

Non-negative lighting and Specular Object Recognition

Sameer Shirdhonkar

July 13, 2005

Abstract

Recognition of specular objects is particularly difficult because their appearance is much more sensitive to lighting changes than that of Lambertian objects. We consider an approach in which we use a 3D model to deduce the lighting that best matches the model to the image. In this case, an important constraint is that incident lighting should be non-negative everywhere. In this paper, we propose a new method to enforce this constraint and explore its usefulness in specular object recognition, using the spherical harmonic representation of lighting. The method follows from a novel extension of Szego's eigenvalue distribution theorem to spherical harmonics, and uses semidefinite programming to perform a constrained optimization. The new method is faster as well as more accurate than previous methods. Experiments on both synthetic and real data indicate that the constraint can improve recognition of specular objects by better separating the correct and incorrect models.

Keywords: Non-negative lighting, specular object recognition, Szego eigenvalue distribution theorem, Szego limit theorem, spherical harmonics, semidefinite programming.

1 Introduction

The appearance of an object varies significantly with a change in incident lighting. Model based recognition approaches simulate this change to reduce sensitivity to lighting variations. This approach has been successful for objects with diffuse or Lambertian reflectances. However, recognizing shiny or specular objects is still difficult because their appearance changes dramatically with even a minor change in lighting. We will explain why it is important to enforce a non-negative lighting constraint when solving this problem for specular objects. We will then describe a new exact and fast method to enforce it.

Model based object recognition is performed by comparing the image to an object model. A model includes a structural description (e.g.: regularly sampled surface normals) and an optical description (surface albedo, BRDF, etc. . .). The comparison is an optimization over all possible lighting conditions and produces an image from the model that is as close as possible to the query image. The object whose model produces the closest image is the most likely one to have produced the query image.

Since lighting intensity is a function of direction, and reflected light is a function of the surface normal, both can be represented as functions on the surface of a sphere. Spherical harmonics provide a basis for these functions that is analogous to a Fourier series expansion for 1D functions (e.g. [1]). With this representation, the set of images that an object can produce lie in a linear subspace, with a dimension that depends on the number of harmonics we use. [1] show that only nine harmonics are needed to recognize convex Lambertian objects, because Lambertian reflectance acts as a low-pass filter. However, specular objects reflect higher frequency light (Thornber and Jacobs[15]), so modeling their appearance requires many more harmonics.

Lighting is everywhere non-negative. With this constraint, a model's images form a convex subset of a linear subspace, making the matching problem more complex. When we use a low-dimensional subspace to represent Lambertian objects, ignoring the non-negative lighting constraint is not too serious. However, as the number of harmonics we use grows, the difference between the images produced by non-negative lighting and linear lighting models grows exponentially.

For example, suppose we try to incorrectly match a uniform albedo sphere to an image of a sphere that has a black dot on it. With low frequency harmonics, which suffice for Lambertian objects, we can never approximate this black dot; it has too many high frequencies. With a high-dimensional representation of light, which we need for specular objects, low frequency lighting can produce smooth shading on a sphere, and high frequencies can create a negative specular highlight that darkens the image in a small spot. To prevent this we must ensure that our optimization does not allow negative light (see Figure 1).

Ramamoorthi [11] points out that the non-negativity constraint also helps reduce high frequency noise, since by limiting the value of a function to be non-negative, we indirectly limit the value of its high frequency components.

To enforce non-negative light we want a constraint on the first spherical harmonic coefficients of light that will ensure that the light is non-negative everywhere. The lower order coefficients need not correspond to a non-negative function, but there should exist a way to add higher order harmonics that will make the function non-negative. Looking at the problem more generally, we want to control the range of the function using only the first few coefficients.

In the analogous 1D case, *i.e.* for a Fourier series, an interesting theorem, due to Gabor Szego [4], addresses this problem. It describes a Toeplitz matrix (see section 3.1) of the first few Fourier coefficients whose eigenvalues are contained in the range of the function. Also, the

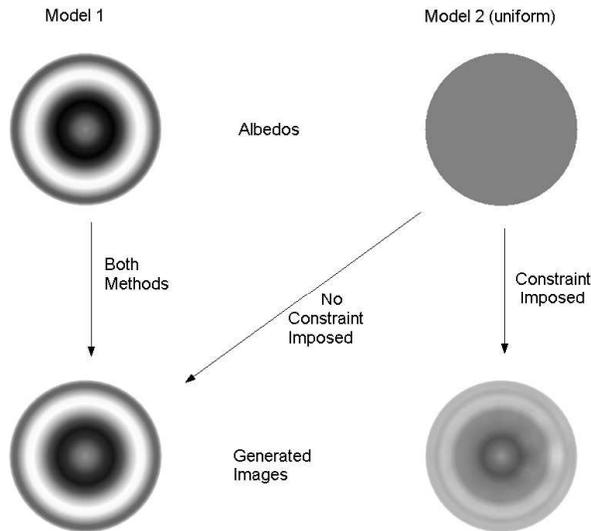


Figure 1: Two different albedo models (both are 4% mirror and 96% Lambertian) and images generated from them while trying to best match the image in the lower left figure. If negative light is allowed, we can get the lower left image from the top right albedo exactly.

Szego Eigenvalue distribution theorem states that as we use more harmonics, the eigenvalues mimic the values taken by the function itself. As we make the Toeplitz matrix larger, the smallest and largest eigenvalues converge to the minimum and maximum values of the function respectively. Negative eigenvalues mean that the coefficients can't be extended to correspond to a non-negative function. So, if we constrain the eigenvalues to be non-negative (*i.e* the matrix to be positive semidefinite), we can exclude all those low frequency functions that can't be extended by adding higher frequencies to become non-negative everywhere.

We extend this theorem to spherical harmonics. In this case, we obtain a much more complicated matrix whose eigenvalues are similar to the function values. To constrain this matrix to be positive semidefinite while minimizing the error between the query and generated images, we use semidefinite programming.

Next, we perform experiments on both synthetic and real data to explore the usefulness of imposing this constraint. We observe that imposing this constraint results in a significantly greater mismatch between the query and incorrect models, for most specular objects. This can improve recognition since now it is harder for the algorithm to get confused by noise in the model or query image.

This paper is divided as follows. First, in section 2, we review some earlier work that has used the non-negativity constraint. In section 3, we present the extension of Szego's eigenvalue distribution theorem to spherical harmonics: the key ingredient in our algorithm. Next, in section 4, we review recovering lighting from an image given an object model, and formulate the problem as a semidefinite program. Finally, section 5 describes some experiments on synthetic as well as real data, which demonstrate the usefulness of the constraint.

2 Past Work

Various approaches to object recognition have used low dimensional linear subspace representations of the set of images produced by an object. For example, Hallinan [6], Murase and Nayar

[8] and Yuille et al [19] have used PCA to model lighting variation and Basri and Jacobs [1] and Ramamoorthi and Hanrahan [12] have used a spherical harmonic representation for an analytic computation of the linear subspace of images. We also use a spherical harmonic representation for images and lighting.

Belhumeur and Kriegman [2] have shown that the set of all possible images of an object under arbitrary lighting is a convex cone, the *illumination cone*. Lighting is represented as a convex combination (to ensure non-negativity) of the extreme rays of the convex cone. Computation and memory requirements can be reduced by projecting the image, the illumination cone and the extreme rays into a low dimensional subspace, although this makes the representation approximate. Calculation of the extremal rays can be avoided by further approximating lighting as a convex combination of rays uniformly sampled from the illumination sphere. They use a non-negative least squares routine to perform the convex optimization.

For Lambertian objects, Basri and Jacobs [1] build on this by expressing the uniformly sampled rays in terms of spherical harmonics. This approach works well for Lambertian objects since they only reflect the diffused (low frequency) components of the incident lighting which are well approximated by a few delta functions. However, since specular objects reflect many more components of light, a very large number of delta functions are needed to represent lighting accurately for them. This method is also approximate since the delta functions are approximated by a few low frequency harmonics and are no longer just positive peaks. Non-negativity of lighting was also enforced by Ramamoorthi et al [11] using a regularization term during optimization. This clearly cannot guarantee that the solution will be non-negative.

There have been many other attempts at recognizing specular objects. Osadchy et al [10] have used specular reflection in recognition by decoupling Lambertian reflection and highlights and using them as separate cues. Sato et al [14] use a physics-based simulator to predict specular features and analyze their detectability and reliability for recognition. Specularity detection is performed using a set of aspects generated from the model by deformable template matching. Gremban et al [3] use multiple views of an object to remove ambiguities due to specularities.

In this paper, we will describe a new, exact method for enforcing non-negativity, as a direct constraint on the spherical harmonic coefficients of light.

3 The non-negativity constraint

We need a condition on the first few spherical harmonic coefficients f_{lm} of a function $f(u) : S^2 \rightarrow \mathbb{R}$ that will imply that we can complete the spherical harmonic expansion of f such that $f(u) \geq 0$ for all u . Here, $u := (\theta, \phi)$ is a point on the surface of the unit sphere, denoted as S^2 . This problem is easier to deal with in 1D, when we need a condition on the Fourier series coefficients f_m of a function $f(\theta) : S^1 \rightarrow \mathbb{R}$ (θ is a point on the unit circle S^1). The condition for non-negativity that we obtain in these two cases is completely analogous; but the expressions are simpler for S^1 and the more familiar Fourier series will help us to understand the method better.

3.1 The Fourier case

Let \mathcal{Q}_n denote the space of functions on S^1 spanned by $\{e^{im\theta} : 0 \leq m \leq n\}$, *i.e* functions that only have low frequency components. The process of low pass filtering a function, so that the output belongs to \mathcal{Q}_n is the same as an orthogonal projection from the space of all (integrable or L^1) functions to \mathcal{Q}_n . We will represent this operation as Q_n . Let f_m denote the m^{th} Fourier coefficient of the function $f \in L^1(S^1)$.

We will now develop some intuitive ideas about the non-negativity condition. First, let's represent the time domain product of two functions f and g , using only their Fourier series coefficients, as a product of a matrix (composed of the coefficients of f) and the vector of coefficients of g , denoted as \hat{g} . We can do this using the convolution theorem, if we consider only the first n coefficients of f and g . Then, the result will be the first n coefficients of the time domain product fg .

Let $[f]$ denote the operator *multiplication by f* . Then, denote the matrix of coefficients of f by $Q_n[f]Q_n$. In this notation, the first Q_n indicates that we are considering only the first n coefficients of f , which is equivalent to applying an ideal low pass filter to f or projecting f into a low dimensional subspace spanned by the first n Fourier basis functions. The resulting time domain function is $f^{(n)}$. The second Q_n indicates the same for the function g . $\hat{g}_n = Q_n\hat{g}$ is a vector of the first n Fourier coefficients of g . Thus, we have,

$$Q_n[f]Q_n\hat{g} = Q_n\widehat{fg} \quad (1)$$

Using the convolution theorem, we arrive at the following form for the matrix $Q_n[f]Q_n$, called a Toeplitz matrix.

$$Q_n[f]Q_n = T_n(f) = \begin{bmatrix} f_0 & f_1 & \cdots & f_n \\ f_{-1} & f_0 & \ddots & \\ \vdots & \ddots & \ddots & f_1 \\ f_{-n} & & f_{-1} & f_0 \end{bmatrix} \quad (2)$$

The $(ij)^{th}$ element of this matrix is f_{j-i} . Now if \hat{g}_n is an eigenvector of the matrix $Q_n[f]Q_n$, with the eigenvalue λ .

$$T_n(f)\hat{g} = \lambda\hat{g} \quad (3)$$

In the time domain, we have

$$f^{(n)}g^{(n)} = \lambda g^{(n)} \quad (4)$$

It is clear from this equation that λ lies in the range of values taken by $f^{(n)}$. Actually, we can show that λ lies in the range of f too. Also, although this is not obvious from our crude treatment, the eigenvalues λ are representative of the values taken by the function f itself. These ideas are made concrete in Szego's eigenvalue distribution theorem [4]. This theorem states that the mean value of any continuous function is the same whether it is applied to the eigenvalues of $Q_n[f]Q_n$ or to the values of the function f , i.e the eigenvalues are "distributed" in the same way as the values of f . Hence, we can constrain the range of f by constraining the eigenvalues.

Before stating the theorem, we need this definition: the *essential lower bound* (or *essential infimum*, denoted by ess inf) of a function $f(x)$ is the largest number m for which the inequality $f(x) \geq m$ holds everywhere, except perhaps in a set of measure zero. The *essential upper bound* (or *essential supremum*, denoted by ess sup) is defined similarly. First, we give another result that states that the eigenvalues of $T_n(f)$ lie in the range of f .

Theorem 1. *Let $f(\theta) \in L^1(S^1)$ be a real valued function and $T_n(f)$ be the Toeplitz matrix of its Fourier series coefficients. $\lambda_i^{(n)}$, $i = 1, \dots, n+1$ are the eigenvalues of $T_n(f)$ arranged in non-decreasing order. Let m and M be the essential lower and upper bounds of $f(\theta)$ respectively. Then,*

$$m \leq \lambda_1^{(n)} \leq \lambda_2^{(n)} \leq \cdots \leq \lambda_{n+1}^{(n)} \leq M \quad (5)$$

The next step is the Szego Limit theorem, a fundamental result in the theory of Toeplitz forms proved by Gabor Szego in 1917, and extended in 1955. Here we only need the original limit theorem, and not its strong form. This can be thought of as a particular case of the main theorem, and is the primary result used in its proof.

Theorem 2 (Szego Limit Theorem). *Let $f(\theta) \in L^1(S^1)$ be a real valued function and $T_n(f)$ be its Toeplitz matrix as defined above. Then*

$$\lim_{n \rightarrow \infty} \frac{\text{tr} \log T_n(f)}{n+1} = \frac{1}{2\pi} \int_{S^1} \log f(\theta) d\theta$$

Note that $\text{tr} \log T_n(f) = \log \det T_n(f) = \sum_{i=1}^{n+1} \log \lambda_i^{(n)}$. Now, we state the main theorem.

Theorem 3 (Szego: Eigenvalue Distribution Theorem). *With notation and conditions as above, and with m and M finite, let $F(\lambda)$ be any continuous function defined in the interval $\lambda \in [m, M]$, then*

$$\lim_{n \rightarrow \infty} \frac{F(\lambda_1^{(n)}) + \dots + F(\lambda_{n+1}^{(n)})}{n+1} = \frac{1}{2\pi} \int_{S^1} F(f(\theta)) d\theta \quad (6)$$

Corollary 1. *With notation and conditions as above,*

$$\lim_{n \rightarrow \infty} \lambda_1^{(n)} = m, \quad \lim_{n \rightarrow \infty} \lambda_{n+1}^{(n)} = M \quad (7)$$

The proofs of these theorems (except that of the Szego Limit theorem) and the corollary are very similar to the proofs for the spherical harmonics case given in the next section. Proofs can also be found in Grenander and Szego [4]. The first theorem and the corollary imply that if the function $f(\theta)$ is non-negative everywhere, then the matrix $T_n(f)$ is positive semidefinite¹ for all n . Also, as we will see from Theorem 7, if $T_n(f)$ is positive semidefinite, there exists a function f with the projection $Q_n f$ that is non-negative everywhere (except perhaps on a set of measure zero). Thus, even though the projection $Q_n f$ that we obtain may not be non-negative everywhere, we are guaranteed that it is the projection of a function that is non-negative everywhere. Also, if $T_n(f)$ is not positive semidefinite, we are guaranteed that the projection $Q_n f$ cannot be extended into a non-negative function f . Thus, using this constraint rules out all those lighting function projections and only those projections that do not correspond to a physical lighting function. Note that although the Dirac delta function is not an element of L^1 , the theorem is valid for it too, since all the eigenvalues of $T_n(\delta)$ are zero, except for $\lambda_{n+1}^{(n)} = n$ which goes to infinity as n increases.

3.2 Spherical Harmonics

Next, we extend the theorem to the case of spherical harmonics, *i.e.* to functions on S^2 . Let \mathcal{P}_L be the space of functions that only have spherical harmonic components of order up to L . Correspondingly, P_L denotes the process of ideal low pass filtering the function f , so that we only retain spherical harmonic components of order at most L . As we go from S^1 to S^2 , we have to restrict the set of functions on which the corresponding theorems are valid to $H^{\frac{1}{2}}(S^2)$, a Sobolev space of functions defined on S^2 . A function is said to belong to a Sobolev space H^k if it has finite norm (the L^2 norm in this case) and also if the norm of all derivatives of

¹The symmetric matrix A is positive semidefinite if $x^t A x \geq 0$ for all vectors x , or equivalently, if all its eigenvalues are non-negative.

the function up to order k is finite. The derivatives need to exist only in a ‘weak’ sense. (see [20] for more on Sobolev spaces and [9] for details on this theorem). $H^{\frac{1}{2}}(S^2)$ is the space of functions such that the norm $\|(I - \Delta)^{\frac{1}{4}}f\|_{L^2}$ is finite. Here, Δ is the Laplace-Beltrami second derivative operator, a generalization of the normal Laplacian operator for manifolds. I is the identity operator. For more details, see [21]. $C(S^2)$ is the space of continuous functions defined on S^2 . Most well-behaved and smooth functions belong to these spaces. First we show that the eigenvalues of $P_L[f]P_L$ are contained in the range of f .

Theorem 4. *Let $f(u) \in L^1(S^2)$ be a real valued function. Let m and M be the essential lower and upper bounds of $f(u)$, respectively, λ_i , $i = 1, \dots, (L+1)^2$ are the eigenvalues of the matrix $P_L[f]P_L$. Then,*

$$m \leq \lambda_1^{(L)} \leq \lambda_2^{(L)} \leq \dots \leq \lambda_{(L+1)^2}^{(L)} \leq M \quad (8)$$

Proof. This proof is similar to the proof for the Fourier case given in [4]. Let \hat{g} be any vector of length $(L+1)^2$ with unit L^2 norm and let $g(u) = \sum_{lm} \hat{g}_{lm} Y_{lm}(u)$. Thus, $\int_{S^2} |g(u)|^2 d\sigma(u) = \sum_{lm} |\hat{g}_{lm}|^2 = 1$. By definition of $P_L[f]P_L$, we have,

$$P_L[f]P_L\hat{g} = \widehat{fg}$$

Consider the following quadratic form :

$$\begin{aligned} \hat{g}^* P_L[f]P_L\hat{g} &= \sum_{lm} \hat{g}_{lm}^* \widehat{fg}_{lm} \\ &= \sum_{lm} \hat{g}_{lm}^* \int_{S^2} f(u)g(u)\overline{Y_{lm}(u)}d\sigma(u) \\ &= \int_{S^2} f(u)g(u) \left(\sum_{lm} \hat{g}_{lm}^* \overline{Y_{lm}(u)} \right) d\sigma(u) \\ &= \int_{S^2} f(u)g(u)\overline{g(u)}d\sigma(u) \\ &= \int_{S^2} f(u)|g(u)|^2 d\sigma(u) \end{aligned}$$

Since, g is normalized, the last expression is simply a weighted mean of the function f and hence it lies in the essential range of f . Hence, we have,

$$m \leq \hat{g}^* P_L[f]P_L\hat{g} \leq M$$

for all vectors \hat{g} with unit norm. If we choose \hat{g} to be an eigenvector corresponding to any eigenvalue λ of $P_L[f]P_L$, we get

$$m \leq \lambda \leq M$$

Thus, all the eigenvalues of $P_L[f]P_L$ are contained in the range of f . □

We will now state the Szego limit theorem which is the main result needed in the proof of the eigenvalue distribution theorem. A more general form of this theorem is stated and proved by Okikiolu in [9].

Theorem 5 (Szego Limit Theorem in S^2). *Let $f \in C(S^2) \cap H^{1/2}(S^2)$ be such that the closed convex hull of the image of f does not contain the origin,² then*

²Since are dealing with real valued functions only, this means that f takes either only positive or only negative values, but not both.

$$\lim_{n \rightarrow \infty} \frac{\text{tr} \log P_n[f] P_n}{(n+1)^2} = \frac{1}{4\pi} \int_{S^2} \log f(u) d\sigma(u)$$

Before stating and proving the eigenvalue distribution theorem, we need the formal definition of the concept of *equal distributions*, as given by H. Weyl [4]. Two sequences of numbers $\{a_i^{(n)}\}_{i=1, \dots, n+1}$ and $\{b_i^{(n)}\}_{i=1, \dots, n+1}$ such that $|a_i^{(n)}| < K$ and $|b_i^{(n)}| < K$ for all i and n are equally distributed in the interval $[-K, K]$ as $n \rightarrow \infty$ if for any continuous function $F(t)$ defined in the interval $[-K, K]$, we have

$$\lim_{n \rightarrow \infty} \frac{\sum_{i=1}^{n+1} [F(a_i^{(n)}) - F(b_i^{(n)})]}{n+1} = 0 \quad (9)$$

Here we use a slightly modified definition in which for each n , the sequences consist of $(n+1)^2$ instead of $(n+1)$ numbers. Roughly, we can say that two sequences are equally distributed if they take on similar values. We also need this test for equally distributed sequences: Two sequences obey the equation (9) for arbitrary continuous functions F , *ie* they are equally distributed, if the equation (9) is satisfied for certain special classes of functions. Two such classes of these functions are $F(t) = \log(1 + zt)$ where z is real and $|z| < K^{-1}$ and $F(t) = t^s$ where $s = 0, 1, 2, \dots$

Theorem 6 (Eigenvalue Distribution Theorem in S^2). *Let $f(u) \in C(S^2) \cap H^{1/2}(S^2)$ be a real valued function. Let m and M be the essential lower and upper bounds of $f(u)$, respectively and assume that m and M are finite. $\lambda_i^{(L)}$, $i = 1, \dots, (L+1)^2$ are the eigenvalues of the matrix $P_L[f]P_L$. If $F(\lambda)$ is any continuous function defined in the finite interval $\lambda \in [m, M]$, then*

$$\lim_{L \rightarrow \infty} \frac{F(\lambda_1^{(L)}) + \dots + F(\lambda_{(L+1)^2}^{(L)})}{(L+1)^2} = \frac{1}{4\pi} \int_{S^2} F(f(u)) d\sigma(u)$$

This is a novel result. The proof of this theorem closely follows the proof of Szego's original theorem and uses Okikiolu's [9] extension of a key lemma used in the theorem's proof - the Szego limit theorem.

Proof. The proof of this result follows Szego's original proof. Using the definition of Riemann integration as a limit, the theorem is equivalent to :

$$\lim_{L \rightarrow \infty} \frac{\sum_{m=1}^{(L+1)^2} F(\lambda_m^{(L)}) - F(f(u_m^{(L)}))}{(L+1)^2} = 0$$

where $u_m^{(L)} = (\frac{2a\pi}{L+2} - \pi, \frac{b\pi}{L+2} - \frac{\pi}{2})$; $a, b = 1, 2, \dots, L+1$; $a + (b-1)L = m$. Using the definition of equal distributions we can restate the limit relation as follows:

The sequences $\lambda_m^{(L)}$ and $f(u_m^{(L)})$ are equally distributed.

Proving this for the special class of functions $F(t) = \log(1 + zt)$ is sufficient. We will now apply the Szego limit theorem to the function $1 + zf(u)$, where $z \in \mathbb{R}$ is such that $|zf(u)| < 1$ for all $u \in S^2$. We can do this since the function $f(u)$ is bounded. This transformation ensures that the closed convex hull of the image of $1 + zf(u)$ does not contain the origin. We have

$$P_L[1 + zf]P_L = I + zP_L[f]P_L$$

since the vector space \mathcal{P}_L is closed *wrt* scaling and shifting. The eigenvalues of $I + zP_L[f]P_L$ are $1 + z\lambda_m^{(L)}$. Hence, using the fact that for any nonsingular matrix A , $\text{tr} \log A = \log \det(A) = \sum_i \log(\lambda_i(A))$, we can write

$$\lim_{L \rightarrow \infty} \frac{\sum_{m=1}^{(L+1)^2} \log(1 + z\lambda_m^{(L)})}{(L+1)^2} = \frac{1}{4\pi} \int_{S^2} \log(1 + zf(u)) d\sigma(u)$$

Again using the definition of Riemann integration, we have

$$\lim_{L \rightarrow \infty} \frac{\sum_{m=1}^{(L+1)^2} \log(1 + z\lambda_m^{(L)}) - \log(1 + zf(u_m^{(L)}))}{(L+1)^2} = 0$$

Thus, the theorem is valid for the set of functions $\log(1 + zt)$ and hence is valid for all continuous functions. \square

We also have the corresponding corollary :

Corollary 2. *With notation and conditions as above,*

$$\lim_{L \rightarrow \infty} \lambda_1^{(L)} = m, \quad \lim_{L \rightarrow \infty} \lambda_{(L+1)^2}^{(L)} = M \quad (10)$$

Proof. In the Fourier case, this is proved using a specific property of Toeplitz matrices (the fact that T_L is a principal submatrix of T_{L+1}) that does not hold in the case of spherical harmonics. Here we give a different proof of this result. From the proof of theorem 4, we have

$$\frac{\hat{g}_L^* T_L(f) \hat{g}_L}{\hat{g}_L^* \hat{g}_L} = \frac{\int_{S^2} f(u) |g_L(u)|^2 d\sigma(u)}{\int_{S^2} |g_L(u)|^2 d\sigma(u)}$$

The subscript L indicates that $g_L \in \mathcal{P}_L$. So, the minimum eigenvalue of $T_L(f)$ is given by

$$\lambda_{min}^{(L)} = \min_{g_L \in \mathcal{P}_L} \frac{\int_{S^2} f(u) |g_L(u)|^2 d\sigma(u)}{\int_{S^2} |g_L(u)|^2 d\sigma(u)}$$

Since f is defined on the closed and bounded interval S^2 and is continuous on it, it attains its essential lower bound. Let u_m be a point where $f(u)$ attains its essential lower bound. Now, choose g_L as the real valued function:

$$g_L = P_L[C_L e^{-L^2 \|u - u_m\|^2}]$$

i.e., g_L is the projection of a Gaussian into \mathcal{P}_L . The variance of the Gaussian decreases as $1/L^2$, but the normalization constant C_L is chosen to keep its area constant. Clearly, as $L \rightarrow \infty$, $g_L \rightarrow C_g \delta(u - u_m)$ and $|g_L|^2 \rightarrow C_{g^2} \delta(u - u_m)$ as well, for some suitable constants C_g and C_{g^2} . C_L can be chosen to make $\int_{S^2} |g_L|^2 d\sigma(u) = C_{g^2} = 1$. Note that we can take g_L as any function in \mathcal{P}_L whose square converges to a delta function. So, we have

$$\begin{aligned} \lim_{L \rightarrow \infty} \lambda_{min}^{(L)} &\leq \lim_{L \rightarrow \infty} \frac{\int_{S^2} f(u) |g_L(u)|^2 d\sigma(u)}{\int_{S^2} |g_L(u)|^2 d\sigma(u)} \\ &= \frac{\int_{S^2} f(u) (\lim_{L \rightarrow \infty} |g_L(u)|^2) d\sigma(u)}{\int_{S^2} (\lim_{L \rightarrow \infty} |g_L(u)|^2) d\sigma(u)} \\ &= \frac{\int_{S^2} f(u) \delta(u - u_m) d\sigma(u)}{\int_{S^2} \delta(u - u_m) d\sigma(u)} \\ &= f(u_m) \\ &= m \end{aligned}$$

Since we must have $\lambda_{min}^{(L)} \geq m$ for all L , we conclude that

$$\lim_{L \rightarrow \infty} \lambda_{min}^{(L)} = m$$

The second limit result can be proved similarly. \square

All we need to do now is calculate the matrix $T_L(f) := P_L[f]P_L$. We will use something similar to the convolution theorem and calculate the (l_1, m_1) th coefficient of the time domain product fg , where $g \in L^2(S^2)$ is any real valued function.

$$\begin{aligned} (P_L[f]P_L\hat{g})_{l_1 m_1} &= (\widehat{fg})_{l_1 m_1} \\ &= \int_{S^2} fg \bar{Y}_{lm} d\sigma(u) \\ &= \sum_{l_2 m_2} \sum_{lm} f_{lm} g_{l_2 m_2} \int_{S^2} Y_{lm} Y_{l_2 m_2} \bar{Y}_{l_1 m_1} d\sigma(u) \\ &= \sum_{l_2 m_2} \sum_{lm} f_{lm} g_{l_2 m_2} G(ll_2 l_1; mm_2 m_1) \\ &= \sum_{l_2 m_2} \sum_{l=|l_1-l_2|}^{l_1+l_2} f_{l, m_1-m_2} G(ll_2 l_1; m_1-m_2, m_2, m_1) g_{l_2 m_2} \end{aligned}$$

Here, $G(l_1 l_2 l_3; m_1 m_2 m_3) = \sqrt{\frac{(2l_1+1)(2l_2+1)}{4\pi(2l_3+1)}} C(l_1 l_2 l_3; m_1 m_2 m_3) C(l_1 l_2 l_3; 000)$ is the Gaunt coefficient. $C(l_1 l_2 l_3; m_1 m_2 m_3)$ are the Clebsch-Gordan (CG) coefficients. Both these coefficients are real constants that arise naturally during the evaluation of integrals of products of spherical harmonics. For more details, please see Appendix A.1 and [13, 7]. Thus the term in position $(l_1 m_1, l_2 m_2) := (l_1(l_1+1) + m_1, l_2(l_2+1) + m_2)$ in the matrix $P_L[f]P_L$ is

$$T_L(f)_{l_1 m_1; l_2 m_2} = \sum_{l=|l_1-l_2|}^{l_1+l_2} f_{l, m_1-m_2} G(ll_2 l_1; m_1-m_2, m_2, m_1) \quad (11)$$

The choice of the subscripts was made so that 1 corresponds to the row number and 2 to the column number in the matrix. The size of the matrix is $(L+1)^2 \times (L+1)^2$, since there are $(L+1)^2$ spherical harmonics of degree less than or equal to L . We can show that this matrix is Hermitian, which is important because optimization software usually needs symmetric matrices as inputs.

$$\begin{aligned} \overline{(P_n[f]P_n)}_{l_1 m_1; l_2 m_2} &= \sum_{l=|l_1-l_2|}^{l_1+l_2} \bar{f}_{l, m_1-m_2} \sqrt{\frac{(2l+1)(2l_2+1)}{4\pi(2l_1+1)}} C(ll_2 l_1; m_1-m_2, m_2, m_1) C(ll_2 l_1; 000) \\ &= \sum_{l=|l_1-l_2|}^{l_1+l_2} (-1)^{m_1-m_2} f_{l, m_2-m_1} \sqrt{\frac{(2l+1)(2l_2+1)}{4\pi(2l_1+1)}} \\ &\quad \times (-1)^{m_2-m_1} \left(\frac{2l_1+1}{2l_2+1} \right) C(ll_1 l_2; m_2-m_1, m_1, m_2) C(ll_1 l_2; 000) \end{aligned}$$

(using properties (21) and (23) of the CG coefficients as listed in Appendix A.1)

$$\begin{aligned}
&= \sum_{l=|l_1-l_2|}^{l_1+l_2} f_{l,m_2-m_1} \sqrt{\frac{(2l+1)(2l_1+1)}{4\pi(2l_2+1)}} C(l l_1 l_2; m_2 - m_1, m_1, m_2) C(l l_1 l_2; 000) \\
&= (P_n[f] P_n)_{l_2 m_2; l_1 m_1}
\end{aligned}$$

If the function f is non-negative everywhere, $T_L(f)$ is positive semi-definite (denoted as $T_L(f) \succeq 0$). We also need to look at the converse problem. Does the positive semidefiniteness of $T_L(f)$ imply that f is non-negative everywhere? Since we deal only with the first few harmonic components of f , we can arbitrarily add higher order harmonics to f . This gives an infinite number of functions f corresponding to the same matrix $T_L(f)$, all of which are obviously not non-negative. However what we are interested in is the existence of at least one function with the given matrix $T_L(f)$, *i.e* with the given lower order harmonics, that is non-negative everywhere. This will ensure that the set of lower order harmonics obtained from the optimization corresponds to a non-negative lighting condition. For the Fourier case, we have the following theorem that answers this question.

Theorem 7. *Given the first n Fourier coefficients of a real valued function $f(\theta)$, $\theta \in S^1$, if the Toeplitz matrix $T_n(f)$ (defined by equation (2)) is positive semidefinite, there exists a unique function with the given lower order Fourier coefficients that is non-negative everywhere. This is the sum of delta functions given by :*

$$f(\theta) = K_0 + \sum_{p=1}^n K_p \delta(\theta - \theta_p) \quad (12)$$

$$(K_0 = 0, K_p \geq 0)$$

Furthermore, if the matrix $T_n(f)$ is strictly positive definite, then there exist infinite functions with the given lower order frequency components that are non-negative everywhere.

Proof. First consider the case when $T_n(f)$ is positive semidefinite. We need to express the given Fourier coefficients as the sum of the Fourier series coefficients of a set of n non-negative delta functions. This can be done by using the theorem of Carathéodory, (Appendix C and [4] section 4.1). Corresponding to the complex constants f_1, f_2, \dots, f_n , we have the unique real numbers $K_p \geq 0$ and $\theta_p \in S^1$; $p = 1, 2, \dots, n$ such that

$$f_\nu = \sum_{p=1}^n K_p e^{i\nu\theta_p}$$

(note that some of the K_p may be zero). Since we already have $f_{-\nu} = \bar{f}_\nu$ and f_0 is such that the Toeplitz matrix formed from them is positive semidefinite, this equation is valid for all the given Fourier series coefficients of f . One possible function that has these initial Fourier series coefficients and is non-negative everywhere is :

$$f(\theta) = \sum_{p=1}^n K_p \delta(\theta - \theta_p)$$

Now suppose there is another function f' that is non-negative everywhere and has the same low frequency components. Then $f' - f$ is composed only of high frequency components and

hence must be negative on a set of finite non-zero measure. Since f is zero everywhere except at finitely many points, the sum of f and $f' - f$ cannot be non-negative everywhere. We thus have a contradiction.

Now, if $T_n(f)$ is positive definite, we can write the following form for $f(\theta)$:

$$f(\theta) = K_0 + \sum_{p=1}^n K_p \delta(\theta - \theta_p)$$

To obtain this representation, note that we can write $f(\theta) = K_0 + \tilde{f}(\theta)$ such that $T_n(\tilde{f})$ is positive semidefinite and $K_0 > 0$. Now we can obtain a representation for $\tilde{f}(\theta)$ as before. To see that there are infinitely many other non-negative functions with the same $T_n(f)$, add any frequency component to $f(\theta)$ of order greater than n and magnitude less than K_0 . The resulting function has the same Toeplitz matrix and is non-negative everywhere. \square

In the case of spherical harmonics, we have the corresponding conjecture:

Conjecture 1. *Given the frequency components of a real valued function $f(u)$, $u \in S^2$, up to order L , if the matrix $T_L(f)$ (defined by equation (11)) is positive semidefinite, there exists a unique function $f(u)$ with the given lower order frequency components that is non-negative everywhere. This is the sum of δ function given by :*

$$f(u) = K_0 + \sum_{p=1}^{N(L)-1} K_p \delta(u - u_p) \quad (13)$$

$$(K_0 = 0, K_p \geq 0, N(L) = (L + 1)^2)$$

Furthermore, if the matrix $T_L(f)$ is strictly positive definite, then there exist infinite functions with the given lower order frequency components that are non-negative everywhere.

We are unable to prove this because, as far as we know, Carathéodory's theorem does not have a spherical harmonics analog. If the existence part is proved, then uniqueness in the positive semidefinite case and the existence of infinite functions in the strictly positive definite case can be proved exactly as in the Fourier case.

4 Recovering Lighting from an Image: Semidefinite Programming

The problem of recovering lighting from the image of an object using its model, can be treated as an optimization problem. We assume a geometric model in the form of surface normals at the various pixel locations and a reflectance model. If we have several models from several different objects, the model that gives the least error corresponds to the object that created the image.

We represent lighting in terms of spherical harmonics and analytically compute the image when the object is illuminated by each individual harmonic. If these images are treated as vectors and stacked as columns of a matrix, we obtain the model matrix M . In rendering the images, we can use any reflectance model, or even a mixture of models. Now if the lighting is described by the spherical harmonic coefficient vector a , the resulting image (as a vector I) is given by $I = Ma$. If the image has N pixels and we use spherical harmonics up to order L to describe the image, M has size $N \times (L + 1)^2$ and a and I are column vectors of lengths $(L + 1)^2$ and N respectively. Then, given the query image $r = I + \text{noise}$, a can be found by

minimizing $\|Ma - r\|$ subject to $T_L(a) \succeq 0$. Since we model frequencies only up to L , the error will usually be non-zero even in the absence of noise. The problem size here is the number of pixels in the image, which can be pretty large. We can reduce this by transforming the problem from the image space to the space of spherical harmonics basis images (see [1]). First, compute the QR decomposition of the matrix M , *i.e.* $M = QR$ where Q is an $N \times (L + 1)^2$ matrix with orthonormal columns ($Q^T Q = I$) and R is an $(L + 1)^2 \times (L + 1)^2$ upper triangular matrix. Next we project both Ma and f into the low $((L + 1)^2)$ dimensional subspace by multiplying with Q^T . We now need to solve the size $(L + 1)^2$ problem:

$$\begin{aligned} \min_a \quad & \|Ra - Q^T r\|^2 \\ \text{subject to} \quad & T_L(a) \succeq 0 \end{aligned}$$

This is an optimization problem with a quadratic objective function and a matrix positive semi-definiteness constraint. Such problems are called semidefinite programming (SDP) problems. The matrix constraint itself is considered linear in SDP since each element of the matrix depends linearly on the vector a (see section 5.1) and is a type of a Linear Matrix Inequality. We can convert this into a linear problem (see [16]) by introducing a slack variable q .

$$\begin{aligned} \min_a \quad & q \\ \text{subject to} \quad & q > \|Ra - Q^T r\|^2 \text{ and } T_L(a) \succeq 0 \end{aligned}$$

Now, the Schur complement property (see [16]) is used to convert the quadratic constraint into the following equivalent linear constraint.

$$\begin{bmatrix} (1 + q)I & (Ra - Q^T r)^{1-q} \\ (Ra - Q^T r)^{1-q}{}^T & (1 + q)I \end{bmatrix} \succeq 0$$

This can be readily converted to a second order cone program (SOCP), which can be solved faster. In an SOCP, the constraint is of the form $\begin{pmatrix} \alpha \\ v \end{pmatrix} \in K_2 := \{ \begin{pmatrix} \alpha \\ v \end{pmatrix} | \alpha \geq \|v\|_2 \}$. K_2 is called the second order cone or Lorentz cone.

$$1 + q \geq \left\| \begin{matrix} 1 - q \\ Ra - Q^T r \end{matrix} \right\|$$

So the final problem to be solved is :

$$\begin{aligned} \min_a \quad & q & (14) \\ \text{subject to} \quad & 1 + q \geq \left\| \begin{matrix} 1 - q \\ Ra - Q^T r \end{matrix} \right\| \text{ and } T_L(a) \succeq 0 \end{aligned}$$

This is a mixed SOCP-SDP problem and can be solved using standard SDP software.

5 Experiments

5.1 Implementation

The entries of the matrix $T_L(a)$ are linear combinations of the entries of the vector a and are given by:

$$T_L(l_1(l_1 + 1) + m_1, l_2(l_2 + 1) + m_2) = \sum_{l=|l_1-l_2|}^{l_1+l_2} a_{l, m_1-m_2} G(l_2 l_1; m_1 - m_2, m_2, m_1) \quad (15)$$

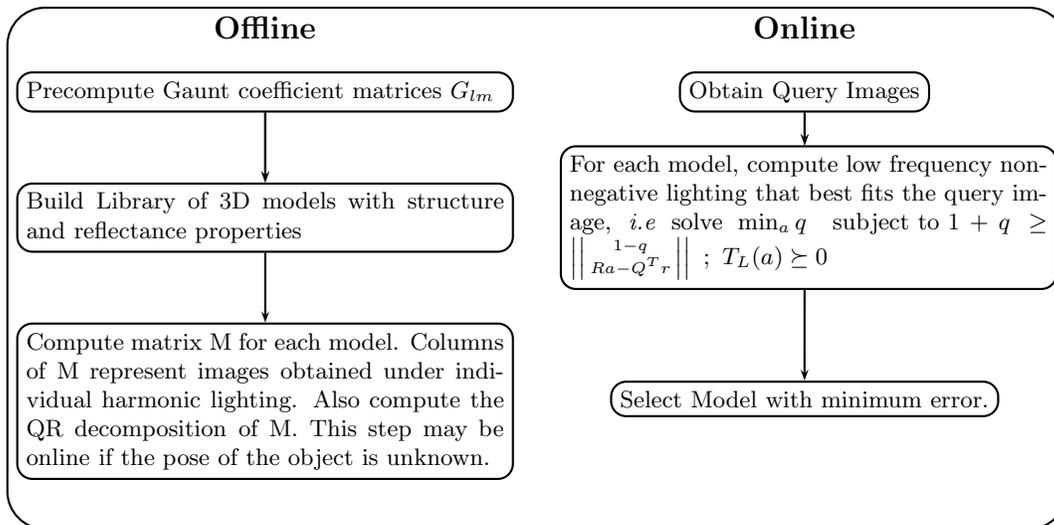


Figure 2: Specular Object Recognition Algorithm

where $G(l_2 l_1; m_1 - m_2, m_2, m_1)$ are the Gaunt coefficients. Since each element of T_L is a linear combination of Gaunt coefficients, we can write it as a linear combination of matrices of Gaunt coefficients, with the elements of a as coefficients.

$$T_L = \sum_{l=0}^L \sum_{m=-l}^l a_{lm} G_{lm} \quad (16)$$

where $G_{lm}(l_1(l_1 + 1) + m_1, l_2(l_2 + 1) + m_2) = G(l_2 l_1; m m_2 m_1)$. Note that $G(l_2 l_1; m m_2 m_1)$ is zero unless $m = m_1 - m_2$. To speed up computation, the matrices G_{lm} are precomputed and stored. Since SDP solvers usually deal only with real valued problems, we use real versions of spherical harmonics instead of the normal complex versions, as described in [7]. The flowchart 2 describes the whole object recognition algorithm.

The SDP is solved in MATLAB 6.5 using the SDPT3 [17] package. Since implementing it directly in SDPT3 is difficult, YALMIP [18] is used for formulating the problem. This is a problem translator that can convert the problem description in its format to that of a wide variety of SDP solvers available for MATLAB. SDPT3 uses a polynomial time predictor-corrector primal-dual infeasible path following algorithm to solve SDP and SOCP problems, and is one of the fastest solvers available for small to medium scale problems. Table 1 compares the times for recovering lighting from an image, using our algorithm (SDP) and the Basri and Jacobs [1] algorithm (Delta), as the number of harmonics used increases. The computer used was a 2.66GHz Pentium 4 with 512MB RAM. [1] use a non-negative combination of delta functions to represent lighting. As the number of delta functions increases, we find that this method produces the same solution as SDP. In this comparison, the number of delta functions is set to obtain less than 1% error. The image size used was 26×51 . Increasing image size just adds the same small time (for the QR decomposition) to both methods. From the table, we see that the time for SDP increases more slowly than that of Delta, and SDP overtakes Delta at around $L = 6$. In our experiments, we have found using $L = 10$ to be sufficiently accurate for a lot of common specular objects. In this case, our method is 35 times faster than Delta, while being exact as well.

| Max L | Number of Harmonics $(L + 1)^2$ | SDP time (s) | Number of Delta functions | Delta time (s) |
|---------|---------------------------------|--------------|---------------------------|----------------|
| 2 | 9 | 0.38 | 32 | 0.03 |
| 4 | 25 | 0.61 | 1922 | 0.52 |
| 6 | 49 | 1.41 | 1922 | 1.41 |
| 8 | 81 | 1.92 | 1922 | 4.28 |
| 10 | 121 | 3.55 | 30722 | 126 |

Table 1: Speed comparison of SDP and Delta function method [1].

5.2 Experiments on Synthetic Images

In these experiments, we evaluate the effect of using the non-negativity constraint (SDP). If the constraint is not imposed, the problem is reduced to simply solving a system of linear equations (LIN).

5.2.1 Variation of error with model specularity and query image frequency

In this experiment, we investigate the effect of the frequency content of the query image and the specularity of the object on the importance of the non-negativity constraint. The model is a varying linear combination of a mirror and a uniform Lambertian albedo (α is the proportion of mirror). The query images are composed of individual harmonics (Y_{l0} for $l = 1, \dots, 30$). These elementary query images will enable us to predict SDP’s behavior qualitatively on normal images, which are a linear combination of individual harmonics. The optimization procedure uses spherical harmonics up to degree $L = 10$. Mirror reflection causes image harmonics of order up to double that of incident light harmonics (*i.e* up to order 20).

Since the images are not produced by any object, we don’t expect any model to have zero error. The magnitude of error will give us an idea of how well the two methods can avoid choosing the wrong model : a larger error means that it is more difficult to fool. The results are shown in figure 3. Firstly, note that SDP has higher error than LIN (which is almost zero) for $L \leq 20$. This is the range of image frequencies that are modeled by the algorithm and the use of SDP should reduce recognition errors here. However, SDP error decreases as the model becomes more specular, and hence the advantage of using SDP decreases. For mirrors and almost-mirrors, using SDP is not likely to help significantly in recognition. Also, we can use the spherical harmonic content of the image to guide us in choosing the number of harmonics required to represent lighting. For example, if most of the harmonic content of the image is of order less than 20, $L = 10$ in the recognition algorithm should suffice.

5.2.2 Fooling LIN

We can use the conclusions drawn from the previous experiment to construct synthetic examples that will clearly show that it is possible for LIN to make a mistake between two objects. Take a sphere with uniform albedo. Obtain a non-negative image from this sphere using low frequency lighting that is negative at some places. Use this image as the albedo of a second sphere. Under low frequency lighting, LIN cannot distinguish between these two objects but SDP can. The example is shown in figure 1.

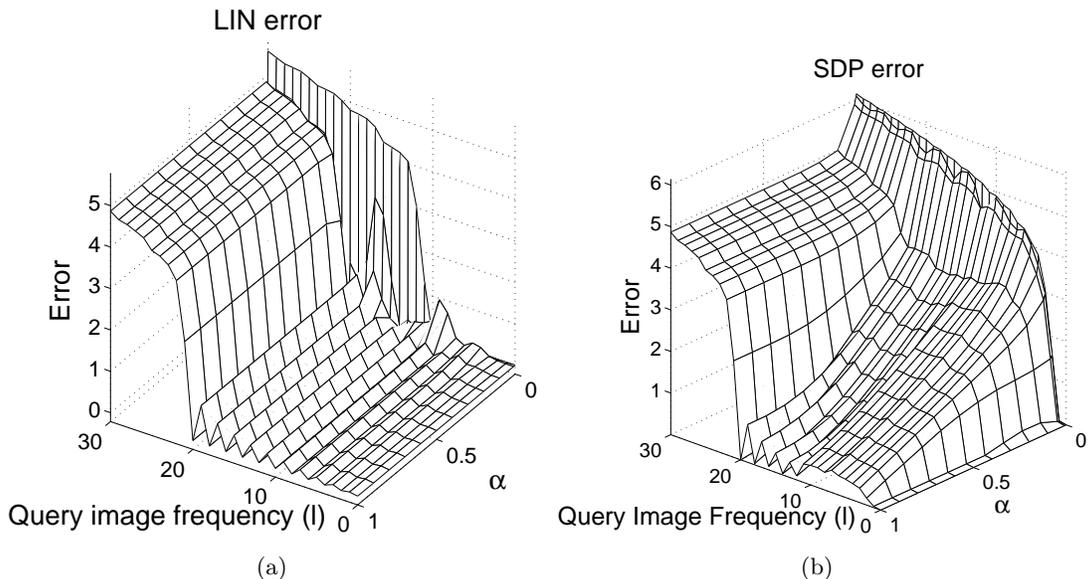


Figure 3: Error vs query image frequency and model specularity: (a) Non-negativity not imposed (LIN), (b) Non-negativity imposed (SDP)

5.3 Experiments on Real Images

Experiments were performed on two real objects to support the results of the synthetic experiments. The first object was a shiny rubber ball, chosen because it was easy to construct its structural model. The second object was a painted ceramic salt shaker. In both these experiments, it is assumed that the the object can be fairly well represented by the mirror + Lambertian model, and that α is constant for the whole surface. These assumptions are not necessary for our method, but they make model construction easier. The first experiment shows that SDP is more robust to noisy models than LIN or a method that simply ignores specularity. First we describe the procedure used for obtaining the surface normals, albedo and α of the surface.

5.3.1 Reflectance Model Construction

The objects used in the experiment were either spherical (ball) or cylindrically symmetric (salt shaker) to enable obtaining surface normals from silhouette images. The full procedure for model construction was as follows :

Surface Normals : First, a silhouette image of the object was taken, by strongly illuminating the background of the object, while keeping the light incident on the object itself to a minimum. On appropriate thresholding, we get a silhouette image. Erroneous pixels in the silhouette image were reduced using a morphological closing operation. The object outline was obtained using the image gradient. This was further smoothed using a Gaussian. For the ball, the center and radius were estimated by fitting a circle to the object outline. These were used to obtain a 3D model and hence surface normals of the sphere. For cylindrically symmetric objects, the axis of symmetry of the object was estimated by fitting a straight line through the outline points. The object outline and the axis of symmetry together gave the 3D structure and hence surface normals of the cylindrically

symmetric object.

Albedo and α : The object was illuminated by a point source of light and its image was captured. To enable reflectance measurement at specular points, 3 images were taken with different exposure times. The first exposure time was set very small (around $\frac{1}{100}$ th– $\frac{1}{20}$ th of a second) so that the specularly caused very little saturation. The second exposure time was set to the largest value so that the Lambertian reflection did not cause saturation (around 10–30 seconds). Finally, the third exposure was set to a suitable intermediate value (around 1 second). 6 such image sets using different directions of the point light source were obtained, taking care to keep the source intensity constant. Next, a single high dynamic range image was constructed from each image set by combining unsaturated pixels from each image, appropriately scaled by the exposure ratio. The region that was not dark in the shortest exposure image was marked as specular. The direction of incident light was calculated using the position of the center of the specular region. The Lambertian region reflectance, corrected for the $\cos(\theta)$ ($\theta =$ angle between surface normal and incident light direction) factor, is proportional to the albedo. For each image set, this is obtained only in the non-specular regions that was illuminated by the points source. To obtain it everywhere on the surface, we calculate the median of the values obtained from the 6 image sets. Since the constant of proportionality doesn't affect our computations, and we simply normalize the obtained albedo by the maximum value. Next, we need an estimate of α . This is the ratio of the specular to total (specular + Lambertian reflection, if the Lambertian albedo was 1) reflection at a point. An estimate of the specular reflection at a point is the sum of grayscale values in the specular area (since lighting is by a point source). Now we need the total Lambertian reflection from this point in all directions, assuming albedo 1. For a sphere, this is simply the sum of the Lambertian reflectance at all points of the illuminated hemisphere, normalized by the albedo. For a cylindrically symmetric object, this has to be estimated using the range of directions of the surface normals that are present. The median of the ratio of specular to total reflection from the 6 image sets gives an estimate of α .

5.3.2 Shiny rubber ball

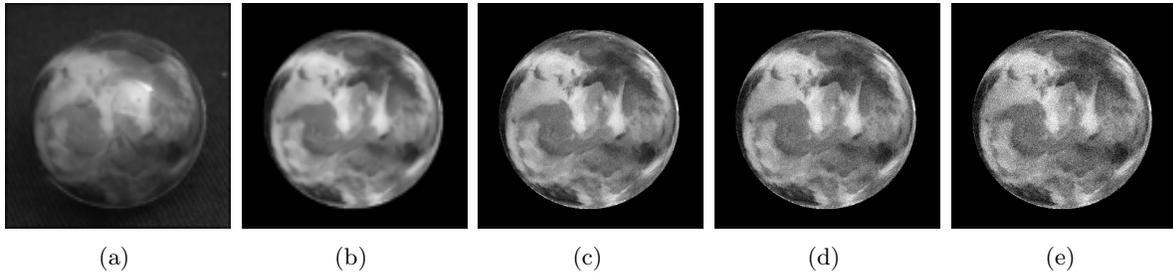


Figure 4: Shiny Rubber Ball: Left to Right: (a) Query Image (b) Measured Albedo (c),(d),(e) Albedo with 8.3%, 15.4% and 24.5% noise levels that fooled LIN, LAMB and SDP respectively.

The experiment consisted of comparing the error difference when lighting is recovered by the correct model, and when it is recovered by a uniform albedo model. The albedo and α for the ball were measured in a separate experiment. A value of $\alpha = 0.04$ was obtained. Next, we repeated the experiment using noisy versions of the albedo, to find out which method gets confused first. For comparison, the same experiments were also done assuming a Lambertian

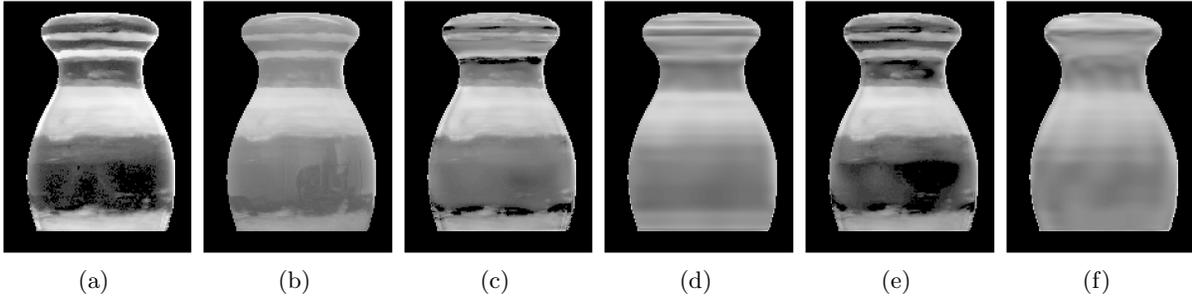


Figure 5: Ceramic shaker (LIN vs SDP): Left to Right: (a) Measured Albedo. (b) Query Image. (c) Best image using correct model and LIN (Error = 8.0%). (d) Best image using uniform model and LIN (Error = 8.7%). (e) Best image using correct model and SDP (Error = 10.3%). (f) Best image using uniform model and SDP (Error = 11.8%). SDP has a higher error difference than LIN between correct and wrong models, and so should be harder to fool.

model (LAMB), not using the non-negativity constraint and only using a 9D subspace ($L = 2$). The results are shown in table 2 and the corresponding images are shown in figure 4. The error difference with SDP is larger than that with LIN or LAMB. Since it is a specular object, LAMB has a much higher error than LIN or SDP, even with the correct model. A more noisy model is needed to confuse SDP, as compared to LIN or LAMB.

| Method | LIN | SDP | LAMB |
|-----------------------------|------|-------|-------|
| Correct Model Error | 6.99 | 10.34 | 19.48 |
| Uniform Model Error | 9.12 | 19.18 | 22.26 |
| Noise needed to fool method | 8.3% | 24.5% | 15.4% |

Table 2: RMS error obtained when matching the query image to the correct model and a uniform albedo model, using various methods. Gaussian noise (with σ as a percentage of correct albedo standard deviation) is added to the model albedo till the method gives the same error as that for the uniform model. LIN and LAMB are fooled more easily. Also, since it is a specular object, LAMB has a large error even for the correct albedo.

5.3.3 Ceramic shaker

The albedo pattern and α of the shaker were obtained exactly as that of the ball. The measured value of α was 0.0031. Although this does not seem much, the shaker was specular enough so that the entire room could be seen in it under normal room lighting. The 3D model of the shaker was also obtained using its cylindrical symmetry. A query image was obtained by using almost uniform lighting (to give a low frequency image). The errors obtained when lighting recovery was attempted using LIN and SDP for a uniform model, as well as the measured model are shown in figure 5, along with the generated images. We can see that SDP produces an error difference between the correct and uniform (incorrect) models that is larger than that produced by LIN. We can expect that in this case too, a noisy model will fool LIN more easily than SDP.

6 Conclusion and Future Work

We have introduced a new method for enforcing the non-negativity constraint of light in lighting recovery and object recognition. The method is based on the extension of Szego's eigenvalue distribution theorem to spherical harmonics. It is exact and faster than the previous method. From the experiments on synthetic as well as real data, it is clear that the non-negativity constraint is indeed helpful in recognition. The SDP method enables better discrimination between the correct and incorrect models, especially in the presence of noise.

The non-negativity constraint enables better recognition by reducing the space of lighting conditions that are possible. We would like to theoretically quantify this reduction in the space of images. Also, we would like to apply this method to various other problems, like environment map creation in computer graphics.

Appendix

Appendix A Spherical Harmonics

Appendix A.1 Complex Spherical Harmonics

The *Surface Spherical Harmonics* [5] are a set of orthonormal basis functions for the set of all functions $f(u)$ defined on the surface of the sphere S^2 , similar to Fourier basis functions on the Circle. They are denoted by Y_{lm} , with $l = 0, 1, 2, \dots$ and $-l \leq m \leq l$.

$$Y_{lm}(u) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_{lm}(\cos \theta) e^{im\phi}$$

$$\theta \in [0, \pi], \quad \phi \in [-\pi, \pi], \quad u = (\theta, \phi) \quad (17)$$

where P_{lm} are the *Associated Legendre Functions*, defined by Rodriguez Formula as:

$$P_{lm}(z) = \frac{(1-z^2)^{m/2}}{2^l l!} \frac{d^{l+m}}{dz^{l+m}} (z^2 - 1)^l$$

A useful relation of spherical harmonics is :

$$\bar{Y}_{lm} = (-1)^m Y_{l,-m} \quad (18)$$

A function $f(u)$ can be expanded in terms of spherical harmonic basis functions, analogous to a Fourier Series expansion of $f(\theta)$ on the circle S .

$$f(u) = \sum_{l=0}^{\infty} \sum_{m=-l}^l f_{lm} Y_{lm}(u) \quad (19)$$

The spherical harmonic coefficients f_{lm} can be computed as

$$f_{lm} = \int_{S^2} f(u) \bar{Y}_{lm}(u) d\sigma(u)$$

$$d\sigma(u) = \sin \theta d\theta d\phi \quad (20)$$

The *Clebsch-Gordon (CG) coefficients* $C(l_1 l_2 l_3; m_1 m_2 m_3)$ [13] are real numbers which appear in many relations involving spherical harmonics. They are zero unless all of these conditions are satisfied :

Next, we can calculate the real Gaunt (R-Gaunt) coefficients [7], or the coupling coefficients for the real spherical harmonics. These can be defined as :

$$G_R(l_1 l_2 l_3; m_1 m_2 m_3) = \int_{S^2} X_{l_1 m_1}(u) X_{l_2 m_2}(u) X_{l_3 m_3}(u) du \quad (26)$$

It is immediately clear that the R-Gaunt coefficients are invariant to any permutation of the $(l_i m_i)$ pairs. Hence, without loss of generality, we can assume that $|m_1| \geq |m_2| \geq |m_3|$. This choice makes their calculation from Gaunt coefficients much simpler. Using the definitions of real spherical harmonics, we can calculate R-Gaunt coefficients as in equation (27).

$$G_R(l_1 l_2 l_3; m_1 m_2 m_3) = \begin{cases} 2G(l_1 l_2 l_3; m_2 + m_3 m_2 m_3) \mathbf{Re}[\bar{U}_{l_1 m_2 + m_3}^{m_1} U_{l_2 m_2}^{m_2} U_{l_3 m_3}^{m_3}] + \\ 2G(l_1 l_2 l_3; m_2 - m_3 m_2 - m_3) \mathbf{Re}[\bar{U}_{l_1 m_2 - m_3}^{m_1} U_{l_2 m_2}^{m_2} U_{l_3 - m_3}^{m_3}] & \text{all } m_i \neq 0 \\ 2G(l_1 l_2 l_3; m_2 m_2 0) \mathbf{Re}[\bar{U}_{l_1 m_2}^{m_1} U_{l_2 m_2}^{m_2}] & m_3 = 0 \\ \delta_{m_1 0} G(l_1 l_2 l_3; 000) & m_2 = m_3 = 0 \end{cases} \quad (27)$$

Appendix B Semidefinite Programming

The basic problem of semidefinite programming [16] in *primal standard form* can be stated as :

$$\begin{aligned} \min_X \quad & C \bullet X \\ \text{subject to} \quad & A_i \bullet X = b_i, \quad i = 1, \dots, m. \\ \text{and} \quad & X \succeq 0 \end{aligned} \quad (28)$$

Here $X \in S\mathbb{R}^{n \times n}$ (the space of real symmetric matrices of size $n \times n$) is the variable and $A_i \in S\mathbb{R}^{n \times n}, C \in S\mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^m$ are given. $A \succ (\succeq) B$ means that $A - B$ is a positive (semi)definite matrix.

The \bullet operator is the inner product operator of two matrices, defined by $A \bullet B = \text{tr}(A^T B)$. The associated norm is the Frobenius norm $\|A\|_F = (A \bullet A)^{\frac{1}{2}}$.

A more convenient formulation of SDP is the *dual standard form* :

$$\begin{aligned} \max_{y, S} \quad & b^T y \\ \text{subject to} \quad & \sum_{i=1}^m y_i A_i + S = C \\ \text{and} \quad & S \succeq 0 \end{aligned} \quad (29)$$

where $y \in \mathbb{R}^m$ and $S \in S\mathbb{R}^{n \times n}$ are the variables and. This is a convenient formulation because we can get rid of the slack variable S and write the problem as :

$$\begin{aligned} \max_{y, S} \quad & b^T y \\ \text{subject to} \quad & \sum_{i=1}^m y_i A_i \preceq C \end{aligned} \quad (30)$$

This is the Linear Matrix Inequality form, which occurs commonly in applications. Note that our constraint $T_L(a) \succeq 0$ is also of this form, with $C = 0$.

A closely related problem is *Second Order Cone Programming* (SOCP). Here, the constraint, instead of positive definiteness of a matrix, is $t \geq \|y\|$, where t is a scalar and $y \in \mathbb{R}^n$. This makes the vector $\begin{pmatrix} t \\ y \end{pmatrix}$ lie inside a second order or Lorentz cone in the space \mathbb{R}^{n+1} . This is a useful formulation because algorithms for solving SOCP problems are faster than those for SDP problems. The following interesting property enables the conversion of some SDP problems into SOCP ones.

Theorem 8 (Schur Complement). *Suppose $U = \begin{bmatrix} A & B \\ B^T & C \end{bmatrix}$ with A and C symmetric and $A \succ 0$. Then,*

$$U \succ 0 (\succeq 0) \quad \text{iff} \quad S = C - B^T A^{-1} B \succ 0 (\succeq 0)$$

The matrix S is called the *Schur complement* of A . For a proof, see [16]. Using this theorem, we can see that

$$\begin{bmatrix} tI & y \\ y^T & t \end{bmatrix} \succeq 0 \text{ is equivalent to } t \geq \|y\|$$

Appendix C Theorem of Carathéodory

This is an important theorem in the study of the Trigonometric Moment Problem.

Theorem 9. [4](Section 4.1) *Let c_1, c_2, \dots, c_n be given complex constants not all zero, $n > 1$. There exists an integer m , $1 \leq m \leq n$, and certain real constants ρ_p, θ_p ; $p = 1, 2, \dots, m$, such that $\rho_p > 0$, $e^{i\theta_p} \neq e^{i\theta_q}$ if $p \neq q$, and*

$$c_\nu = \sum_{p=1}^m \rho_p e^{i\nu\theta_p} \tag{31}$$

The integer m and the constants ρ_p and $e^{i\theta_p}$ are uniquely determined.

If we define $c_{-\nu} := \bar{c}_\nu$ and c_0 is such that the Toeplitz matrix formed from c_ν $\{\nu = -n, -n+1, \dots, n\}$ is positive semidefinite, then the equation (31) is valid for all $\nu = -n, -n+1, \dots, n$.

References

- [1] R. Basri and D. W. Jacobs, *Lambertian Reflectance and Linear Subspaces*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 25, no. 2. (Feb. 2003), pp. 218-233.
- [2] P. Belhumeur, D. Kriegman, *What is the Set of Images of an Object Under All Possible Lighting Conditions?*, International Journal of Computer Vision, 28(3), 1998, pp. 245-260.
- [3] K.D. Gremban and K. Ikeuchi, *Planning Multiple Observations for Specular Object Recognition*, IEEE Conference on Robotics and Automation, Vol. 2, May, 1993, pp. 599604.
- [4] U. Grenander, G. Szego, *Toeplitz Forms and their Applications*, University of California Press, 1958
- [5] H. Groemer, *Geometric Applications of Fourier Series and Spherical Harmonics*, Cambridge University Press, 1996

- [6] P. Hallinan, *A Low-Dimensional Representation of Human Faces for Arbitrary Lighting Conditions*, Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, pp. 995-999, 1994.
- [7] Herbert H.H. Homeier, E. Otto Steinborn, *Some properties of the coupling coefficients of real spherical harmonics and their relation to Gaunt coefficients*, Journal of Molecular Structure (Theochem), Vol. 368 (1996) pp.31-37.
- [8] H. Murase and S. Nayar, *Visual Learning and Recognition of 3D Objects from Appearance*, International Journal of Computer Vision, vol. 14, no. 1, pp. 5-25, 1995.
- [9] K. Okikiolu, *The Analogue of the Strong Szego Limit Theorem on the 2 and 3-Dimensional Spheres*, Journal of the American Mathematical Society, Vol 9, No. 2 (April 1996), pp. 345-372.
- [10] Margarita Osadchy, David Jacobs, Ravi Ramamoorthi *Using Specularities for Recognition*, Proceedings of IEEE International Conference on Computer Vision, 2003.
- [11] Ravi Ramamoorthi, *A Signal-Processing Framework for Forward and Inverse Rendering*, PhD Thesis, Stanford University, Aug 2002.
- [12] Ravi Ramamoorthi and Pat Hanrahan. *On the relationship between radiance and irradiance: determining the illumination from images of a convex Lambertian object*, Journal of Optical Society of America, 18(10):24482459, 2001
- [13] M. E. Rose, *Elementary Theory of Angular Momentum*, John Wiley & Sons, 1957
- [14] K. Sato, K. Ikeuchi, and T. Kanade, *Model-Based Recognition of Specular Objects Using Sensor Models*, Processing of IEEE Workshop on Directions in Automated CAD-Based Vision, June, 1991, pp. 210.
- [15] K. K. Thornber and D. W. Jacobs, *Broadened, Specular Reflection and Linear Subspaces*, Technical Report, NEC Research Institute, Inc., Princeton, NJ.
- [16] M.J. Todd, *Semidefinite Optimization*, Acta Numerica, Vol 10 (2001), pp. 515-560.
- [17] R.H. Tutuncu, K.C. Toh, M.J. Todd, *Solving semidefinite-quadratic-linear programs using SDPT3*, Mathematical Programming, Vol 95, No. 2 (Feb. 2003), pp. 189-217.
- [18] <http://control.ee.ethz.ch/~joloef/yalmip.msql>
- [19] A. Yuille, D. Snow, R. Epstein, and P. Belhumeur, *Determining Generative Models of Objects Under Varying Illumination: Shape and Albedo from Multiple Images Using SVD and Integrability*, International Journal of Computer Vision, vol. 35, no. 3, pp. 203-222, 1999.
- [20] L. C. Evans, *Partial Differential Equations*, American Mathematical Society, 1998.
- [21] Erwin Kreyszig, *Introduction to Differential Geometry and Riemannian Geometry*, University of Toronto Press, 1968.