of viewing directions, given by  $L_{p,\text{out}}(\vec{v})$ . We can record the values of  $L_{p,\text{out}}(\vec{v})$  as viewed along sampled viewing directions. Meaningful viewing directions are on the upper hemisphere. Hence, we can once again express the reflected radiance  $L_{p,\text{out}}(\vec{v})$  as a weighted sum of SH basis functions over the hemisphere.

In the image-based relighting, each pixel is associated with its own ABRDF [8] [23] [13]. If the viewpoint is fixed, an ABRDF reduces to a spherical function in a global coordinate system, given by  $P(\vec{s})$ . For certain situations such as outdoor natural illumination, meaningful light vector samples are restricted to elevational angle smaller than  $\frac{\pi}{2}$ . In this case, we can approximate an ABRDF as a weighted sum of the SH basis functions over the hemisphere.

To obtain the SH coefficients of a spherical function  $f(\vec{s})$ , we need to sample  $f(\vec{s})$ . In the above and many

<sup>3</sup>The details about the rotation of the environment and the pre-computation of the environment can be found in [7] [10].




Fig. 3. Three tested examples throughout the paper.

other rendering applications, only the upper hemisphere  $\vec{s} \in \Omega_h$  or even scattered samples are available. When we say hemispherical sampling, we refer to either uniform (Fig. 2(a)) or regular (Fig. 2(b)) sampling on hemisphere. When we say the sampling is scattered, the samples may distribute on only a part of the sphere (Fig. 2(c)). For example, the BRDF of a real surface is usually captured over the upper hemisphere [25]. In the image-based approach of reflectance measurement [26], the BRDF is only captured at some parts of the sphere due to the shape of the object and the relative orientation of the camera.

### III. EXPERIMENTAL SETTING

#### A. Rendering Examples

Throughout this paper, we use three examples to illustrate the problem and verify our proposed methods (Fig. 3). The first one is the image-based relighting while the other two examples show reflection on 3D objects. It should be noted that, in all tested examples, the functions being fitted are hemispherical. These examples are mainly for illustration purposes and may be different from actual applications.

In the image-based relighting example, “ring” (Fig. 3(a)), the viewpoint is fixed and reference images are recorded under multiple lighting directions. Since we are interested in the lighting of the shown objects above the ground, only those light vectors on the upper hemisphere are captured in this example. There is a set of  $M$  reference images  $\{f_{x,y}(\vec{s}_1), \dots, f_{x,y}(\vec{s}_M)\}$  captured under lighting directions  $\{\vec{s}_1, \dots, \vec{s}_j, \dots, \vec{s}_M\}$ . Hence, a pixel at spatial position  $(x, y)$  is associated with a hemispherical function of light vector. All pixels share a common coordinate system.

Example “torus” (Fig. 3(b)) shows a 3D diffuse torus with self-shadowing being accounted for. The torus is represented as a triangular mesh. Each vertex  $i$  is associated with a hemispherical function sampled in the lighting dimension.

Example “dama” (Fig. 3(c)) shows a glossy 3D sculpture. The lighting condition is assumed fixed. At each

vertex of the mesh, we record the reflected radiances as viewed from sampled viewing directions. Since the object is opaque and each vertex has its own local coordinate frame, a hemispherical function of viewing direction is stored at each vertex.

#### B. Further Compression

As each vertex/pixel may have its own local illumination properties (such as visibility and BRDF), there can be tens of thousands of SH vectors to store. Further compression is unavoidable. To exploit the data correlation of SH coefficients between adjacent vertices/pixels, the same-order coefficients from different vertices/pixels are grouped to form one data source (or a map if vertices/pixels are organized in grid structure). Hence if the SH vectors are  $n$ -dimensional,  $n$  data sources (or maps) are formed. To demonstrate the vulnerability of the SH coefficients estimated by ULS methods, we compress these SH coefficients by some standard compression methods. Four types of data reduction/compression methods are used throughout this paper, they are IEEE half-float representation, S3TC-like texture compression, uniform quantization, and wavelet-based compression. Other compression techniques, such as vector quantization and clustered PCA (CPCA), would have similar results.

1) *Half-float*: The 16-bit IEEE half-float representation is popular due to its wide support by modern GPUs. In most graphics applications, no significant loss in visual quality is observable, except in representing large-magnitude SH coefficients.

2) *S3TC*: S3TC [27] is designed for compressing color textures in game applications. However, current implementations only support 8-bit integer data. At the time we prepare this paper, the floating-point version of S3TC is not yet available. Since the original SH coefficients are in the 32-bit floating point format, we implemented our own S3TC-like coding method for floating point data. It effectively represents each SH coefficient with 4.667 bits.

3) *Uniform Quantization*: Uniform quantization is a standard compression method used in most compression schemes. In our experiments, SH coefficients are quantized to 8 bits per coefficient. The quantization interval of each of the  $n$  data sources is calculated based on the operational rate distortion function [28].

4) *Wavelets*: Wavelet-based coding receives much attention in graphics research [29] [30] [31] due to its multi-resolution nature and progressiveness. It has also been adopted in the JPEG2000 image coding standard as the core engine [32]. We use the 9/7 wavelet to



decompose each SH coefficient map into dyadic structure with 10 subbands. After that, we apply the standard uniform quantization to these subbands [33].

#### IV. UNCONSTRAINED LEAST SQUARE METHOD AND NOISE MODEL

##### A. Unconstrained Least Square

Before we continue, let us first investigate why the ULS methods lead to significant visual artifacts. Let  $f(\vec{s})$  be the hemispherical function being approximated, where  $\vec{s} \in \Omega_h$ . Given  $M$  sampled radiance values  $\vec{\beta} = [f(\vec{s}_1), \dots, f(\vec{s}_M)]^T$  under incident light directions  $\{\vec{s}_1, \dots, \vec{s}_M\}$ , the SH coefficients can be obtained by the least square solution of the following linear equation

$$\vec{\beta} = \mathbf{Y} \vec{c}, \quad (7)$$

where

$$\mathbf{Y} = \begin{bmatrix} y_1(\vec{s}_1) & \cdots & y_n(\vec{s}_1) \\ \vdots & \ddots & \vdots \\ y_1(\vec{s}_M) & \cdots & y_n(\vec{s}_M) \end{bmatrix}.$$

We can use some ULS methods, such as generalization inverse or recursive least square [34], [35], to obtain the least square solution, given by

$$\vec{c} = A^{-1} \vec{b} \quad (8)$$

where  $A = \mathbf{Y}^T \mathbf{Y}$  is proportional to the autocorrelation matrix of the SH basis functions on the hemisphere and  $\vec{b} = \mathbf{Y}^T \vec{\beta}$  is related to the hemispherical SH projection.

Fig. 4-6 show the annoying artifacts in all examples after further compressing the noise-sensitive SH coefficients. All their SH coefficients are estimated by the ULS. In example “ring”, the artifact is image-wise due to its image-based nature. Even though the image-based scene is only reduced to 16-bit half-float per SH coefficient (considered as high bit-rate) in Fig. 4(b), artifacts are observable near the silhouette of the rings. With lower bit-rates (Fig. 4(c)-(d)), the artifacts become unacceptable. Similar artifacts are found in other examples, “torus” (Fig. 5) and “dama” (Fig. 6), when the SH coefficients are further compressed. It is shown that these SH coefficients are highly sensitive to noise.

##### B. Noise Model

The artifacts come from the compression process. It is interesting to investigate why common compression methods generate serious artifacts. The answer is that quantization or compression noise depends on the data magnitude. In this subsection, we first model the noise mathematically. Next, we explain why, together with




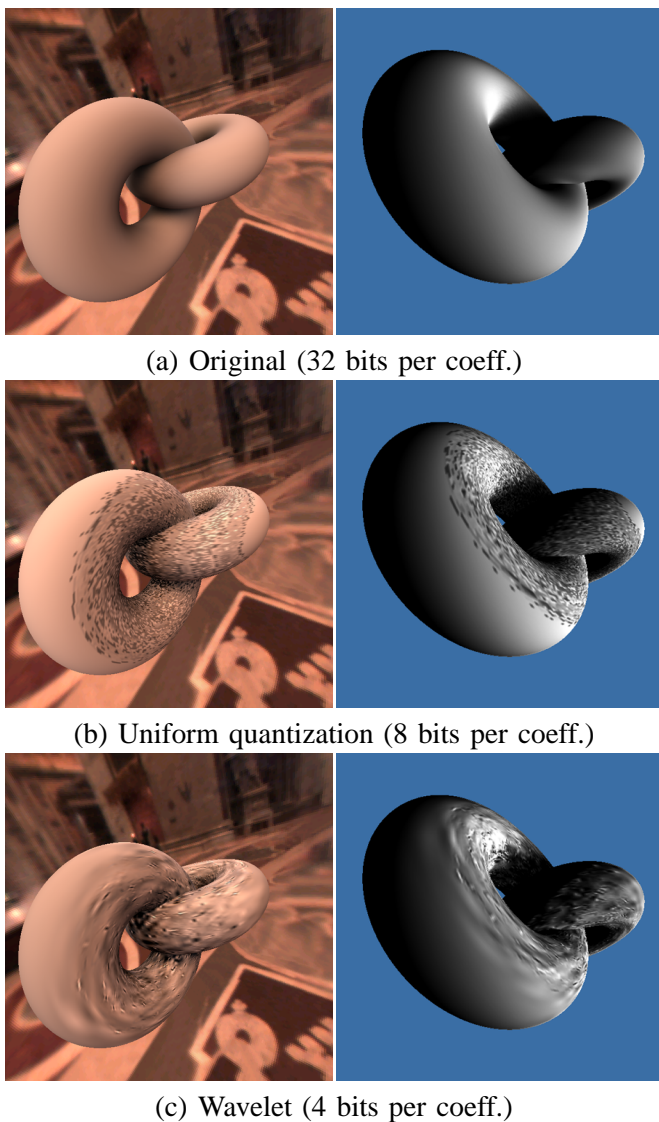
Fig. 4. Relit images of the “ring” example after further compressing SH coefficients estimated by ULS.

the large-magnitude SH coefficients, the noise induces serious artifacts. We model the quantization effect as

$$\hat{c}_i = c_i + \delta_i c_i \quad \forall i \quad (9)$$

where  $\delta_i$ 's are independently identical random variables with zero mean and variance  $\sigma_\delta^2$ . Their density function has the symmetric property. The variance  $\sigma_\delta^2$  can be considered as the percentage error on the SH coefficients. In this model, we assume low-frequency and high-frequency coefficients are quantized with the same bit-rate. With this setting, we can ensure that high-frequency effects (shadow and highlight) are preserved.

In (9), the term  $\delta_i c_i$  represents the quantization noise. We assume that the quantization noise is proportional to the magnitude of the SH coefficients. This assumption is reasonable as it matches several standard quantization schemes. For example, in the low-precision floating point representation, the precision error is proportional to the magnitude of the value. Another example is the uniform quantization with entropy coding. From rate distortion theory [36], the quantization error of a data source under uniform quantization with entropy coding is proportional to the variance of the data source, as well as, the magnitude of the data source. Other compression methods, such as vector quantization and clustered PCA, also have similar behavior.



(a) Original (32 bits per coeff.)

(b) Uniform quantization (8 bits per coeff.)

(c) Wavelet (4 bits per coeff.)

Fig. 5. Rendering results of “torus” after compressing SH coefficients estimated by ULS. The images on the left column show the illumination under a HDR distant environment while the right column shows the scene illuminated by a directional light source.

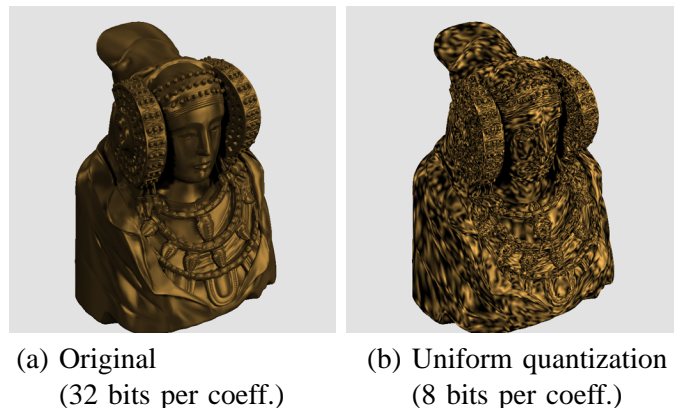
From (9), the average extra error  $E_q$  on the spherical function due to the quantization is given by

$$E_q = \sum_{i'=1}^n \sum_{i=1}^n E \left[ \delta_{i'} \delta_i c_{i'} c_i \int_{\Omega_h} y_{i'}(\vec{s}) y_i(\vec{s}) d\vec{s} \right], \quad (10)$$

where  $E[\cdot]$  is the expectation operator. Since  $\delta_i$ 's are independently identical random variables with zero mean and variance  $\sigma_\delta^2$ , the expectation  $E[\delta_{i'} \delta_i]$  is equal to zero for  $i \neq i'$ . Hence, the average extra error is rewritten as

$$\begin{aligned} E_q &= \sum_{i=1}^n \sigma_\delta^2 c_i^2 \int_{\Omega_h} y_i^2(\vec{s}) d\vec{s} \\ &= \frac{1}{2} \sigma_\delta^2 \sum_{i=1}^n c_i^2. \end{aligned} \quad (11)$$

Note that  $\int_{\Omega_h} y_i^2(\vec{s}) d\vec{s} = \frac{1}{2}$  for uniformly or regularly sampled data. Equation (11) means that the extra error is



(a) Original (32 bits per coeff.)

(b) Uniform quantization (8 bits per coeff.)

Fig. 6. Rendering results of “dama” after compressing SH coefficients estimated by ULS.

proportional to the magnitude of  $c_i$ 's. Hence, controlling the magnitude of the estimated SH coefficients is crucial even though the sampled input data is noise-free.

### C. Magnitude of SH Coefficients

For a linear system in (7), the magnitude of SH coefficients estimated by the ULS depends on the singularity of  $A$ . Considering the eigen value decomposition (SVD) [37] on the matrix  $A$ ,

$$A = USU^T, \quad (12)$$

where  $U$  is a unitary matrix and  $S$  is a diagonal matrix containing the non-negative eigen values in a descending order. As mentioned above, the ULS solution is given by

$$\vec{c} = A^{-1} \vec{b} = US^{-1}U^T \vec{b}. \quad (13)$$

Since  $U$  is a unitary matrix, the norm of  $\vec{c}$  depends on  $S^{-1}$  and  $\vec{b}$ . In the case of hemispherical SH basis (or only part of sphere is covered), the higher order eigen values are very small. That means the norm of  $S^{-1}$  is very large and hence also the magnitude of  $\vec{c}$ .

The problem can be also explained from the view of linear dependence. If some eigen values of  $A$  is very small, the SH basis functions on the hemisphere are nearly linear dependent. That means some unit norm linear combinations of SH basis functions have almost zero values on the hemisphere. These linear combinations can also be easily constructed from the SVD of  $A$ . One of such unit norm linear combinations is given by

$$u_{n,1}y_1(\vec{s}) + u_{n,2}y_2(\vec{s}) + \dots + u_{n,n}y_n(\vec{s}), \quad (14)$$

where  $u_{n,i}$  is the last column of  $U$  or the eigen vector of  $A$  associated with the smallest eigen value. Hence, in the ULS fitting, the coefficients of these combination cannot be well determined and may be assigned with large magnitudes. According to (11), the large magnitude



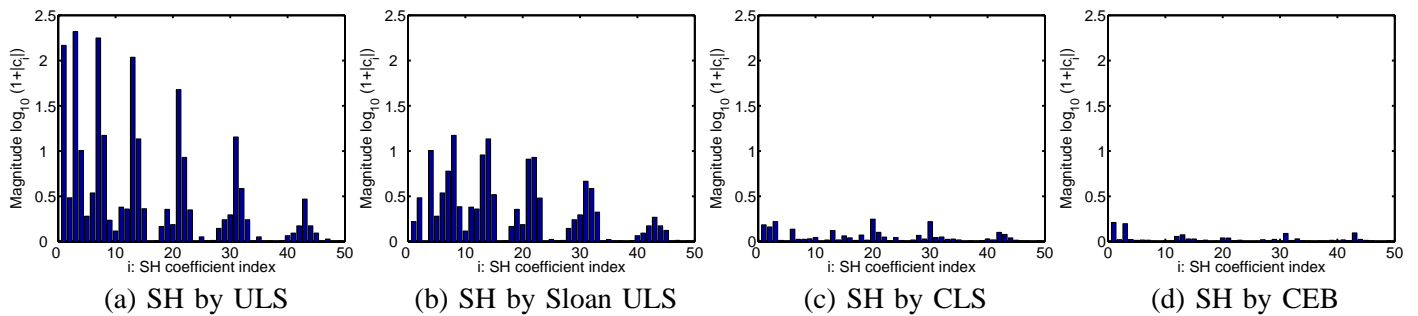


Fig. 7. The magnitude of SH coefficients of a typical pixel in example “ring”.

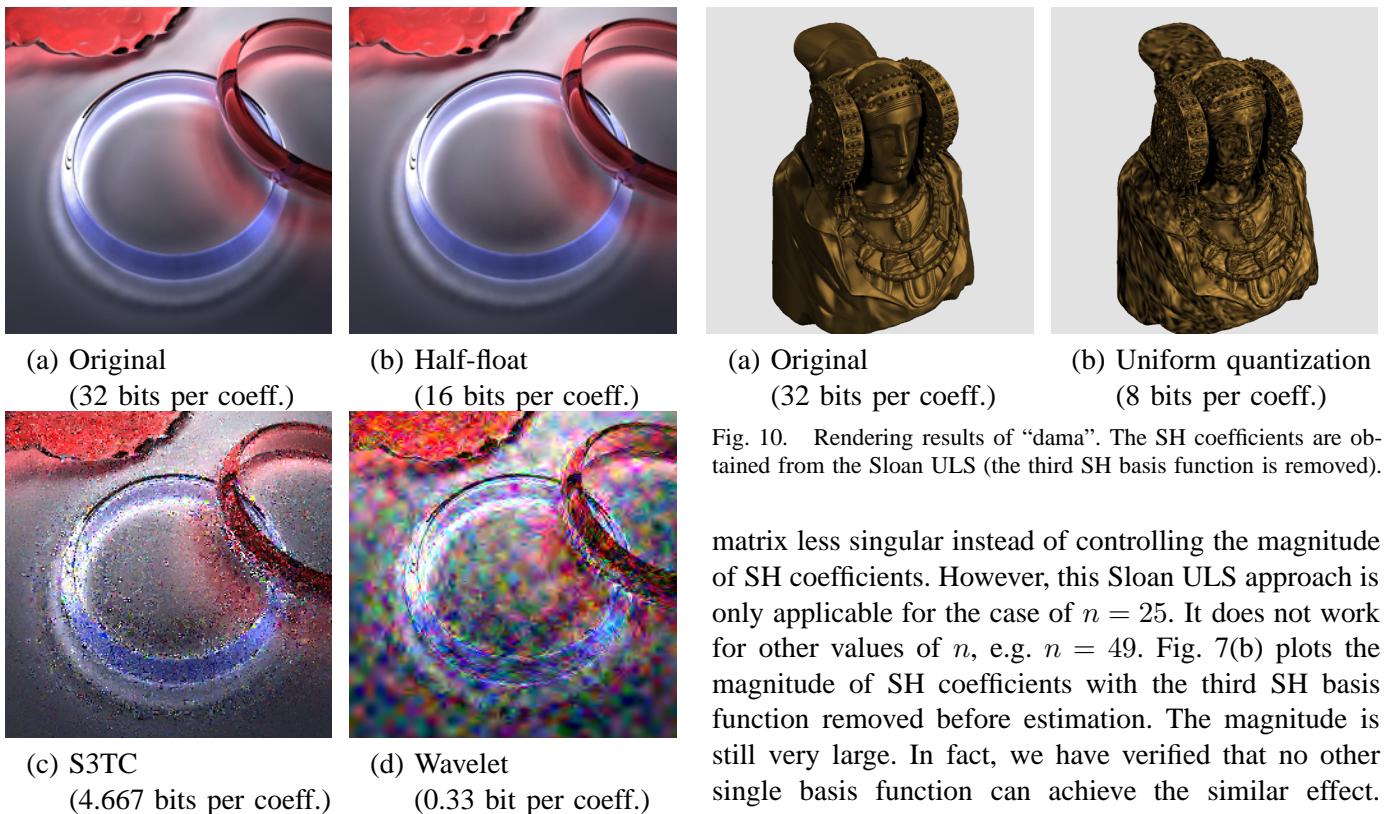


Fig. 10. Rendering results of “dama”. The SH coefficients are obtained from the Sloan ULS (the third SH basis function is removed).

Fig. 8. Relit images of the “ring” example. The SH coefficients are obtained from the Sloan ULS (the third SH basis function is removed).

together with the introduced quantization noise will lead to large distortion in rendering.


Fig. 7(a) plots the magnitude of SH coefficients of a typical pixel in the example “ring” (Fig. 4). It can be observed that the magnitude of SH coefficients is very large which means they are very noise sensitive as implied by (11).

#### D. Sloan Unconstrained Least Square

In [11], Sloan *et al.* observed that the matrix  $A$  is nearly singular. To make the matrix less singular or, in other words, to reduce the condition number of the matrix, they proposed to delete the third SH basis function for  $n = 25$  before estimating the SH coefficients. It should be noted that their motivation is to make the

matrix less singular instead of controlling the magnitude of SH coefficients. However, this Sloan ULS approach is only applicable for the case of  $n = 25$ . It does not work for other values of  $n$ , e.g.  $n = 49$ . Fig. 7(b) plots the magnitude of SH coefficients with the third SH basis function removed before estimation. The magnitude is still very large. In fact, we have verified that no other single basis function can achieve the similar effect. Fig. 8-10 show the rendering results of SH coefficients obtained with Sloan ULS. The visual artifact has been slightly suppressed, especially in the cases of half-float (Fig. 8(b)). However, annoying artifact still exists in other cases. They visually confirm the noise sensitivity of ULS methods including Sloan ULS.

One may argue that we can delete multiple SH basis functions. However, the selection of SH basis functions for deletion is very computationally intensive and data dependent. We have tested this idea for a number of hemispherical functions. Let us consider deleting 2 out of the first 49 basis functions. We need to calculate the condition numbers of  $A$  for all combinations. Among all combinations, there are more than 100 *candidate* combinations with very similar small condition numbers (10 significant digits are the same). Selecting candidates based on such similar condition numbers is risky because their differences in condition number may be due to the numerical errors. Therefore, the condition number is



(a) Original (32 bits per coeff.) (b) Uniform quant. (8 bits per coeff.) (c) Wavelet (4 bits per coeff.)

Fig. 9. Rendering results of “torus”. The SH coefficients are obtained from the Sloan ULS (the third SH basis function is removed).

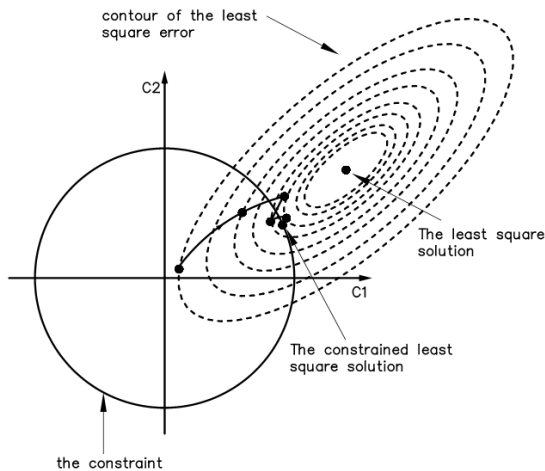


Fig. 11. The illustration of the constraint and the intermediate solutions during searching.

not sufficient to measure the noise sensitivity. Different candidates produce different SH vectors with substantially different magnitudes. For instance, a candidate produces a SH vector with norm of 6.51 while another produces a SH vector with norm of 42. Unless we compute all the SH vectors (from  $\vec{b}$ ) and their norms from these 100 candidates, there is no way to tell which combination is optimal. That means the computation is data dependent. Note that there are hundreds of thousands of hemispherical functions in the rendering applications. Such computation is prohibitively large. If we consider deleting more than two basis functions, this brute-force search becomes impractical even for computing the condition numbers of all combinations. For deleting 4 out of the first 49 basis functions, the number of combinations is more than 200,000.

## V. CONSTRAINED SOLUTIONS

In this section, we propose two efficient methods for estimating the SH coefficients from hemispherical samples. We first impose a constraint on the SH vector. The constraint is related to the energy of the sampled radiance values. We then propose two fast methods to

efficiently control the magnitude of SH coefficients for any value of  $n$ . Finally, we extend our method to general scattered sampling.

### A. The Constraint

In (10), the extra square error can be rewritten as

$$E_q = \frac{1}{2} \sigma_\delta^2 \vec{c}^T \vec{c}. \quad (15)$$

The parameter  $\sigma_\delta^2$  can be considered as the percentage error on the SH coefficients. Now, we impose a constraint that the extra square error  $E_q$  due to the quantization noise should be less than or equal to  $\sigma_\delta^2$  times the energy of the spherical function  $f(\vec{s})$ . That means

$$\frac{1}{2} \sigma_\delta^2 \vec{c}^T \vec{c} \leq \sigma_\delta^2 \times \text{Energy of the function} \quad (16)$$

$$\vec{c}^T \vec{c} \leq E_c, \quad (17)$$

where the energy can be estimated from the sampled values of  $f(\vec{s})$ .

### B. Constrained Least Square

With the constraint (17), we can set up the following constrained least square (CLS) cost function [37]

$$J_2(\vec{c}) = \|\vec{\beta} - \mathbf{Y} \vec{c}\|^2 + \lambda(\vec{c}^T \vec{c} - E_c), \quad (18)$$

where  $\lambda > 0$ . Its constrained solution is given by

$$\vec{c} = (A + \lambda I_{n \times n})^{-1} \vec{b}, \quad (19)$$

where  $I_{n \times n}$  is a  $n \times n$  identity matrix. According to the constraint

$$\vec{c}^T \vec{c} = \vec{b}^T (A + \lambda I_{n \times n})^{-2} \vec{b} \leq E_c, \quad (20)$$

$\lambda$  should satisfy the following inequality

$$\vec{b}^T (A + \lambda I_{n \times n})^{-2} \vec{b} \leq E_c. \quad (21)$$

If we set  $\lambda = 0$ , the least square solution is obtained. If we set  $\lambda \rightarrow \infty$ ,  $\vec{c}$  is a zero vector. As  $\lambda$  increases, the



	PSNR (dB)
ULS	32.21
CLS	31.32

TABLE I

THE AVERAGE PSNR BETWEEN PAIR OF REFERENCE AND RELIT IMAGES.

norm of  $\vec{c}$  decreases monotonically. To understand this property, we take the derivative on  $\vec{c}^T \vec{c}$  with respect to  $\lambda$ , we get

$$\frac{d}{d\lambda}(\vec{c}^T \vec{c}) = -2\vec{b}^T (A + \lambda I_{n \times n})^{-3} \vec{b}. \quad (22)$$

Since  $(A + \lambda I_{n \times n})$  is positive definite when  $\lambda \geq 0$ , the derivative is negative and the square norm of  $\vec{c}$  is a decreasing function of  $\lambda$ . From (18),  $J_2$  is an increasing function of  $\lambda$ . Our goal is to find out the smallest value of  $\lambda$  such that (21) is satisfied. We use an iterative approach to determine  $\lambda$  by first assigning an initial value of  $\lambda$  and then using a simple binary search to estimate a suitable value of  $\lambda$  based on (21).

Fig. 11 illustrates our idea in 2D. We need to search for the smallest  $\lambda$  that falls in the constraint region (enclosed by the solid line). If the ULS method is used, the solution has the smallest error (Table I) but relatively large magnitude (Fig. 7(a)). Table I shows the error in the relighting example “ring” obtained by measuring the peak-signal-noise-ratio (PSNR) between every pair of reference and relit images. When the CLS method is used, the PSNR decreases a little bit only (Table I) but the magnitude of the solution is greatly reduced (Fig. 7(c)). From Fig. 12-14, we can observe that the proposed CLS can effectively suppress the rendering artifact due to the quantization noise.

### C. Speed Up

In (21), we need to calculate a matrix inverse with size  $n \times n$ . If we only need to estimate the SH coefficients of one spherical function, the computation is not so heavy. However, in rendering applications, there are massive number of spherical functions (one spherical function per vertex/pixel). As every spherical function is associated with the same  $A$ , the searching process for  $\lambda$  can be sped up by considering the symmetric property of  $A$ . Consider the SVD on  $A$  is  $A = USU^T$ , where  $U$  is a unitary matrix and  $S$  is a diagonal matrix containing the non-negative eigen values in a descending order. The constrained solution (19) is also a solution of the following equation

$$\vec{\beta} = \mathbf{Y}' \vec{c}', \quad (23)$$

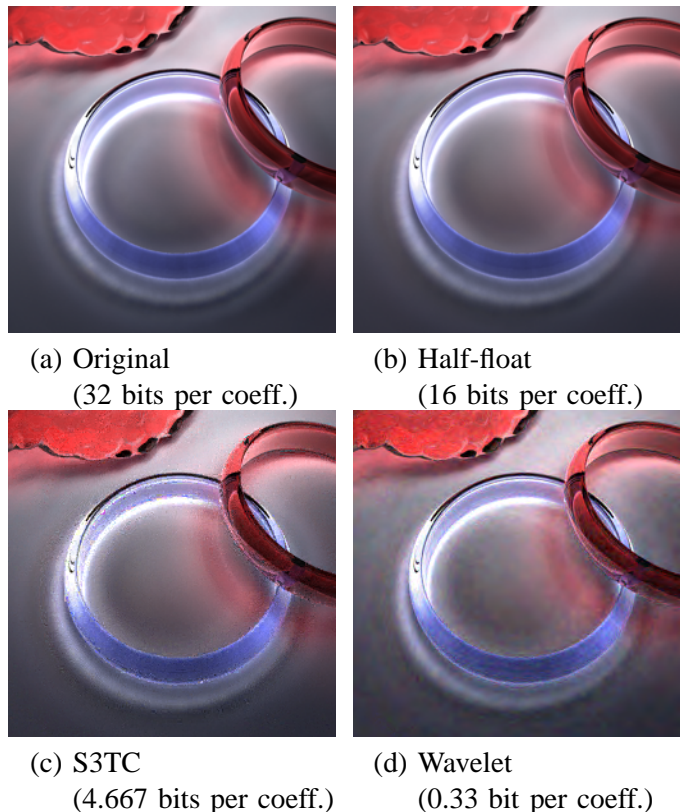


Fig. 12. Relit images of the “ring” example after further compressing the noise-resistant SH coefficients estimated by CLS.

where  $\mathbf{Y}' = \mathbf{Y}U$  and  $\vec{c}' = U^T \vec{c}$ . The solution is given by

$$\vec{c}' = (S + \lambda I_{n \times n})^{-1} \vec{b}' \quad (24)$$


where  $\vec{b}' = \mathbf{Y}'^T \vec{\beta}$ . The constraint (21) can be rewritten as

$$\vec{b}'^T (S + \lambda I_{n \times n})^{-2} \vec{b}' \leq E_c. \quad (25)$$


Note that every spherical function is associated with the same  $A$ , as well as the same  $S$ , which can be pre-computed. The calculation in (25) is simpler than that of (21) because the matrix  $S$  is a diagonal matrix while  $A$  is not. With the above formulation, the complexity of searching is  $O(n)$  for each iteration. The complexity for calculating  $\vec{c}$  from the searched  $\lambda$  is  $O(n^2)$ , which is same as that of the ULS. The complexity of the hemispherical SH projection (the evaluation of  $\vec{b}$ ) is equal to  $O(Mn)$ . Therefore, the total time complexity of the CLS is  $O(k_\lambda n + n^2 + Mn)$ , where  $k_\lambda$  is the total number of iterations in searching for  $\lambda$ . Typical value of  $k_\lambda$  is equal to 32. In the binary search, the initial value is set to 1 and the initial lower bound is set to zero.

### D. Extension to Scattered Sampling


Although our derivation so far assumed the sampling is hemispherical, we can extend to the derivation of scattered sampling and obtain similar equations. For



(a) Original (32 bits per coeff.)



(b) CLS, uniform quantization (8 bits per coeff.)



(c) CLS, wavelet (4 bits per coeff.)

Fig. 13. Rendering results of “torus” after compressing the noise-resistant SH coefficients estimated by CLS.

scattered sampling, sample values cover only part of the sphere  $\Omega'$ . In this case, the integration  $\int_{\Omega'} y_i^2(\vec{s}) d\vec{s} = \kappa_i$  are different for different  $i$ . In this case, the objective function is

$$J'_2(\vec{c}) = \|\vec{\beta} - \mathbf{Y}\vec{c}\|^2 + \lambda(\vec{c}^T \mathbf{D}_\kappa \mathbf{D}_\kappa \vec{c} - E'_c), \quad (26)$$

where


$$\mathbf{D}_\kappa = \begin{bmatrix} \sqrt{\kappa_1} & 0 & 0 & \cdots & 0 \\ 0 & \sqrt{\kappa_2} & 0 & \cdots & 0 \\ \vdots & 0 & \ddots & \vdots & 0 \\ 0 & \cdots & \cdots & \cdots & \sqrt{\kappa_n} \end{bmatrix}$$

and  $E'_c$  is the energy of the function  $f(\vec{s})$ . We can apply a linear transform on  $\vec{c}$  and  $\mathbf{Y}$ , given by

$$\vec{c}^* = \mathbf{D}_\kappa \vec{c} \quad (27)$$

$$\mathbf{Y}^* = \mathbf{Y} \mathbf{D}_\kappa^{-1} \quad (28)$$

$$\mathbf{A}^* = \mathbf{Y}^{*T} \mathbf{Y}^*. \quad (29)$$



(a) Original (32 bits per coeff.)

(b) Uniform quantization (8 bits per coeff.)

Fig. 14. Rendering results of “dama” after compressing the noise-resistant SH coefficients estimated by CLS.

The objective function becomes

$$J'_2(\vec{c}^*) = \|\vec{\beta} - \mathbf{Y}^* \vec{c}^*\|^2 + \lambda(\vec{c}^{*T} \vec{c}^* - E'_c), \quad (30)$$

Now, we can use the method mentioned in Section V-C to solve  $\vec{c}^*$ . After that, we can obtain  $\vec{c} = \mathbf{D}_\kappa^{-1} \vec{c}^*$ .

### E. Constrained Eigen Basis

The work of Sloan *et al.* [11] suggests that we can remove some less important SH basis functions while the constraint is still held. However, as mentioned in Section IV, the complexity of directly removing SH basis functions is too high. Instead of deleting SH basis functions, we can develop a novel and efficient method to delete basis functions in a *transformed domain*. The method is called the constrained eigen basis (CEB). It estimates the constrained SH coefficients tailored for uniform hemispherical sampling. It does not require a search of  $\lambda$  as in the CLS.

Consider a SVD version of ULS, given by

$$\vec{c} = \mathbf{A}^{-1} \vec{b} = \mathbf{U} \mathbf{S}^{-1} \mathbf{U}^T \vec{b}. \quad (31)$$

With the above definition, the function  $f(\vec{s})$  can be approximated as

$$f(\vec{s}) \approx \vec{y}^T(\vec{s}) \vec{c} = \begin{bmatrix} y_1(\vec{s}) \\ \vdots \\ y_n(\vec{s}) \end{bmatrix}^T \mathbf{U} \mathbf{S}^{-1} \mathbf{U}^T \vec{b}. \quad (32)$$

In the above equation, the term

$$\begin{bmatrix} y_1(\vec{s}) \\ \vdots \\ y_n(\vec{s}) \end{bmatrix}^T \mathbf{U} = \begin{bmatrix} e_1(\vec{s}) \\ \vdots \\ e_n(\vec{s}) \end{bmatrix}^T \quad (33)$$

define a new basis, namely the eigen basis, given by

$$\begin{bmatrix} e_1(\vec{s}) \\ \vdots \\ e_n(\vec{s}) \end{bmatrix} = \mathbf{U}^T \begin{bmatrix} y_1(\vec{s}) \\ \vdots \\ y_n(\vec{s}) \end{bmatrix}. \quad (34)$$

The eigen basis is obtained by applying a unitary transform, defined by eigen vectors of  $A$ , on the original SH basis functions. As the matrix  $A$  is proportional to the autocorrelation matrix of the SH basis functions on the hemisphere, the higher order eigen basis functions are less important to reconstruct the SH basis functions. We expect that deleting such higher order eigen basis functions in (32) does not degrade the approximation. When some eigen basis functions are removed, it is identical to estimate the SH vector with less number of elements in  $S^{-1}$ , given by

$$\vec{c} = U' S'^{-1} U'^T \vec{b}. \quad (35)$$

where

$$U' = \begin{bmatrix} u_{1,1} & \cdots & u_{1,k} \\ u_{2,1} & \cdots & u_{2,k} \\ \vdots & \vdots & \vdots \\ u_{n,1} & \cdots & u_{n,k} \end{bmatrix}$$

and

$$S'^{-1} = \begin{bmatrix} s_{1,1}^{-1} & 0 & \cdots & 0 \\ 0 & s_{2,2}^{-1} & \vdots & 0 \\ \vdots & \cdots & \ddots & \vdots \\ 0 & 0 & \cdots & s_{k,k}^{-1} \end{bmatrix}.$$

Since higher order  $s_{i,i}$ 's have the small eigen values, removing them in (31) can improve the magnitude of SH coefficients. In other words, we can apply an ordered construction approach to the eigen basis in order to estimate the SH coefficients based on (35). Since the order of eigen basis functions and the SVD are the same to every spherical function, they can be pre-computed. Moreover, due on the orthonormality of  $U$ , the construction can be done incrementally. Note that the number of removed eigen basis functions can be different for different spherical functions being approximated. By transforming the coefficients of the eigen basis back to the hemispherical SH domain, we can ensure that all spherical functions are represented by the same number of SH coefficients. Such regularity of SH dimension is a desired property for further data compression. For each spherical function, the construction procedure can be summarized by the following pseudocode.

```

Compute the projection  $\vec{b}$ 
Set sum= 0 and  $k = 1$ 
while sum <  $E_c$ 
     $z_k = s_{k,k}^{-1} \vec{u}_k^T \vec{b}$ , where  $\vec{u}_k$  is
        the  $k$ -th column vector of  $U$ .
    sum = sum +  $z_k^2$ 
     $k = k + 1$ 
 $\vec{c} = [\vec{u}_1 \cdots \vec{u}_{k-1}]^T [z_1 \cdots z_{k-1}]$ 

```

In above pseudocode, since vectors  $\vec{u}_k$ 's are an orthonormal basis, the norm of the SH vector is equal to that of the vector  $\vec{z}$ . Hence, we can use a variable sum to accumulate the  $z_k^2$  until the constraint is not satisfied. The accumulator sum is equal to  $\vec{c}^T \vec{c}$  when the first  $k$  eigen basis functions are used. Since all hemispherical functions share the same eigen basis functions  $e_i(\vec{s})$ 's, the complexity of the above method is  $O(n^2)$  only. No searching of  $\lambda$  is required. The total complexity is  $O(n^2 + Mn)$ .

Note that this efficient algorithm only works for hemispherical sampling. For scattered sampling, we need to re-scale SH coefficients by  $\kappa_i$ 's. Variable sum no longer equals the magnitude of SH coefficients and hence cannot be used for checking the constraint. We need to reconstruct the SH coefficient vector  $\vec{c}$ . The complexity of reconstructing  $\vec{c}$  from  $z$  is  $O(n^2)$ . Hence, for scattered sampling, the total complexity of the CEB is  $O(n^3 + Mn)$  and not efficient. Therefore, the CEB should be adopted for hemispherical sampling while the CLS should be adopted for scattered sampling.

Although both CPCA and CEB use the concept of the eigen domain, they actually process the data in the different domains. CPCA works on the SH vectors estimated and hence it cannot control the magnitude of SH vectors. Note that reducing the dimensionality of the estimated SH vectors (based on mean square error) is not equivalent to reducing the magnitude. According to the property of PCA, the sum of the squared first few principal components is close to the sum of the original vector. On the other hand, our CEB works on the hemispherical basis functions. The rendering results of using CEB is quite similar to that of using CLS. Due to the page limit, the rendering results of using CEB are not included in this article. Readers are referred to the webpages (listed at the end of this article) for all visual results.

## VI. COMPARISON AND RESULTS

Table II summarizes the property of various approaches that we have discussed. All approaches have the similar complexity when hemispherical data are considered. The CEB is slow for scattered sampling. The ULS and Sloan ULS (deleting the third SH basis function) are not noise-resistant.

To compare the noise-resistance of four different SH estimation methods (ULS, Sloan ULS, CLS, and CEB), we use the example "ring" (Fig. 3(a)). For each estimation method, the first 49 SH coefficients are estimated for each pixel. The same SH coefficients from different pixels form a SH coefficient map. We then further



	Complexity for Hemispherical Data	Complexity for Scattering Data	Noise Resistance
ULS	$O(n^2 + Mn)$	$O(n^2 + Mn)$	NO
Sloan ULS	$O(n^2 + Mn)$	$O(n^2 + Mn)$	NO
CLS	$O(k_\lambda n + n^2 + Mn)$	$O(k_\lambda n + n^2 + Mn)$	YES
CEB	$O(n^2 + Mn)$	$O(n^3 + Mn)$	YES

TABLE II

COMPARISON OF VARIOUS APPROACHES FOR FITTING SPHERICAL SAMPLES.

	Plain SH coefficients (dB)	S3TC-compressed SH coefficients (dB)
ULS	32.21	15.28
Sloan ULS	31.67	24.31
CLS	31.32	29.73
CEB	30.75	29.79

TABLE III

PSNR OF THE RELIT IMAGES IN "RING". A FLOATING-POINT S3TC METHOD IS USED FOR FURTHER COMPRESSING SH COEFFICIENTS.

compress these SH coefficient maps with S3TC and wavelet-based coding.

The statistical results are shown in Table III and Fig. 15 in the form of peak-signal-to-noise ratio (PSNR). From Table III, if the estimated SH coefficients are not further compressed with S3TC (4.667 bits per coefficient), the PSNR values of different estimation methods are similar. However, when SH coefficients are further compressed, the performance of ULS is drastically reduced. Although Sloan ULS is better than ULS, it cannot achieve the quality as CLS and CEB do. In Fig. 15, we plot the PSNR of various methods against the bit-rate of the wavelet-based compression on the SH coefficients. Even we use a higher bit-rate to compress SH coefficients, say, 2 bits per SH coefficient, PSNR values of ULS are still very low. The low performance of ULS is consistent for various bit-rates. Sloan ULS is better, but still far from satisfactory. On the other hand, PSNR values of CLS and CEB are only slightly affected by quantization noises. They demonstrate the noise-resistance of the estimated SH coefficients.

## VII. CONCLUSIONS AND FUTURE DIRECTIONS

This paper discusses how the magnitude of the SH coefficients affects the rendering quality. By introducing a noise model, we show that the extra error from the quantization noise is proportional to the magnitude of the estimated SH coefficients. We demonstrate that SH coefficients estimated by existing ULS methods are sensitive to the quantization noise. Two effective and fast noise-resistant methods, CLS and CEB, for estimating the SH coefficients are presented to minimize the effect of quantization noise. Statistical and visual comparisons of different SH estimation methods have also been given.




Fig. 15. PSNR of relit images in "ring". A wavelet-based coding is used for further compressing SH coefficients with various bit-rates.

To sum up, the two proposed methods effectively noise-proof the SH coefficients in both image-based re-lighting and 3D rendering applications. Statistics shows that the CLS is slightly better than the CEB. On the other hand, the complexity of using CEB for hemispherical data is much smaller than that of the CLS because CEB does not require the iterative search of constraint parameter  $\lambda$ .

Currently, our methods only handle one constraint. When multiple constraints are required, modification and even redesign of algorithm is needed. This multiple-constraint situation is worth for study in the future as modern GPUs impose various limitations on the input data.

## WEB AVAILABILITY

Due to the page limit, not all visual results can be shown in this article. Readers are referred to the following webpages for a complete visual comparison of various SH estimation methods. Source codes of the proposed methods are also available for download.

<http://www.cse.cuhk.edu.hk/~ttwong/papers/shfit/shfit.html>  
<http://www.ee.cityu.edu.hk/~csleung/>

## ACKNOWLEDGMENT

The authors would like to acknowledge all valuable comments and suggestions made by anonymous reviewers and associate editor, Anselmo Lastra. Thank to Paul Debevec for providing the environment map. The work is supported by a research grant from City University of Hong Kong (Project No. 7001703).

## REFERENCES

- [1] T. MacRobert, *Spherical harmonics: an elementary treatise on harmonic functions with applications*, Dover Publications, 1948.
- [2] R. Courant and D. Hilbert, *Methods of Mathematical Physics*, Interscience Publisher, Inc., New York, 1953.

- [3] K. F. Evans, "The spherical harmonics discrete ordinate method for three-dimensional atmospheric radiative transfer," *Journal of the Atmospheric Sciences*, vol. 55, pp. 429–446, 1998.
- [4] B. Cabral, N. Max, and R. Springmeyer, "Bidirectional reflection functions from surface bump maps," in *SIGGRAPH '87 Proceedings*, July 1987, vol. 21, pp. 273–281.
- [5] F. X. Sillion, J. R. Arvo, S. H. Westin, and D. P. Greenberg, "A global illumination solution for general reflectance distributions," in *SIGGRAPH '91 Proceedings*, July 1991, vol. 25, pp. 187–196.
- [6] R. Ramamoorthi and P. Hanrahan, "Frequency space environment map rendering," *ACM Trans. on Graphics*, vol. 21, pp. 517–526, 2002.
- [7] J. Kautz, P. Sloan, and J. Snyder, "Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics," in *Proc. of the Thirteenth Eurographics Workshop on Rendering*, 2002, pp. 291–297.
- [8] T.-T. Wong, P.-A. Heng, S.-H. Or, and W.-Y. Ng, "Image-based rendering with controllable illumination," in *Proceedings of the 8-th Eurographics Workshop on Rendering (Rendering Techniques'97)*, June 1997, pp. 13–22, Springer-Verlag.
- [9] T. T. Wong, C. W. Fu, P. A. Heng, and C. S. Leung, "The plenoptic illumination function," *IEEE Transactions on Multimedia*, vol. 3, pp. 361–371, 2003.
- [10] P. Sloan, J. Kautz, and J. Snyder, "Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments," *ACM Trans. on Graphics*, pp. 527–536, 2002.
- [11] P. Sloan, J. Hall, J. Hart, and J. Snyder, "Clustered principal components for precomputed radiance transfer," *ACM Trans. on Graphics*, vol. 22, pp. 382–391, 2003.
- [12] P. Sloan, X. Liu, H. Y. Shum, and J. Snyder, "Bi-scale radiance transfer," *ACM Trans. on Graphics*, vol. 22, pp. 370–375, 2003.
- [13] T. T. Wong and C. S. Leung, "Compression of illumination-adjustable images," *IEEE Transactions on CSVT*, vol. 13, pp. 1107–1118, 2003.
- [14] P.-M. Ho, T.-T. Wong, and C.-S. Leung, "Compressing the illumination-adjustable images with principal component analysis," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 15, no. 3, pp. 355–364, March 2005.
- [15] Microsoft DirectX 9.0, [http://msdn.microsoft.com/library/default.asp?url=/library/en-us/directx9\\_c/directx/graphics/programmingguide/advancedtopics/prt/prt.asp](http://msdn.microsoft.com/library/default.asp?url=/library/en-us/directx9_c/directx/graphics/programmingguide/advancedtopics/prt/prt.asp), 2004.
- [16] S. Westin, J. Arvo, and K. Torrance, "Predicting reflectance functions from complex surfaces," in *Computer Graphics (SIGGRAPH '92 Proceedings)*, 1992, pp. 255–264.
- [17] P. Gautron, J. Krivanek, S. Pattanaik, and K. Bouatouch, "A novel hemispherical basis for accurate and efficient rendering," in *EGSR*, 2004.
- [18] R. Ramamoorthi, "Analytic PCA construction for theoretical analysis of lighting variability, including attached shadows, in a single image of a convex lambertian object," in *PAMI*, 2003, pp. 1322–1333.
- [19] K. Gorski, "On determining the spectrum of primordial inhomogeneity from the COBE DMR sky maps I. method," *Astrophysical Journal*, vol. 430, pp. L85–L88 (ApJ Homepage), 1994.
- [20] I. T. Jolliffe, *Principal component analysis*, New York : Springer-Verlag, 2002.
- [21] J. M. Shapiro, "Embedded image coding using zerotree of wavelet coefficients," *IEEE Transactions on Signal Processing*, vol. 41, pp. 3445–3462, 1993.
- [22] A. Said and W. A. Pealman, "A new fast and efficient image codec based on set partition in hierarchical trees," *IEEE Transactions on CSVT*, vol. 6, pp. 243–250, 1996.
- [23] T. T. Wong, C. W. Fu, and P. A. Heng, "Interactive relighting of panoramas," *IEEE Computer Graphics and Applications*, vol. 21, pp. 32–41, 2001.
- [24] J. T. Kajiya, "The rendering equation," in *Computer Graphics (SIGGRAPH '86 Proceedings)*, 1986, pp. 143–150.
- [25] J. F. Murray-Coleman and A. M. Smith, "The automated measurement of BRDFs and their application to luminaire modeling," *Journal of the Illuminating Engineering Society*, pp. 87–99, Winter 1990.
- [26] S. R. Marschner, S. H. Westin, E. P. F. LaFortune, K. E. Torrance, and D. P. Greenberg, "Image-based brdf measurement including human skin," in *Proceedings of 10th Eurographics Workshop on Rendering*, June 1999, pp. 139–152.
- [27] T. A. Moller and E. Haines, *Real-time rendering*, A.K. Peters, 2003.
- [28] N. S. Jayant and P. Noll, *Digital Coding of Waveforms*, Prentice-Hall, Englewood Cliffs, NJ, 1984.
- [29] P. H. Christensen, E. J. Stollnitz, D. H. Salesin, and T. D. DeRose, "Global illumination of glossy environments using wavelets and importance," *ACM Transactions on Graphics*, vol. 15, no. 1, pp. 37–71, January 1996.
- [30] A. Certain, J. Popovic, T. DeRose, T. Duchamp, D. H. Salesin, and W. Stuetzle, "Interactive multiresolution surface viewing," in *Proceedings of SIGGRAPH'96*, August 1996.
- [31] X. Gu, S. Gortler, and H. Hoppe, "Geometry images," in *Proceedings of ACM SIGGRAPH 2002*, 2002, pp. 355–361.
- [32] B. E. Usevitch, "A tutorial on modern lossy wavelet image compression: Foundations of JPEG 2000," *IEEE Signal Processing Magazine*, pp. 22–35, September 2001.
- [33] R. L. Joshi, H. Jafarkhani, J. H. Kasner, T. R. Fischer, N. Farvardin, M. W. Marcellin, and R. H. Bamberger, "Comparison of different methods of classification in subband coding of images," *IEEE Transactions on Image Processing*, vol. 6, no. 11, pp. 1473–1486, Nov. 1997.
- [34] S. Haykin, *Adaptive filter theory*, Prentice Hall, Englewood Cliffs, NJ, 1994.
- [35] C. S. Leung, A. C. Tsoi, and L. W. Chan, "Two regularizers for recursive least squared algorithms in feedforward multilayered neural networks," *IEEE Transactions on Neural Networks*, vol. 12, pp. 1314–1332, 2001.
- [36] A. Gersho and R. M. Gray., *Vector quantization and signal compression*, Kluwer Academic, Boston, 1992.
- [37] G. Golub and C. Van Loan, *Matrix Computations*, Johns Hopkins Univ. Press, MD, 2nd edition, 1990.

PLACE  
PHOTO  
HERE

**Ping-Man Lam** received the B.Eng. degree in electronic from City University of Hong Kong in 2002. He is currently a PhD student in the Department of Electronic Engineering, City University of Hong Kong. His research interest is computer graphics.

PLACE  
PHOTO  
HERE

**Chi-Sing Leung** received the B.Sci. degree in electronics, the M.Phil. degree in information engineering, and the PhD. degree in computer science from the Chinese University of Hong Kong in 1989, 1991, and 1995, respectively. He is currently an Associate Professor in the Department of Electronic Engineering, City University of Hong Kong. His research interests include neural computing, data mining, and computer graphics. In 2005, he received the 2005 IEEE Transactions on Multimedia Prize Paper Award for his paper titled, "The Plenoptic Illumination Function" published in 2002.

PLACE  
PHOTO  
HERE

**Tien-Tsin Wong** received the B.Sci., M.Phil., and PhD. degrees in computer science from the Chinese University of Hong Kong in 1992, 1994, and 1998, respectively. Currently, he is an Associate Professor in the Department of Computer Science & Engineering, Chinese University of Hong Kong. His main research interest is computer graphics, including image-based rendering, natural phenomena modeling, and multimedia data compression. He received *IEEE Transactions on Multimedia Prize Paper Award 2005* and *Young Researcher Award 2004*.