Smooth and non-smooth wavelet basis for capturing and representing light

Cameron Upright, Dana Cobzas and Martin Jagersand
CS, University of Alberta, Canada
{upright,dana,jag}@cs.ualberta.ca

Abstract

Indirectly estimating light sources from scene images and modeling the light distribution is an important, but difficult problem in computer vision. A practical solution is of value both as input to other computer vision algorithms and in graphics rendering. For instance, photometric stereo and shape from shading requires known light. With estimated light such techniques could be applied in everyday environments, outside of controlled lab conditions. Light estimated from images is also helpful in augmented reality in order to consistently relight an artificially introduced object. While algorithms that recover light as individual point light sources work for simple illumination environments, it has been shown that a basis representation achieves better results for complex illumination. In this paper we propose a light model that uses Daubechies wavelets and a method for recovering light from cast shadows and specular highlights in images. We assume that the geometry is known for part of the scene. In everyday images, one can often obtain a CAD model of man-made objects (e.g. a car), but the rest of the scene is unknown. Experimentally, we tested our method for difficult cases of both uniform and textured objects and under complex geometry and light conditions. We evaluate the stability of estimation and quality of scene relighting using our smooth wavelet representation compared to a non-smooth Haar basis and two other popular light representations (a discrete set of infinite light sources and a global spherical harmonics basis). We show good results using the proposed Daubechies basis on both synthetic and real datasets.

1. Introduction

Light representation and recovery is an important problem in computer vision. Vision techniques such as photometric stereo or shape from shading rely on light information. In many situations simplified assumptions about light or object reflectance are made in order to solve the shape problem. These assumptions limit the applicability of the reconstruction techniques to carefully controlled laboratory setups. Therefore light reconstruction (referred as *inverse light*) has significant importance among vision problems.

In this paper we consider the problem of estimating light from images where some 3D shape is known, (i.e. the shape of one object) but not necessarily the whole scene. One advantage of capturing the lighting on an existing object as opposed to introducing a light probe is that we not only recover the scene's lighting but also the reflectance parameters of the object. This allows us to render the scene under the same or different lighting conditions as well as introduce synthetic objects into an augmented reality application. Another advantage of using existing shapes in a scene instead of a light probe is that the recovered light minimizes the error in the current scene, allowing a potentially higher quality rendering for moderate view changes. Furthermore, reflectance probes are not always practical to use, and they cannot be inserted afterwords into existing imaging. By contrast, many images contain some man-made object for which either a CAD model can be obtained, or an identical copy of the object scanned post-hoc to obtain the required shape information, hence a technique that recovers light from known shape has more general applicability.

Inverse light techniques can be grouped in two major categories. Some techniques recover a discrete collection of point light sources[16, 2, 9] or a discrete sample of the illumination hemisphere [14, 9]. Others recover light as projection on a global basis defined on the surface of the illumination hemisphere (usually spherical harmonics basis [12, 1]). The methods from the first category are designed for recovering sharp light effects while the ones from the second category work best only for diffuse (Lambertian) scenes.

To model both diffuse and sharp light effects people have proposed basis that provide local support in both spatial domain (here image dimension) and frequency domain. Wavelets are an example of such a basis. In particular, for sharp effects (specular highlights, sharp shadows), Ng et al. [8] showed that representing the illumination hemisphere with coefficients of a Haar wavelet basis gives a better approximation than spherical harmonics coefficients. Recently Okabe et al. [10] used a Haar wavelet basis for light representation. Hara et al. [5] used another exam-

ple of basis with local support (Mises-Fisher distributions) for representing illumination in specular scenes. Our work proposes a new light representation based on Daubechies wavelets and a general method for estimating the light basis from images. We investigate the advantages of using a smooth basis compared to the non-smooth Haar basis and show the superiority of the smooth basis for both synthetic and real datasets.

While there are several works on inverse light reconstruction methods, we do not know of any study comparing the stability and quality of the different representations. However, it is known that inverse light is in general an ill-conditioned problem. For example, Marschner and Greenberg [7] empirically observed that under diffuse reflectance inverse light tends to be ill-conditioned. For a global basis, Ramamoorthi and Hanrahan theoretically proved that only the low frequency light components can be reconstructed from diffuse scenes [12]. They later showed that for a specular scene, estimating the spherical harmonics light coefficients is well-conditioned only up to an order related to the surface roughness [13]. Okabe et al. [10] showed that the inverse problem becomes well conditioned when using shadows and a spherical harmonics light representation under diffuse light conditions. However no study has been made for a basis with local support, like the wavelet, for both high and low frequency lights as well as different perturbations in the input data.

In this paper we study the stability, efficiency and quality of the reconstruction of a smooth Daubechies and non-smooth Haar wavelet basis and compare them with two other popular choices for the basis (discrete set of infinite light sources and spherical harmonics basis). We provide upper bounds on error propagation using classical Wilkinson condition number numerical analysis[4], and experimentally we give real world practical results for simple and complex scenes, diffuse and specular objects, different illumination conditions, different levels of noise in real and synthetic images.

Our inverse light method uses shadows and specular highlights on textured objects. Sato et al. [14] previously used shadows on diffuse objects for recovering light as a discrete illumination hemisphere. We extend the method to incorporate specular highlights. Other works [15, 6] incorporate multiple cues in the light estimation but without giving insights on which cue most helps the light reconstruction . Here we study the influence of shadows vs. specular highlights on the quality and stability of the reconstruction.

The remainder of this paper is organized as follows. In the next section we define the problem from a theoretical viewpoint. Section 3 presents the wavelet based light model and Section 4 its use in the formulation of the rendering (reflectance) equation. In Section 5, we present the light recovery method, first for one view, uniform albedo and
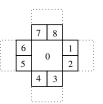

Figure 1. The nine faces of a cubemap representing the illumination hemisphere

then its extension to the multiview, varying albedo case, followed by the system description and implementation details in Section 6 . Finally, Section 7 presents experimental results with both synthetic and real image as well as an extended comparison between the wavelet light models, the spherical harmonics model and the discretized hemisphere model.

2. Problem definition

The light recovery (or the inverse light) problem can be stated as follows: Given a set of calibrated images $I_i, i = 1 \ldots N$ of an object with known geometry and reflectance $S = \{\mathbf{x}_j, j = 1 \ldots M\}$, recover the light that illuminates the scene.

The inverse light problem formulation relies on a chosen light model as well as an image formation (rendering) model. The light is represented as a linear combination of basis functions (here we propose the use of a Daubechies wavelet basis but the formulation is valid for any basis). The image formation model involves choosing a camera projection model that defines how the object geometry is mapped on the image and an object reflectance model that defines how the object appearance is formed given the light conditions and camera position. We chose to use a projective camera model and compare two parametric reflectance models, Lambertian and specular Phong.

It is known that given only one image it is not possible to disambiguate light color from albedo. We first present the inverse light solution for a single view of an object with uniform known albedo. Next we generalize the method for objects with general unknown textured albedo by using multiple images taken by a camera moving around the scene.

3. Wavelet-based light model

We represent the light using a wavelet basis defined over an illumination hemisphere that is represented as a cube map. We chose to use a cube map for its simplicity, and because it is efficiently used in computer graphics for representing environment maps. Each pixel on the cube map represents the direction of a light at infinity. The hemisphere is divided up into nine square faces, a large one on the top and eight small ones on the sides (See Figure 1). We choose to
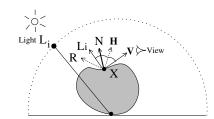
Figure 2. The vectors used in lighting calculations.

ignore the lower part of the cube map assuming that the image shading is mostly caused by light coming from above the object. The light basis is defined over each of these faces. While the model is valid for any type of basis, our proposed model uses a wavelet model (Daubechies). The intensity of a light is found by summing together all basis which influence that specific cube map pixel:

$$l_i = \sum_{k=1}^{K} w_k b_{ki} \tag{1}$$

where l_i is the intensity of light i, b_{ki} is the value of the k'th basis function at location i, and w_k is the corresponding weight for the basis $\mathbf{b}_k$. Rewriting in vector form we get:

$$\mathbf{l} = B\mathbf{w} \tag{2}$$

4. Reflectance model

We next show how the light model is incorporated in the image formation equation. Throughout the paper we assume that the cameras are calibrated and we denote the image projection of a 3D point $\mathbf{x}_j$ by $I(\mathbf{x}_j) = I_j$. The appearance of a point on a surface depends on the material properties of that surface. The Bidirectional Reflectance Distribution Function (BRDF) defines how light interacts with a surface. The simplest reflectance model valid for diffuse scenes is the Lambertian model, which has a constant BRDF k_d. The intensity of a pixel corresponding to the projection of the object point $\mathbf{x}_j$ on an object is given by

$$I(\mathbf{x}_j) = k_{dj} V(\mathbf{x}_j) \sum_{i}^{N} l_i S_i(\mathbf{x}_j) \langle \mathbf{L_i}, \mathbf{N}_j \rangle \omega_i \tag{3}$$

where k_{dj} is the albedo for the point $\mathbf{x}_j$, and $V(\mathbf{x}_j)$ denotes visibility of the object point with respect to the view direction. The discrete lights on the cube map are indexed by i. A binary value $S_i(\mathbf{x}_j)$ represents the shadow and is equal to one when the i'th light is visible from the point $\mathbf{x}_j$, and zero otherwise. The intensity, direction and solid angle of the i'th discrete light is given by l_i, $\mathbf{L_i}$ and ω_i respectively. The normal of the surface at point $\mathbf{x}_j$ is given by $\mathbf{N}_j$ (see Figure 2). $\langle .,. \rangle$ denotes the dot product.
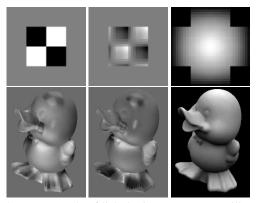

Figure 3. Top: example of light basis, Bottom: scene illuminated with the light basis; From left to right: Haar diagonal basis at the lowest level; corresponding Daubechies; first order spherical harmonic.

To represent general reflectance we chose a dichromatic model that has a Lambertian term for the diffuse part and a specular term given by the Phong model. While not physically based, it works very well in practice and it is one of the most popular reflectance models used for rendering in the computer graphics community. The intensity of a pixel corresponding to $\mathbf{x}_j$ is:

$$I(\mathbf{x}_j) = V(\mathbf{x}_j) \sum_{i}^{N} l_i S_i(\mathbf{x}_j)(k_{dj} \langle \mathbf{L_i}, \mathbf{N}_j \rangle + k_s \langle \mathbf{L_i}, \mathbf{R}_j \rangle^n) \omega_i \tag{4}$$

where k_s is the specular value, $\mathbf{R}_j$ is the reflection of the viewing direction with respect to the normal, and n is the shininess of the material. If the object is made from more than one material, per-point specular parameters have to be specified. A common variant of the Phong model is the Blinn-Phong model. The Blinn-Phong model replaces the $\langle \mathbf{L_i}, \mathbf{R}_j \rangle$ term with $\langle \mathbf{H}_i, \mathbf{N}_j \rangle$, where $\mathbf{H}_i = \frac{\mathbf{L_i} + \mathbf{V}}{|\mathbf{L_i} + \mathbf{V}|}$ is the bisector between light and view directions. The geometry behind the Lambertian and Blinn-Phong models can be seen in Figure 2. Incorporating the wavelet based light model from Equation 1 into Equation 4 and performing a little rearranging we get

$$I(\mathbf{x}_j) = V(\mathbf{x}_j) \sum_{k=1}^{K} w_k \sum_{i=1}^{N} \\ S_i(\mathbf{x}_j) b_{ik}(k_d \langle \mathbf{N}_j, \mathbf{L}_i \rangle + k_s \langle \mathbf{H}_i, \mathbf{N}_j \rangle^n) \omega_i \tag{5}$$

Equation 5 can be interpreted as a sum of "basis images", and can be efficiently implemented on graphics hardware. Examples of these basis images can be seen in Figure 3. The top row displays the basis while the second row an example of a scene illuminated with the basis.

5. Inverse light

5.1. Inverse light from a single image

When given just a single image of a scene we are unable to distinguish between lighting and albedo, only the product of the two can be found. The same is true for a specular object. Due to this limitation, with one image we are only able to reconstruct lighting for objects with uniform albedo. Without loss of generality we can assume that the object is white, with $k_d^\lambda = 1$. We also assume known reflectance, setting $k_s^\lambda = 1$ and n to a fixed value. Equation 5 can be written in a simple linear form:

$$\mathbf{I} = A_s \mathbf{w} \qquad (6)$$

where the components of A_s are:

$$a_{jk} = V(\mathbf{x}_j) \sum_{i=1}^{N} S_i(\mathbf{x}_j) b_{ik} (k_d \langle \mathbf{N}_j, \mathbf{L}_i \rangle + k_s \langle \mathbf{H}_i, \mathbf{N}_j \rangle^n) \omega_i$$

Non-negativity constraints are needed on the l_i variables to prevent negative lighting

$$\forall i \sum_{k=1}^{K} w_k b_{ki} \geq 0$$

or compactly

$$B\mathbf{w} \geq 0 \qquad (7)$$

5.2. Inverse light from multiple images

We now generalize the light and reflectance estimation method for objects with varying albedo. We make the assumption that the lighting doesn't change and that the object remains stationary with respect to the lighting (only the camera moves). In this case, when viewing a fixed point on an object from different camera angles, the difference in the image intensity of that point is due to the specular component in Equation 5.

$$I_j - I_j' = \sum_{k=1}^{K} w_k (c_{jk} - c_{jk}') \qquad (8)$$

$$c_{jk} = V(\mathbf{x}_j) \sum_{i=1}^{N} S_i(\mathbf{x}_j) b_{ki} k_s \langle \mathbf{H}_i, \mathbf{N}_j \rangle^n \omega_i \qquad (9)$$

$$c_{jk}' = V'(\mathbf{x}_j) \sum_{i=1}^{N} S_i(\mathbf{x}_j) b_{ki} k_s \langle \mathbf{H}_i', \mathbf{N}_j \rangle^n \omega_i \qquad (10)$$

When only one point is specular, one of the terms drop (the color difference represents the specular part of the specular pixel). When more than one image is available a similar
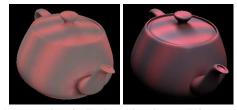

Figure 4. Example of albedo estimation. (left) reconstructed albedo (right) one of the original images

constraint is formed for every pair. The equation has the same form as in the single view case:

$$\hat{\mathbf{I}} = A_m \mathbf{w} \qquad (11)$$

and again we have to impose non-negativity constraints $B\mathbf{w} \geq 0$.

After an initial light is calculated from corresponding specular pixels in pairs of images, per-point albedo is recovered from the diffuse pixels. We then refine the light using the full equations 6,11 with the estimated albedo. Figure 4 shows on the left the reconstructed albedo for the object on the right side.

5.3. Numerical light estimation and regularization

Inverse light is an inherently ill-conditioned problem, especially when using real images and noisy data. To alleviate this problem we used a smoothness regularization on the gradient of the light map:

$$E_s = \lambda_s \sum_i \|\nabla l_i\|^2 \qquad (12)$$

where λ_s is a positive constant that controls the magnitude of the smoothness penalty. By rewriting Equation 12 as a quadratic form and combining it with light basis representation (Equation 2), we get the following :

$$E_s = \lambda_s \mathbf{l}^T Q_{sl} \mathbf{l} = \lambda_s \mathbf{w}^T B^T Q_{sl} B \mathbf{w} = \lambda_s \mathbf{w}^T Q_s \mathbf{w} \quad (13)$$

The light estimation with known reflectance is therefore a linear least squares problem with linear inequality constraints for both the one image with uniform albedo and the multi-image with varying albedo cases (Equations 6,11,13,7). It can be re-formulated as a quadratic programming problem.

$$\min_{\mathbf{w}} \left(\mathbf{w}^T (A^T A + \lambda_s Q_s) \mathbf{w} - 2\mathbf{I}^T A \mathbf{w} \right) \qquad (14)$$

subject to

$$B\mathbf{w} \geq 0 \qquad (15)$$

where A is the light transport matrix, $\mathbf{w}$ is the vector of basis function coefficients and $\mathbf{I}$ is a vector of the observed pixel values.

As shown by Ramamoorthi and Hanrahan [11] in the context of efficient rendering, a good approximation of an environmental map can be achieved by zeroing small wavelet coefficients. Here, in the case of the inverse light problem, the areas of the cubemap that need more detail are initially unknown. We start with a coarse basis that is then locally refined based on projected image residual.

6. System and implementation details

Algorithm 1. Lighting Reconstruction

Require: n calibrated input images $I_1 \ldots I_n$
M points on the object $\mathbf{x}_1 \ldots \mathbf{x}_M$
specular params k_s^λ, n
initial basis representation

1: solve for lighting coefficients $\mathbf{w}_k$ (Equation 6 or 11 for single and multi-view cases)
2: project image residual onto remaining unused basis functions, add the top 75% to the basis representation
3: solve for lighting coefficients $\mathbf{w}_k$ again (Equation 6 or 11 for single and multi-view cases)
4: **if** multi-view reconstruction **then**
5:     estimate albedo
6:     solve for lighting coefficient $\mathbf{w}_k$ one last time, using the reconstructed albedo and Equation 6
7: **end if**

The inverse light algorithm (outlined above) starts with an initial coarse basis representation. In the case of the wavelet basis functions, only the scaling functions are used plus the first level of wavelets on the top cubemap face. We first solve for the lighting coefficients $\mathbf{w}_k$ using Equation 6 or Equation 11 for the single and multiple view cases respectively. Next we calculate the residual between the original images and rendered images that use the current light estimate. For refining the wavelet basis, we project the residual onto the unused basis' of the light transport matrix A and add new basis vectors corresponding to the top 75% of residuals. In the case of the Spherical Harmonic and Discrete basis', no refinement is needed. We then reestimate the lighting coefficients using the new basis (notice that while the wavelets are orthonormal in the lightmap space, orthonormality is not preserved in image space). In the mutiview case we also estimate the albedo by solving Equation 4. This albedo is then used to estimate the lighting coefficients one last time.

7. Experiments

We experimentally compared the proposed smooth Daubechies wavelet light representation with the Haar [10] and spherical harmonics [10, 13] based light maps as well as a simple uniform discrete map defined over the cube map [14]. We performed two types of experiments. First we evaluate the quality of the light reconstruction in real and synthetic images. Next we evaluate the stability of the inverse light method for the different representations in noisy conditions (for both scene and image noise).

7.1. Quality of the reconstruction light basis

To compare the *quality of the reconstructed light map* we generated synthetic images lit with two environmental maps courtesy of P. Debevec [3]. The results are displayed in Figure 5 (left) for the St. Peter's Basilica lightmap and in Figure 5 (right) for Ufizzi lightmap. We chose these two light map as examples of an environment with few sharp lights (St. Peter) and an environment with a large area light (Ufizzi). For all light representations we first tested the single view inverse light method (Section 5.1) on a white teapot model (the specular object is shown in Figure 6, first row). Next we tested the multiview inverse light method using a colored teapot (Figure 6, second row).

The images in Figure 5 show that the spherical harmonics representation tends to blur the whole light map but all other 3 basis show similar results. For St. Peter's light map that has sharp lights Haar works better than the Daubechies while for the Ufizzi lightmap that has one big area light the smooth basis gives a more accurate reconstruction. This is expected since the smooth basis offers better area support than the non-smooth one. In both cases the discrete hemisphere also shows good performance. Among different objects tested, the white specular one (row 2) gives better results. The multiview case (row 3) is more numerically difficult as the estimated non-uniform albedo creates some artifacts in the light map.

Next we tested the *quality of the reconstructed images*. It is known that in the case of lambertian scenes only the low frequency light components are recovered but they are sufficient for reproducing the input images [12]. The result was extended to specular scenes [13] showing that only the spherical harmonics light coefficients up to an order (related to the surface roughness) are well-conditioned. Here in addition to spherical harmonics we also tested the wavelet and discrete hemisphere lights. We chose a view that was not included in the training set used for the light estimation and rendered the image with the recovered light (and albedo for the mutiview case). We performed experiments for the two inverse light methods (single and multi-view case) for both synthetic and real views. For the synthetic case we used the same teapot object with the two environmental maps like in the previous experiment. For the real scene, we used images of two identical glossy ceramic ducks. One was painted white for the one view case and the other left colored for the multiview case. We obtained the geometric model of the duck using a laser scanner and we registered it with the image using corresponding points (manually selected).

| *Haar* | *Daubechies* | *Discrete hemisphere* | *Spherical harmonics* | *Haar* | *Daubechies* | *Discrete hemisphere* | *Spherical harmonics* |

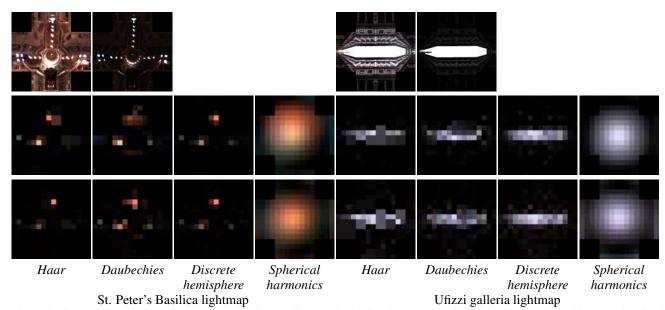St. Peter's Basilica lightmap   Ufizzi galleria lightmap

Figure 5. Lightmaps reconstructed by the proposed inverse light method using 4 example basis: Haar wavelet, Daubechies wavelet, simple discrete cubemap and spherical harmonics. The first row presents the original lightmap and a scaled version that shows the important lights. The second presents the results on the one view specular object with uniform texture and the last row presents the result of the multi-view textured scene. All experiments are performed with the teapot object from Figure 6.

The images were calibrated with respect to the camera using the dotted calibration pattern shown in Figure 6 (3rd row left). We notice that the inverse light is not sensitive to small changes in surface reflectance parameters and therefore we used a heuristic approximation of the reflectance parameters for the real images. A selection of the reconstructed images are shown in Figure 6 while the numerical errors are shown in Table 1. All image errors are normalized on $[0:255]$. Figure 6, first row displays the error images corresponding to the second row in the same figure as well as the second row in Table 1.

From this experiment we noticed that the Daubechies basis gives the best over all performance. The spherical harmonics give very poor perceptual quality (looking at the images in Figure 6) but the measured error is surprisingly better than the Haar or discrete hemisphere case (see Table 1 and first row in Figure 6). This is due to the fact that the spherical harmonics perform better for the diffuse/non-shadowed parts of the object but it tends to blur any specular highlights or shadows that have big influence on the perceptual appearance. The Haar and discrete hemisphere lightmaps give similar results with good perceptual quality on the reconstructed images but poor image error. The Daubechies basis gives the compromise between visual quality and numerical accuracy.
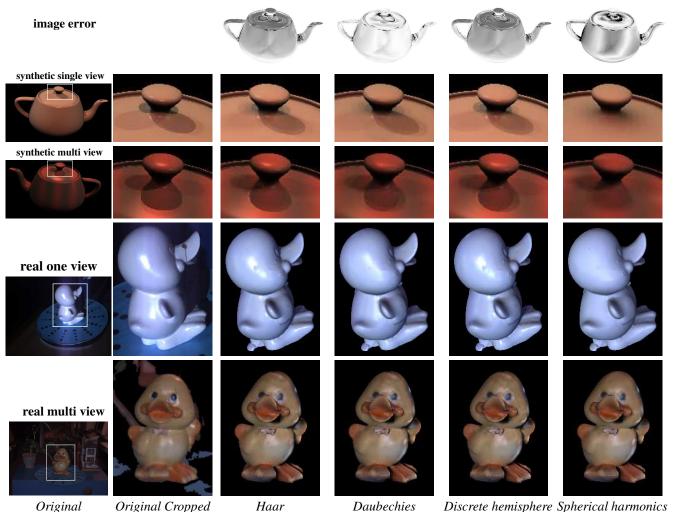
Despite having bigger numerical errors [1] the real scene

Lightmap	Object	Image Error			
		Haar	*DB4*	*SH*	*DH*
	Teapot Diffuse	8.2	2.6	6.4	8.3
St.Peter's	*Teapot Specular*	6.9	2.3	5.6	7.0
	Teapot Multi	3.7	2.0	3.2	3.7
	Teapot Diffuse	9.0	1.6	8.8	9.0
Ufizzi	*Teapot Specular*	9.2	1.8	9.1	9.0
	Teapot Multi	5.6	2.1	3.3	5.7
White Duck Indoor		21.3	21.3	21.4	21.3
Yellow Duck Indoor		25.6	25.6	25.6	25.6

Table 1. Results for quality of light recosntruction for different basis. All image errors are normalized to $[0:255]$. The corresponding recosntructed images are shown in Figure 6

reconstruction was visually quite close to the original with the Daubechies basis giving the best results. This shows that when using real (and thus non-perfect) images the smooth wavelet representation is more robust and spreads the error more evenly on the objects giving a better appearance in the reconstructed view. If one looks carefully at the light map and reconstructed scene from the shiny white duck test (3th row Figure 6), they can notice the three fluorescent lights above the duck and some anomalous blue lights on the horizon. These lights are there to make up for the reflection of the blue calibration pattern, which gives the bottom of the duck a blue appearance. In the muti-view case (yellow duck - last row Figure 6) the results are quite similar.

---

[1] The larger numerical errors for the real images are likely due to difficulty in calibration. It is important for the light estimation to get the rays exactly right from light via reflection to camera. This is a much more difficult problem than camera calibration alone.

Original Original Cropped Haar Daubechies Discrete hemisphere Spherical harmonics

Figure 6. Results with the 4 representations (Haar, Daubechies, discrete hemisphere and spherical harmonics) for synthetic and real images. We tested on a novel view, not included in the light reconstruction. The 2nd and 3rd rows show the single view reconstruction for a white and colored specular teapot and the last two rows are real experiments for a white and colored shinny duck. The corresponding lightmaps for the synthetic case are shown in Figure 5 (left) 3rd, 4th rows. The numerical errors are presented in Table 1. An illustration of the image errors for the single view case is shown in the first row. White represents low errors and black big errors.

7.2. Stability of the inverse light method for different light basis

In a second set of experiments we analyzed the stability of the inverse light reconstruction in the presence of noise. We again compared all 4 basis representations (Haar-H, Daubechies-DB4, discrete hemisphere-DH and spherical harmonics-SH). All tests are performed on the single view method with the white teapot and the St. Peter's Basilica environmental map. We introduced three types of errors: noise in the object geometry (by perturbing the normals), noise in camera calibration and image noise. Finally we compared the contribution of shadow cues vs. specular highlight cues in the reconstruction method. Numerical results are presented in Table 2.

Given a fixed setup, the image intensities are coupled to the light sources by the light transport matrix. An error in calibration or geometry manifests itself as a perturbation of the matrix $\tilde{A}_s = A_s + E$. Likewise an image error/noise can be modeled as $\tilde{I} = I + \Delta I$. Now classical Wilkinson perturbation analysis [4] gives a bound on the calculated light as

$$\frac{\|\Delta \mathbf{w}\|}{\|\mathbf{w}\|} \leq \varkappa(A_s) \left(\frac{\|E\|}{\|A_s\|} + \frac{\|\Delta I\|}{\|I\|} \right)$$

The above formula assumes that the angle between the residual and the solution is small. The condition number $\varkappa$ thus gives an upper bound on how difficult it is to recover light. Typical scenes can give lower errors as shown

Test	Condition numb. $\varkappa$ ($*10^5$)			
	Haar	*DB4*	*SH*	*DH*
Geometry Error	1.2	1.1	2.5	1.2
Calibration error	1.2	0.82	2.5	1.3
Image Noise	1.2	0.83	2.5	1.3
No Shadow cues	3.4	1.1	9.6	6.8
No Specular cues	6.5	7.4	3.1	10.9
No Spec. & No Shadows	11893	5541	13.5	18920

Table 2. Condition number for stability experiments.

in Table 2. The condition number also indicates relative "hardness" of the reconstruction (higher $\varkappa$ = harder).

Among the different bases, the spherical harmonics are the most stable. This can be explained by the fact that they uses global functions that smooth any high frequency perturbations. Additionally we noticed that in practice the Daubechies and Haar are influenced more by the errors in the camera position than the errors in object geometry or image noise. This is due to the fact that camera calibration errors result in large consistent image misalignments compared to the noise due to perturbed normals that is more evenly distributed on the object (we noticed this difficulty in the real image experiments from Figure 6). Looking at the influence of shadows vs. specular reflections we notice that the specular reflectance make the inverse light more stable than just the shadows (except for the spherical harmonic basis). As expected and in accordance to previous results [12, 7] when no shadows and no specular reflections are present the light estimation from only shading becomes ill-conditioned. Only the spherical harmonics are able to reconstruct from diffuse light effects.

It should be noted that the number of basis functions used in our experiments differed between each basis. The discrete basis used the most (the same as the number of lights), Haar and Daubechies used 75% of the available basis functions, while Spherical Harmonics only used the top 9. One would expect that the discrete basis to outperform the others, since there are more degrees of freedom available. While this basis does appear to outperform the others in Figure 5, it gives higher reconstructed image errors in Table 1. Another disadvantage of the discrete basis is that it doesn't scale too well. When using a 16x16 cubemap for example, there are 768 variables to solve for. Even if running time wasn't a problem in this case, the conditioning is so bad that the constrained least square solver fails.

8. Discussion

We have presented a new light representation based on a Daubechies wavelet basis and an inverse light method that uses the proposed representation to recover light from calibrated images of a known object using shadows and spec-

ular highlights. We compared the new representation with three other light bases: spherical harmonics, Haar wavelet and a discrete hemisphere. We shown that the Daubechies basis give the most accurate reconstruction for several types of real and synthetic scenes. We also compared the stability on the four representations in the presence of errors (geometric, camera and image noise). While the spherical harmonics gave the most stable light reconstruction, the Daubechies proved to be comparably stable.

References

[1] R. Basri and Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2), 2003. 1

[2] C. Bouganis and M. Brookes. Multiple light source detection. *PAMI*, 26(4):509–514, April 2004. 1

[3] P. Debevec. Light probe image gallery - http://www.debevec.org/probes/. 5

[4] P. Deuflhard and A. Hohmann. *Numerical Analysis in Modern Scientific Computing: An Introduction*. Springer-Verlag New York, Inc., 2003. 2, 7

[5] K. Hara, K. Nishino, and K. Ikeuchi. Multiple light sources and reflectance property estimation based on a mixture of spherical distributions. In *ICCV*, pages 1627–1634, 2005. 1

[6] Y. Li, S. Lin, H. Lu, and H.-Y. Shum. Multiple-cue illumination estimation in textured scenes. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, page 1366, 2003. 2

[7] S. R. Marschner and D. P. Greenberg. Inverse lighting for photography. In *Proceedings of the Fifth Color Imaging Conference*, 1997. 2, 8

[8] R. Ng, R. Ramamoorthi, and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.*, 22(3):376–381, 2003. 1

[9] K. Nishino, Z. Zhang, and K. Ikeuchi. Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis. In *ICCV*, 2001. 1

[10] T. Okabe, I. Sato, and Y. Sato. Spherical harmonics vs. haar wavelets: basis for recovering illumination from cast shadows. In *CVPR*, pages 50–57, 2004. 1, 2, 5

[11] R. Ramamoorthi and P. Hanrahan. An efficient repres. for irradiance environment maps. In *Siggraph*, 2001. 5

[12] 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), 2001. 1, 2, 5, 8

[13] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *Siggr.*, 2001. 2, 5

[14] I. Sato, Y. Sato, and K. Ikeuchi. Illumination from shadows. *PAMI*, 25(3):290–300, 2003. 1, 2, 5

[15] Y. Wang and D. Samaras. Estimation of multiple directional light sources for synthesis of augmented reality images. *Graph. Models*, 65(4):185–205, 2003. 2

[16] Y. Yang and A. Yuille. Sources from shading. In *CVPR*, pages 534–439, 1991. 1