

A Theory Of Frequency Domain Invariants: Spherical Harmonic Identities for
BRDF/Lighting Transfer and Image Consistency

Dhruv Mahajan, *Columbia University*, Ravi Ramamoorthi, *Columbia University*,
and Brian Curless, *University of Washington*

Abstract

This paper develops a theory of frequency domain invariants in computer vision. We derive novel identities using spherical harmonics, which are the angular frequency domain analog to common spatial domain invariants such as reflectance ratios. These invariants are derived from the spherical harmonic convolution framework for reflection from a curved surface. Our identities apply in a number of canonical cases, including single and multiple images of objects under the same and different lighting conditions. One important case we consider is two different glossy objects in two different lighting environments. For this case, we derive a novel identity, independent of the specific lighting configurations or BRDFs, that allows us to directly estimate the fourth image if the other three are available. The identity can also be used as an invariant to detect tampering in the images.

While this paper is primarily theoretical, it has the potential to lay the mathematical foundations for two important practical applications. First, we can develop more general algorithms for inverse rendering problems, which can directly relight and change material properties by transferring the BRDF or lighting from another object or illumination. Second, we can check the consistency of an image, to detect tampering or image splicing.

Index Terms

Frequency Domain Invariants, Spherical harmonic identities, Convolution, Inverse rendering, Relighting, Tampering, Image Forensics.

I. INTRODUCTION

In this paper, we develop a theory of frequency domain invariants in computer vision. This new class of invariants can address complex materials in complex lighting conditions, for applications like inverse rendering, image forensics and relighting. Our work extends the widely used spatial domain theory of invariants [NRN03], [NB96], [JSY03], [DYW05], developing the frequency domain analogs.

Our analysis is based on the spherical convolution theorem for reflection of distant lighting from curved objects [BJ03], [RH01]. This theory shows that the reflected light in the frequency domain is a product of the spherical harmonic coefficients of the lighting signal and BRDF filter. This product relationship is similar to the spatial product of albedo and irradiance for textured objects, that has been the basis for a variety of spatial domain invariants such as reflectance ratios [NB96] and photometric invariants [NRN03]. By exploiting the product form

of the frequency domain relations, we can derive analogous frequency-domain invariants, but now for general lighting and reflectance properties.

This paper also describes one of the first applications in computer vision of the spherical harmonic analysis for complex non-Lambertian materials. In earlier work, [BJ01], [BJ03], [RH01] have shown that the set of all Lambertian reflectance functions (the mapping from surface normals to intensities) lies close to a 9D linear subspace for convex objects of known shape lit by complex distant illumination. This result often enables computer vision algorithms, previously restricted to point sources without attached shadows, to work in general complex lighting. There has been considerable work on novel algorithms for lighting-insensitive recognition, photometric stereo and relighting [BJ03], [BJ01], [HS05], [SFB03], [ZWS05], [WLH03] In graphics, the general convolution formulae have been used for rendering with environment maps [RH02], and insights have been widely adopted for forward and inverse rendering (e.g., [RH01], [SKS02]).

However, there has been relatively little work in vision on using the convolution formulae for glossy objects, even though the frequency analysis [RH01] applies for general materials. The main goal of this paper is to derive new formulae and identities for direct frequency domain spherical (de)convolution. Specifically, we make the following theoretical contributions:

Derivation of New Frequency Domain Identities: Our main contribution is the derivation of a number of new theoretical results, involving a class of novel frequency domain identities. We study a number of setups, including single (sections IV and V) and multiple (section VI) images under single and multiple lighting conditions. For example, one important case we consider (section VI-C) is that of two different glossy¹ materials in two different lighting environments (figure 1). Denote the spherical harmonic coefficients by $B_{lm}^{light,material}$, where the subscripts refer to the harmonic indices, and the superscripts to the lighting (1 or 2) and object or material (again 1 or 2). We derive an identity for the specular component, $B_{lm}^{1,1} B_{lm}^{2,2} = B_{lm}^{1,2} B_{lm}^{2,1}$, directly from the properties of convolution, *independent* of the specific lighting configurations or BRDFs.

Analogy between Spatial and Frequency Domain Invariants: By definition, invariants are insensitive to certain appearance parameters like lighting. They usually transform images

¹Parts of the theory (in sections IV and VI) address only purely specular (or purely Lambertian) objects. However, as discussed in the paper and shown in our results, the theory and algorithms can be adapted in practice to glossy objects having both diffuse and specular components. Hence, we use the term “glossy” somewhat loosely throughout the paper.

to a simple feature space where more accurate algorithms can be developed for the task at hand. Invariants have been previously used mostly for material and lighting insensitive recognition [NRN03], [NB96]. There has been a substantial body of previous work in developing spatial domain invariants. Jin et al. [JSY03] derive a constraint on the rank of the radiance tensor field to do stereo reconstruction for non-Lambertian objects. Nayar and Bolle [NB96] compute the ratio of intensities at adjacent pixels to derive lighting independent reflectance ratios. Davis et al. [DYW05] derive a similar BRDF independent ratio. Narsimhan et al. [NRN03] consider a summation of multiple terms (diffuse plus specular), where each term is a product of material and geometry. However most of the above methods are limited to point sources [NRN03], [DYW05] or consider textured Lambertian objects only [NB96].

We show (section VII) that the class of identities derived in this paper can be considered the analog in the frequency domain of fundamental spatial domain invariants. We consider curved homogeneous glossy objects instead of textured Lambertian objects. We also assume radially symmetric BRDFs, a good approximation for most specular reflectance. Moreover, we consider general complex lighting; by contrast, much of the previous spatial domain theory is limited to single point sources. Conversely, while our identities operate globally needing the full range of reflected directions, spatial domain invariants involve mostly local pixel-based operations.

Analysis of Diffuse Irradiance in Reflected Parameterization: Another major contribution of the paper is the analysis of diffuse irradiance in the reflected parameterization. This analysis allows us to study objects with both diffuse and specular components in a unified framework. We show that even with the parameterization by reflected direction, the effects of diffuse irradiance are limited to low frequencies. To our knowledge, this is the first such combined diffuse plus specular theory and is likely to have broader implications for other problems in vision.

The theory and novel identities presented in the paper have potential applications in many areas of vision and graphics like inverse rendering, consistency checking, BRDF-invariant stereo and photometric stereo or lighting-insensitive recognition. In particular, this paper is motivated by the following three important practical applications, and seeks to lay the mathematical foundations in these areas.

Inverse Rendering: Estimation of the BRDF and lighting has been an area of active research in vision and graphics. Inverse rendering deals with measuring these rendering attributes from

photographs. Rendering synthetic images by using these measurements from real objects greatly enhances the visual realism of the rendered images. For example, we estimate illumination from a single image of a glossy material with known BRDF. By the convolution theorem, a glossy material will reflect a blurred version of the lighting. It is appealing to sharpen or deconvolve this by dividing in the frequency domain by the spherical harmonic coefficients of the BRDF. The basic formula is known [RH01], but cannot be robustly applied, since BRDF coefficients become small at high frequencies. Our contribution is the adaptation of Wiener filtering [GW03], [Wie42] from image processing to develop robust deconvolution filters (figures 4 and 12). We are able to amplify low frequencies to recover the lighting and reduce noise simultaneously.

BRDF/Lighting Transfer: Given images of an object under a sparse set of lighting conditions, relighting it with novel lighting is an interesting problem. Current methods [MWLT00], [MG97], [SSI99] require explicit estimation of lighting and BRDF from the images of a scene. [ZWS05], [WLH03] also use spherical harmonics for face relighting, but they assume Lambertian faces. This paper presents more general algorithms, which directly relight and change material properties by transferring the BRDF or lighting from another object or illumination. For example, consider a simple case where we have images of two different objects in two different lighting conditions. We derive an identity that enables us to render the fourth light/BRDF image, given the other three, *without explicitly estimating any lighting conditions or BRDFs*. A common example (figure 1) is when we observe two objects in one lighting, and want to insert the second object in an image of the first object alone under new lighting. It is difficult to apply conventional inverse rendering methods in this case, since none of the illuminations or BRDFs are known.

Image Consistency Checking and Tampering Detection: The final, newer application, is to verify image consistency and detect tampering (Johnson and Farid [JF05], Lin et al. [LWTS05]). The widespread availability of image processing tools enables users to create “forgeries”, e.g., by splicing images together (one example is shown in figure 13). Moreover, watermarking is not usually a viable option in many applications, such as verifying authenticity for news reporting. However, (in)consistencies of lighting, shading and reflectance can also provide valuable clues. Most previous work has focused on checking consistency at a signal or pixel level, such as the camera response [LWTS05], or wavelet coefficients (Ng et al. [NCS04]). But most of these methods do not exploit consistencies of lighting, shading and reflectance. Johnson and Farid [JF05] detect inconsistencies in lighting to expose forgeries, but their method is limited to




Fig. 1. One application of our framework. We are given real photographs of two objects of known geometry (shown in inset; note that both objects can be arbitrary, and one of them is a sphere here only for convenience). The two objects have different (and unknown) diffuse and specular material properties. Both objects are present in the first image under complex lighting, but the cat is not available in the second image, under new lighting. *Unlike previous methods, none of the lighting conditions or BRDFs are known* (lightings on left shown only for reference). Our method enables us to render or relight the cat, to obtain its image in lighting 2 (compare to actual shown on the right). This could be used for example to synthetically insert the cat in the second image.

point light sources. This paper takes an important first step in laying the theoretical foundations for this new research direction, by deriving a new class of identities which can be checked to detect tampering and consistency of lighting and shading in a complex lighting environment. A limitation of our approach is that our identities require the knowledge of 3D model/geometry of the object, though such geometry could be available through prior acquisition or estimated from the images, e.g., based on known shape distributions [BV99].

The rest of this paper is organized as follows. Section 2 briefly explains the spherical convolution and signal processing framework. Section 3 demonstrates the use of deconvolution to estimate lighting. In sections 4 and 5, we introduce identities for the simple case of a single image of an object. Section 6 derives more identities for the case of multiple images. In section 7 we discuss the implications of our theory and its relation to spatial domain invariants. Section 8 gives experimental validation of our theory and shows potential applications. Finally, we conclude our discussion in section 9 and talk about the future research directions that this work makes possible. This paper is an extended and detailed version of a paper that was presented at ECCV

2006 [MRC06].

II. BACKGROUND

We now briefly introduce the spherical convolution and signal-processing framework [BJ03], [RH01] needed for our later derivations. We start with the Lambertian case,

$$B(\mathbf{n}) = \int_{S^2} L(\omega) \max(\mathbf{n} \cdot \omega, 0) d\omega, \quad (1)$$

where $B(\mathbf{n})$ denotes the reflected light as a function of the surface normal. B is proportional to the irradiance (we omit the albedo for simplicity), and $L(\omega)$ is the incident illumination. The integral is over the sphere S^2 , and the second term in the integrand is the *half-cosine* function. The equations in this paper do not explicitly consider color; the (R,G,B) channels are simply computed independently. A similar mathematical form holds for other radially symmetric BRDFs, such as the Phong model for specular materials. In the specular² case, we reparameterize by the reflected direction \mathbf{R} (the reflection of the viewing ray about the surface normal), which takes the place of the surface normal. For the Phong model, the reflection equation becomes:

$$B(\mathbf{R}) = \frac{s+1}{2\pi} \int_{S^2} L(\omega) \max(\mathbf{R} \cdot \omega, 0)^s d\omega, \quad (2)$$

where s is the Phong exponent, and the BRDF is normalized (by $(s+1)/2\pi$).

If we expand in spherical harmonics $Y_{lm}(\theta, \phi)$, using spherical coordinates $\omega = (\theta, \phi)$, \mathbf{n} or $\mathbf{R} = (\alpha, \beta)$, and $\rho(\theta)$ for the (radially symmetric) BRDF kernel, we obtain

$$L(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l L_{lm} Y_{lm}(\theta, \phi) \quad B(\alpha, \beta) = \sum_{l=0}^{\infty} \sum_{m=-l}^l B_{lm} Y_{lm}(\alpha, \beta) \quad \rho(\theta) = \sum_{l=0}^{\infty} \rho_l Y_{l0}(\theta). \quad (3)$$

It is also possible to derive analytic forms and good approximations for common BRDF filters ρ . For the Lambertian case, almost all of the energy is captured by $l \leq 2$. For Phong and Torrance-Sparrow models of specular reflection, good approximations [RH01] are Gaussians: $\exp[-l^2/2s]$ for Phong, and $\exp[-(\sigma l)^2]$ for Torrance-Sparrow, where σ is the surface roughness parameter in the Torrance-Sparrow model, and s is the Phong exponent.

In the angular (vs. angular frequency) domain, equations 1 and 2 represent rotational convolution of lighting with BRDF. The BRDF can be thought of as the filter, while the lighting

²“Specular” will always be used to mean generally glossy, including but not restricted to mirror-like.

is the input signal. This allows us to relate them multiplicatively in the angular frequency domain (convolution theorem). In the frequency domain, the reflected light B is given by a simple product formula or spherical convolution (see [BJ03], [RH01] for the derivation and an analysis of this convolution),

$$\boxed{B_{lm} = \Lambda_l \rho_l L_{lm} = A_l L_{lm}}, \quad (4)$$

where for convenience, we define the normalization constant Λ_l as

$$\Lambda_l = \sqrt{\frac{4\pi}{2l+1}} \quad A_l = \Lambda_l \rho_l. \quad (5)$$

It is also possible to extend these results to non-radially symmetric general isotropic BRDFs [RH01]. For this case, we must consider the entire 4D light field, expressed as a function of both orientation and outgoing direction,

$$B_{lmpq} = \Lambda_l \rho_{lq,pq} L_{lm}, \quad (6)$$

where the reflected light field is now expanded in a mixed basis of representation matrices and spherical harmonics, and has four indices because it is a 4D quantity. The 3D isotropic BRDF involves an expansion over both incoming and outgoing directions. The new indices p and q correspond to the spherical harmonic indices for the expansion over outgoing angles (analogous to the indices l and m used for the lighting).

The remainder of this paper derives new identities and formulae from equation 4, $B_{lm} = A_l L_{lm}$. Most glossy BRDFs (such as Torrance-Sparrow) are approximately radially symmetric, especially for non-grazing angles of reflection [RH01], [RH02]. Most of the theory in this paper also carries over to general isotropic materials, as per equation 6, if we consider the entire light field. Another reason to focus on equation 4, is that it is simple and allows practical spherical harmonic computations from only a *single image*—a single view of a sufficiently curved object (assuming a distant viewer) sees all reflected directions.³

³In case we do not have the full range of normals, we can use multiple cameras. As we move the camera (viewer) the same point on the object now corresponds to a different reflected direction. Hence we can get all the reflected directions even if the object has only a partial set of normals by the careful placement of cameras.

III. KNOWN BRDF: DECONVOLUTION TO ESTIMATE LIGHTING

Lighting estimation is a specific example of the general inverse rendering problem. Given a single image and BRDF of known geometry and homogenous material, we want to estimate the directional distribution of the incident light. This information can then be used to insert new objects in the scene, alter the lighting of the object or check lighting consistency between two objects. Since reflected light is a spherical convolution of lighting and BRDF, it makes sense to deconvolve it to estimate lighting. We present a deconvolution algorithm for curved surfaces under complex lighting. Section III-A describes the basic deconvolution idea and introduces an ideal deconvolution filter. We then discuss the properties of this filter for Phong-like BRDFs in section III-B. Section III-C describes the Wiener filter used to regularize the inverse filter so that it can be used for practical purposes. Finally, we show the results of applying this filter in section III-D.

A. Deconvolution - Basic Idea

Given a single image of a curved surface, we can map local viewing directions to the reflected direction, determining $B(\mathbf{R})$, and then B_{lm} by taking a spherical harmonic transform. If the material includes a diffuse component as well as specular, we use the dual lighting estimation algorithm of Ramamoorthi and Hanrahan [RH01], which estimates the specular B_{lm} consistent with the diffuse component. As per equation 4, B_{lm} will be a blurred version of the original lighting, filtered by the glossy BRDF.

From equation 4 in the spherical harmonic domain, we derive

$$L_{lm} = \frac{B_{lm}}{A_l} = A_l^{-1} B_{lm}, \quad (7)$$

where the last identity makes explicit that we are convolving (in the angular domain) with a new radially symmetric kernel A_l^{-1} , which can be called the inverse, sharpening or deconvolution filter. A_l^{-1} effectively amplifies high frequencies to recover blurred out details.

B. Analysis of Inverse Phong Filter

We now discuss the properties of the angular form of the inverse filter. Surprisingly, not much work has been done to analyze this filter in detail. For simplicity, we will use the Fourier

transform rather than spherical harmonics. We will illustrate that the properties discussed in the Fourier domain are also valid for spherical harmonics.

We use the inverse Phong filter for our analysis. As mentioned earlier, a Gaussian $\exp[-l^2/2s]$ gives a good approximation for Phong reflection, where s is the Phong exponent. So, the inverse Phong filter can be approximated by $\exp[l^2/2s]$. Note that this analysis also applies to the Torrance-Sparrow approximation by substituting $s = 1/2\sigma^2$. Since the filter becomes large at high frequencies, leading to amplification of the noise, we need to truncate it first to a cut-off frequency r . The inverse Fourier transform of this truncated filter is

$$f(x; r, s) = \int_{-r}^r e^{\frac{u^2}{2s}} e^{2\pi i x u} du \quad (8)$$

Putting $u = \sqrt{2s}v$

$$f(x; r, s) = \sqrt{2s} \int_{-\frac{r}{\sqrt{2s}}}^{\frac{r}{\sqrt{2s}}} e^{v^2} e^{2\sqrt{2s}\pi i v x} dv$$

$$f(x; r, s) = \sqrt{2s} g(\sqrt{2s}x, \frac{r}{\sqrt{2s}}), \quad (9)$$

$$g(x; k) = \int_{-k}^k e^{t^2} e^{2\pi i t x} dt \quad (10)$$

$g(x; k)$ is the inverse Fourier transform of the canonical filter $\exp[t^2]$ truncated at k and is independent of Phong exponent s . Going from f to g is just the application of the Fourier Scale Theorem. Let $H(u)$ be the Fourier transform of $h(x)$.

$$h(x) \leftrightarrow H(u)$$

Then, the Fourier scale theorem states that

$$h(ax) \leftrightarrow \frac{1}{|a|} H\left(\frac{u}{a}\right)$$

In our case $a = \frac{1}{\sqrt{2s}}$. The frequencies u of the canonical filter $\exp[u^2]$ get scaled by $\frac{1}{\sqrt{2s}}$. By the Fourier scale theorem, this means that x gets scaled by $\sqrt{2s}$ in the spatial domain. Hence $f(x; r, s)$ is just the spatially scaled version of $g(x; k)$. $g(x; k)$ can be further analyzed to give

$$g(x; k) = \frac{2\pi e^{k^2}}{k} n\left(kx, \frac{\pi}{k}\right), \quad (11)$$

$$n(\alpha, \beta) = \int_{\alpha}^{\infty} e^{\beta^2(\alpha^2 - u^2)} \sin(2\pi u) du \quad (12)$$




Fig. 2. **Left:** f , g and n functions for two different values of frequency cut-off r . As r increases, f becomes more and more oscillatory with period $\frac{1}{r}$. The period of n however does not change. **Right:** (g) shows that the amplitude of f (at $x = 0$) increases exponentially with r^2 . The log-log plot (h) of amplitude of f vs. x is a straight line with slope -1, showing that the filter falls off as $\frac{1}{x}$.

A detailed derivation is given in Appendix A. Plots for $f(x; r, s)$, $g(x; k)$ and $n(\alpha, \beta)$ are shown in figure 2(a)-(f). Here k and r are related to each other by $k = \frac{r}{\sqrt{2s}}$.

$n(\alpha, \beta)$ can be considered as a normalized form of the inverse filter and is independent of both Phong exponent s and cut-off frequency r . We now make some important empirical observations about $n(\alpha, \beta)$. For fixed β , it has the shape of a damped sinusoid with a period of 1 in α . This insight comes from a large number of plots, only two representative examples of which we have shown (figure 2(c,f)). We have also found out that the variation in peak amplitude of n (at $\alpha = 0$) is small for different β . Moreover, the amplitude of n falls off as $\frac{1}{\alpha}$.

We now discuss some important properties of $f(x; r, s)$.

1) *Periodicity*: : Since n has period 1 in α , $g(x; k)$ has period $\frac{1}{k}$ (from equation 11). From equation 9,

$$\begin{aligned}
 \text{Period of } f(x; r, s) &= \frac{1}{\sqrt{2s}} * \text{Period of } g(x; k) \\
 &= \frac{1}{\sqrt{2s}} * \frac{1}{k} \\
 &= \frac{1}{\sqrt{2s}} * \frac{\sqrt{2s}}{r} \\
 &= \frac{1}{r}
 \end{aligned} \tag{13}$$

So as cut-off frequency r increases, the filter becomes more and more oscillatory (figure 2(a,d)).

2) *Peak Amplitude*: : We now discuss the effect of cut-off frequency r and Phong exponent s on the peak amplitude of the filter f (which occurs at $x = 0$). From equations 9 and 11, the peak amplitude introduced due to frequency cut-off r and Phong exponent s is $\sqrt{2s} \frac{e^{k^2}}{k} * \text{Peak Amplitude of } n$. Since the variation in peak amplitude of n is small for different β , we can neglect it in comparison to other terms. Hence the peak amplitude is approximately

$$\sqrt{2s} \frac{e^{k^2}}{k} \approx \frac{2se^{\frac{r^2}{2s}}}{r} \tag{14}$$

As r increases, the peak amplitude grows almost exponentially, as can be seen from figure 2(g).

3) *Amplitude Fall-off*: : The amplitude fall-off in x of $f(x; r, s)$ is the same as that of $n(\alpha, \beta)$ in α . Figure 2(h) shows that the log-log plot of amplitude fall-off for $f(x; r, s)$ is a straight line with *slope* = -1 . Hence the amplitude of $f(x; r, s)$ falls off as $\frac{1}{x}$.

C. Wiener Regularization

Section III-B shows that it is difficult to apply equation 7 directly and that we need regularization. From the analysis of section III-B and figure 2, it is clear that simply cutting off high frequencies makes the filter more oscillatory and causes an increase in amplitude, resulting in substantial ringing and amplification of noise. Choosing a lower cut-off results in a loss of all the high frequencies, even if they could be somewhat recovered, and we still have substantial residual ringing.

⁴We call them ‘‘periodic’’, in the sense of the periodicity of the damped sinusoid shape they have.




Fig. 3. **Left:** Wiener filters in the frequency domain. (a) shows the Wiener filters for different values of K for a Phong BRDF ($s = 100$). Note that the maximum occurs at $l^* = \sqrt{\log(\frac{1}{Ks})}$, the value being $A_l^{*max} = \frac{1}{2\sqrt{K}}$. $K = 0$ (magenta graph) is the ideal inverse Phong filter and can be approximated by $\exp[l^2/2s]$. Note that this filter attains very large values for large frequencies. (b) shows the result of applying the Wiener filters back to the original Phong filter (blue graph in b) to see which frequencies are let through. Most frequencies are let through without attenuation, while very high frequencies are filtered out. Without Wiener filtering, it should be one everywhere. **Right:** Wiener filters in the angular domain. Note the decrease in oscillations as we increase the value of K (c,e and d,f). Also the period of the filter decreases with increasing Phong exponent s (c,d and e,f).

These types of problems have been well studied in image processing, where a number of methods for deconvolution have been proposed. We adapt Wiener filtering [GW03], [Wie42] for this purpose. Assuming spectrally white noise, we define a new inverse filter,

$$A_l^* = \frac{1}{A_l} \left(\frac{|A_l|^2}{|A_l|^2 + K} \right) = \frac{A_l}{|A_l|^2 + K} \quad L_{lm} = A_l^* B_{lm}, \quad (15)$$

where K is a small user-controlled constant⁵. When $|A_l|^2 \gg K$, the expression in parentheses on the left is close to 1, and $A_l^* \approx A_l^{-1}$. When $|A_l|^2 \ll K$, $A_l^* \approx A_l/K$.

Figure 3 shows the Wiener filter in the spatial and frequency domains for a Phong BRDF with different Phong exponents and K values. Note the smooth fall-off of the filter in the frequency

⁵Wiener filters are widely used in image restoration. There, K is the ratio of the power spectral density (PSD) of the undegraded image (without blur) and the PSD of the noise. In our case, the BRDF plays the role of the blurring function and hence K needs to be defined suitably. In fact, it needs to be estimated from the image. Finding the optimal K value is one of the difficult issues in applying Wiener filters. For now, we do not attempt to estimate it, and instead use a user-specified constant.




Fig. 4. (a): Original synthetic image (Phong BRDF with exponent $s = 100$, diffuse $K_d = 2$ and specular $K_s = 1$) with noise—close examination of (a),(b) will reveal the noise. **Top row:** We recover (c) much of the “ECCV” text in the original lighting (d). Previous techniques (b) can estimate only a blurred result. Note that top and bottom rows show a closeup of the sphere. **Bottom row:** We can use the recovered illumination to create a new rendering of a high-frequency material (f). This compares well with the actual result (g); a previous method (e) creates a very blurred image.

domain. Differentiating equation 15 with respect to A_l reveals that A_l^* attains its maximum value of $A_l^{*max} = \frac{1}{2\sqrt{K}}$ at $A_l^{max} = \sqrt{K}$. We can think of the corresponding value of l at this maximum as the cut-off frequency l^* of the filter. For a Phong filter approximated as $\exp[-l^2/2s]$, this corresponds to the cut-off frequency $l^* = \sqrt{\log(\frac{1}{K^s})}$. $K = 0$ (magenta graph) is the ideal inverse Phong filter and can be approximated by $\exp[l^2/2s]$. Note that this filter attains very large values for large frequencies. For a given s , as K increases, l^* decreases and hence more and more of the higher frequencies get truncated. (b) shows the result of applying the Wiener filters back to the original Phong filter (blue graph in b) to see which frequencies are let through. Most frequencies are let through without attenuation, while very high frequencies are filtered out. Note that without Wiener filtering, it should be equal to 1 everywhere. (c)-(f) shows these filters in the angular domain. Increasing the value of K (c,e and d,f) decreases the amplitude of the filter and makes it less oscillatory, thus decreasing the ringing effects. A similar behavior was noted when choosing a lower cut-off frequency in section III-B; however, with Wiener filtering the ringing is substantially more damped. Increasing the Phong exponent s (c,d and e,f) decreases the periodicity of the filter.




Fig. 5. Lighting estimation for different values of K with the synthetic ECCV sphere of figure 4. The ideal inverse Phong filter is very high for large frequencies resulting in the amplification of noise ($K = 0$ case). As K increases, more and more high frequencies get attenuated, resulting in a decrease in ringing due to noise. However, note that the amount of blurring increases with K .

D. Lighting Estimation in Frequency Domain

The top row in figure 4 shows the results (on the synthetic noisy sphere in (a)) of deconvolution (c)—the “ECCV” text used in the lighting (d) can be recovered fairly clearly. One interesting point is the effect of noise. In our case, the image in a glossy surface is already low pass filtered (because of the BRDF), while any noise usually has much higher frequency content, as seen in the original synthetic image (a). The filter in equation 15 is a low-pass filter, though the cutoff can be set high for low-noise. The amplification at very low frequencies is small. Mid range frequencies are amplified substantially while high frequencies are reduced (because of the inherent regularization). Hence, we can *simultaneously deconvolve the lighting and suppress noise* (compare the noise in (c) with that in (a) or (b)). Figure 5 shows the lighting estimation for different values of K . The ideal inverse Phong filter is very high for large frequencies resulting in the amplification of noise ($K = 0$ case). As K increases, more and more high frequencies get attenuated, resulting in a decrease in ringing due to noise. However, note that the amount of blurring increases with K . Figure 12 shows an application of our method with real data and a geometrically complex object.

It is also interesting to compare our results to previous techniques. Angular-domain approaches are usually specialized to point lights, use higher-frequency information like shadows (Sato et al. [SSI99]) or recover large low-frequency lighting distributions (Marschner and Greenberg [MG97]). Even the more precise dual angular-frequency lighting estimation technique of Ramamoorthi and Hanrahan [RH01] can obtain only a blurred estimate of the lighting (b). The result of applying the latter approach is clearly seen in the bottom row of figure 4, where [RH01] produces a blurred image (e) when trying to synthesize renderings of a new high-frequency

material, while we obtain a much sharper result (f).

IV. THEORETICAL ANALYSIS: SINGLE IMAGE OF ONE OBJECT WITH SPECULAR BRDF

We now carry out our theoretical analysis and derive a number of novel identities for image consistency checking and relighting. We structure the discussion from the simplest case of a single image of one object in this section, to more complex examples in section VI—two objects in the same lighting, the same object in two lighting conditions, and finally two (or many) objects in two (or many) lighting conditions. Deconvolution, discussed in section III is a special single image case where we know the BRDF of the object but lighting is unknown. In this section we discuss the converse case, where the lighting is known, but the BRDF is unknown. The objects are assumed to be purely specular. We then present a general theory for objects with both diffuse and specular components in the next section.

We show that for radially symmetric specular BRDFs, described using equation 4, we can *eliminate* the BRDF to derive an identity that must hold and can be checked *independent of the BRDF*. This is the first of a number of frequency domain identities we will derive in a similar fashion. First, from equation 4, we can write

$$A_l = \frac{B_{lm}}{L_{lm}}. \quad (16)$$

This expression could be used to solve for BRDF coefficients⁶. However, we will use it in a different way. Our key insight is that the above expression is independent of m , and must hold for all m . Hence, we can eliminate the (unknown) BRDF A_l , writing

$$\frac{B_{li}}{L_{li}} = \frac{B_{lj}}{L_{lj}} \quad (17)$$

for all i and j . Moving terms, we obtain our first identity,

$$B_{li}L_{lj} - B_{lj}L_{li} = 0. \quad (18)$$

In effect, we have found a redundancy in the structure of the image, that can be used to detect image tampering or splicing. The lighting L and image B are functions on a 2D

⁶Since natural lighting usually includes higher frequencies than the BRDF, we can apply equation 16 directly without regularization, and do not need to explicitly discuss deconvolution. However, note that we have assumed a purely specular BRDF. The next section does derive a new robust formula (equation 23) for BRDF estimation when *both* diffuse and specular components are present.




Fig. 6. **Left:** The synthetic images used. These correspond to closeups of specular spheres rendered with “ECCV” and “ICCV” lighting. To the naked eye, the two images look very similar. **Middle and Right:** The graphs show that our identity can clearly distinguish consistent image/lighting pairs (lower line) from those where lighting and image are inconsistent (upper line).

(spherical) domain. However they are related by a 1D radially symmetric BRDF, leading to a 1D redundancy⁷, that can be used for consistency checking in equation 18.

To normalize identities in a $[0..1]$ range, we always use an error of the form

$$\text{Error} = \frac{|B_{li}L_{lj} - B_{lj}L_{li}|}{|B_{li}L_{lj}| + |B_{lj}L_{li}|}.$$

There are many ways one could turn this error metric into a binary consistency checker or tamper detector. Instead of arbitrarily defining one particular approach, we will show graphs of the average normalized error for each spherical harmonic order.

Figure 6 applies our theory to synthetic data of an ideal Phong BRDF, with noise added. We show closeups of spheres generated with “ECCV” and “ICCV” lighting. To the naked eye, these look very similar, and it is not easy to determine if a given image is consistent with the lighting. However, our identity in equation 18 clearly distinguishes between consistent (i.e., the image is consistent with the lighting [ECCV or ICCV] it is supposed to be rendered with) and inconsistent illumination/image pairs. As compared to Johnson and Farid [JF05], we handle general complex illumination. Moreover, many of the identities in later sections work directly with image attributes, not even requiring explicit estimation or knowledge of the illumination. However, all our identities require the explicit knowledge of 3D models/geometry of the object.

Figure 7 shows the results on a synthetic sphere with Blinn-Phong BRDF (specular lobe). In general, the convolution theorem of equation 4 does not hold for Blinn-Phong because it

⁷The frequency space identity in this section (equation 18) cannot be derived for the known BRDF case, since the lighting is not radially symmetric and therefore cannot be eliminated.




Fig. 7. **Left:** Grace cathedral lighting environment. **Middle:** Sphere rendered with Blinn-Phong BRDF, $(N.H)^s$, $s = 500$. **Right:** Single image identity value for the rendered sphere. For low frequencies (shown here), the BRDF filter is essentially symmetric, and the identity values are small.

is not symmetric about the reflected direction. However, it can be shown that the BRDF filter is essentially symmetric for low frequencies l . Equation 18 holds for small frequencies but breaks down for high frequencies. So the identities in this and later sections are robust to small dissymmetries in the BRDF (e.g., low frequency symmetry), but with narrower operating range.

Our framework could be used to blindly (without watermarking) detect tampering of images, making sure a given photograph (containing a homogeneous object of known shape) is consistent with the illumination it is captured in.⁸ To the best of our knowledge, ours is the first theoretical framework to enable these kinds of consistency checks. Example applications of tamper detection on real objects are shown in figures 11 and 13.

Finally, it should be noted that if we are given the full light field (all views) instead of simply a single image, a similar identity to equation 18 holds for general BRDFs that need not be radially symmetric. In particular, based on equation 6, a similar derivation gives

$$B_{lipq}L_{lj} - B_{ljpq}L_{li} = 0. \quad (19)$$

For the rest of this paper, we will not explicitly write out the form of the identities for general light fields, but it should be understood that similar properties can be derived for general isotropic BRDFs and light fields for most of the formulae we discuss here.

⁸Our identities are “necessary” conditions for image consistency, under our assumptions and in the absence of noise. They are not theoretically “sufficient”. For example, if an unusual material were to zero out a certain frequency, tampering at that frequency might go undetected. Also note that noise tends to add high frequencies, while materials tend to filter out high frequencies, causing the consistency errors to rise (become less reliable) with harmonic order.

V. SINGLE IMAGE: COMBINING DIFFUSE AND SPECULAR

We now consider the more general case of an unknown glossy BRDF with both specular *and* Lambertian (diffuse) reflectance. To our knowledge, this is the first such combined diffuse plus specular theory of the single image case, and the analysis (such as equations 20 and 23) is likely to have broader implications for other problems in vision, such as photometric stereo and lighting-insensitive recognition.

A. Common Parameterization

The major technical difficulty is that while both diffuse (Lambertian) and specular components are radially symmetric, they are so in different parameterizations (normal vs reflected direction). An important technical contribution of this paper is to express the diffuse irradiance in the reflected parameterization,

$$B_{lm} = K_d D_{lm} + A_l^{\text{spec}} L_{lm}. \quad (20)$$

The parameters of reflectance are the diffuse coefficient K_d and the specular BRDF filter coefficients A_l (we drop the superscript from now on). D_{lm} are the spherical harmonic coefficients of the irradiance written in the reflected parameterization. They depend linearly on the lighting coefficients L_{lm} (assumed known) as $D_{lm} \approx \sum_{n=0}^2 A_n^{\text{Lamb}} L_{nm} T_{lmn}$, with $T_{lmn} = \int_{S^2} Y_{nm}(\frac{\alpha}{2}, \beta) Y_{lm}^*(\alpha, \beta) d\Omega$. The $\alpha/2$ in the first term converts from normal to reflected parameterization.⁹

The coefficients T_{lmn} can be determined analytically or numerically, since the formulae for Y_{nm} and Y_{lm}^* are well known. Plots for D_{lm} and T_{lmn} are shown in figure 8 for a particular complex natural lighting environment. Since n ranges from 0 to 2 for Lambertian reflectance, m varies from -2 to $+2$, so we can safely neglect terms with $|m| > 2$ or $|n| > 2$. Moreover, for $l \geq 2$, we find that T_{lmn} either vanishes or falls off rapidly as $l^{-3/2}$ or $l^{-5/2}$. Hence, though somewhat more complex, Lambertian effects in the reflected parameterization are still relatively simple and low frequency. Please see Appendix B for a more detailed derivation.

⁹We would like to emphasize that the reflected parameterization is not directly related to Rusinkiewicz's half-angle parameterization [Rus98]. In fact, the convolution theorem does not hold for the half-angle parameterization.




Fig. 8. (a),(b): D_{lm} plots for low and high frequencies. Note that D_{lm} coefficients are small for $|m| > 2$ and hence can be safely neglected. (c): T_{lmn} plot. T_{lmn} falls off rapidly as $l^{-3/2}$ or $l^{-5/2}$ for $l \geq 2$.

B. Determining K_d and Image Consistency:

We now seek to eliminate A_l from equation 20 to directly estimate K_d for inverse rendering and reflectance estimation.

As before, A_l can be eliminated by considering different values of m ,

$$\frac{B_{li} - K_d D_{li}}{L_{li}} = \frac{B_{lj} - K_d D_{lj}}{L_{lj}} \implies K_d = \frac{B_{li} L_{lj} - B_{lj} L_{li}}{D_{li} L_{lj} - D_{lj} L_{li}}. \quad (21)$$

Since the above equation is true for all l, i, j , we also get an identity that must hold for any l , i and j , and can be used for image consistency checking,

$$\frac{B_{l_1 i} L_{l_1 j} - B_{l_1 j} L_{l_1 i}}{D_{l_1 i} L_{l_1 j} - D_{l_1 j} L_{l_1 i}} = \frac{B_{l_2 m} L_{l_2 n} - B_{l_2 n} L_{l_2 m}}{D_{l_2 m} L_{l_2 n} - D_{l_2 n} L_{l_2 m}}. \quad (22)$$

C. Determining A_l and Image Consistency:

Equivalently, we can eliminate K_d ,

$$\frac{B_{li} - A_l L_{li}}{D_{li}} = \frac{B_{lj} - A_l L_{lj}}{D_{lj}} \implies A_l = \frac{B_{li} D_{lj} - B_{lj} D_{li}}{L_{li} D_{lj} - L_{lj} D_{li}}. \quad (23)$$

This can be used to directly estimate the specular BRDF coefficients, irrespective of the diffuse coefficient K_d . As a sanity check, consider the case when $K_d = 0$. In this case, $B_{li} = A_l L_{li}$, so the expression above clearly reduces to A_l . Hence, equation 23 can be considered a new robust form of reflectance estimation that works for both purely specular and general glossy materials. Further note that we estimate an accurate *non-parametric* BRDF representation specified by general filter coefficients A_l .




Fig. 9. **Left:** Synthetic sphere image with both diffuse (K_d set to 1) and specular (taken from measurements of a real material) components. **Middle:** Image consistency checks (equation 24) can distinguish small inconsistencies between illumination and image (“ECCV” vs “ICCV” lighting). **Right:** For estimation of A_l , our approach gives accurate results, outperforming a parametric estimation technique.

Since the formula above is true for all i, j , we get an identity for image consistency,

$$\frac{B_{li}D_{lj} - B_{lj}D_{li}}{L_{li}D_{lj} - L_{lj}D_{li}} = \frac{B_{lm}D_{ln} - B_{ln}D_{lm}}{L_{lm}D_{ln} - L_{ln}D_{lm}}. \quad (24)$$

Figure 9 shows these ideas applied to a synthetic sphere with both diffuse and specular components. In this case, we used as input A_l from measurements of a real material, and they *do not* correspond exactly to a Phong BRDF. Hence, our technique recovers the specular BRDF somewhat more accurately than a comparison method that simply does nonlinear estimation of Phong parameters. We also show image consistency checks similar to those in the previous section, using equation 24. As in the previous section, we can distinguish small inconsistencies between lighting and image. An application to detect splicing for a real object is shown in the left graph of figure 13.

VI. THEORETICAL ANALYSIS: TWO MATERIALS AND/OR LIGHTING CONDITIONS

Section IV analyzed the single object, single image case. In this section¹⁰, we first consider two different objects (with different materials) in the same lighting. Next, we consider one object imaged in two different lighting conditions. Then, we consider the two lighting/two BRDF case corresponding to two images (in different lighting conditions), each of two objects with distinct BRDFs. In the next section, we will discuss some broader implications.

¹⁰This section will primarily discuss the purely specular case. For consistency checking, we have seen that in the reflective reparameterization, the diffuse component mainly affects frequencies D_{lm} with $|m| \leq 2$. Therefore, it is simple to check the identities for $|m| > 2$. Diffuse relighting is actually done in the spatial domain, as discussed in section VII. Section VIII provides experimental validation with objects containing both diffuse and specular components.

A. Two Objects/BRDFs: Same Lighting

We consider a single image (hence in the same lighting environment) of two objects, with different BRDFs. Let us denote by superscripts 1 or 2 the two objects,

$$B_{lm}^1 = A_l^1 L_{lm} \quad B_{lm}^2 = A_l^2 L_{lm}. \quad (25)$$

From these, it is possible to eliminate the lighting by dividing,

$$\frac{B_{lm}^2}{B_{lm}^1} = \frac{A_l^2}{A_l^1} = \gamma_l. \quad (26)$$

We refer to γ_l as the *BRDF transfer function*. Given the appearance of one object in complex lighting, multiplication of spherical harmonic coefficients by this function gives the appearance of an object with a different material. γ_l is *independent* of the lighting condition, and can be used in any (unknown) natural illumination. Also note that this function is *independent of m* , so we can average over all m , which makes it very robust to noise—in our experiments, we have not needed any explicit regularization for the frequencies of interest. Moreover, we do not need to know or estimate the individual BRDFs. It is not clear that one can derive such a simple formula, or bypass explicit lighting/reflectance estimation, in the spatial/angular domain. Section VI-C will explore applications to rendering.

It is also possible to use these results to derive a frequency space identity that *depends only on the final images, and does not require explicit knowledge of either the lighting condition or the BRDFs*. We know that equation 26 should hold for all m , so

$$\frac{B_{li}^2}{B_{li}^1} = \frac{B_{lj}^2}{B_{lj}^1} \implies B_{li}^2 B_{lj}^1 - B_{li}^1 B_{lj}^2 = 0. \quad (27)$$

This identity can be used for consistency checking, making sure that two objects in an image are shaded in consistent lighting. This enables detection of inconsistencies, where one object is spliced into an image from another image with inaccurate lighting. Also note that the single image identity (equation 18) is just a special case of equation 27, where one of the objects is simply a mirror sphere (so, for instance, $B^1 = L$).

B. Two Lighting Environments: Same Object/BRDF

We now consider imaging the same object in two different lighting environments. Let us again denote by superscripts 1 or 2 the two images, so that,

$$B_{lm}^1 = A_l L_{lm}^1 \quad B_{lm}^2 = A_l L_{lm}^2. \quad (28)$$

Again, it is possible to eliminate the BRDF by dividing,

$$\frac{B_{lm}^2}{B_{lm}^1} = \frac{L_{lm}^2}{L_{lm}^1} = L'_{lm}. \quad (29)$$

We refer to L'_{lm} as the *lighting transfer function*. Given the appearance of an object in lighting condition 1, multiplication of spherical harmonic coefficients by this function gives the appearance in lighting condition 2. L'_{lm} is *independent* of the reflectance or BRDF of the object. Hence, the lighting transfer function obtained from one object can be applied to a different object observed in lighting condition 1. Moreover, we never need to explicitly compute the material properties of any of the objects, nor recover the individual lighting conditions.

The relighting application does not require explicit knowledge of either lighting condition. However, if we assume the lighting conditions are known (unlike the previous subsection, we need the lighting known here since we cannot exploit radial symmetry to eliminate it), equation 29 can be expanded in the form of an identity,

$$B_{lm}^2 L_{lm}^1 - B_{lm}^1 L_{lm}^2 = 0. \quad (30)$$

This identity can be used for consistency checking, making sure that two photographs of an object in different lighting conditions are consistent, and neither has been tampered.

C. Two Materials and Two Lighting Conditions

Finally, we consider the most conceptually complex case, where both the lighting and materials vary. This effectively corresponds to two images (in different lighting conditions), each containing two objects of different materials. We will now use two superscripts, the first for the lighting and the second for the material.

	Lighting 1	Lighting 2
BRDF 1	$B_{lm}^{1,1} = A_l^1 L_{lm}^1$	$B_{lm}^{2,1} = A_l^1 L_{lm}^2$
BRDF 2	$B_{lm}^{1,2} = A_l^2 L_{lm}^1$	$B_{lm}^{2,2} = A_l^2 L_{lm}^2$

Simply by multiplying out and substituting the relations above, we can verify the basic identity discussed in the introduction to this paper,

$$B_{lm}^{1,1} B_{lm}^{2,2} = B_{lm}^{1,2} B_{lm}^{2,1} = A_l^1 A_l^2 L_{lm}^1 L_{lm}^2, \quad (31)$$

or for the purposes of consistency checking,

$$\boxed{B_{lm}^{1,1} B_{lm}^{2,2} - B_{lm}^{1,2} B_{lm}^{2,1} = 0.} \quad (32)$$

An interesting feature of this identity is that we have completely eliminated all lighting and BRDF information. Consistency can be checked based simply on the final images, without estimating any illuminations or reflectances. Note that if the second object is a mirror sphere, this case reduces to the two lightings, same BRDF case in equation 30.

Equation 31 also leads to a simple framework for estimation. The conceptual setup is that we can estimate the appearance of the fourth lighting/BRDF image (without loss of generality, say this is $B_{lm}^{2,2}$), given the other three, *without explicitly computing any illumination or reflectances*. Clearly, this is useful to insert the second object into a photograph where it wasn't present originally, assuming we've seen both objects together under another lighting condition. From equation 31, we have

$$B_{lm}^{2,2} = \frac{B_{lm}^{1,2} B_{lm}^{2,1}}{B_{lm}^{1,1}} \quad (33)$$

$$= B_{lm}^{1,2} \left(\frac{B_{lm}^{2,1}}{B_{lm}^{1,1}} \right) = B_{lm}^{1,2} L'_{lm} \quad (34)$$

$$= B_{lm}^{2,1} \left(\frac{B_{lm}^{1,2}}{B_{lm}^{1,1}} \right) = B_{lm}^{2,1} \gamma_l. \quad (35)$$

This makes it clear that we can visualize the process of creating $B_{lm}^{2,2}$ in two different ways. Figure 10 further illustrates the two approaches. One way (a) is to start with *another* object in the *same* lighting condition, i.e. $B_{lm}^{2,1}$ and apply the BRDF transfer function γ_l . The BRDF transfer function is found from the image of both objects in lighting condition 2. Alternatively (b), we start with the *same* object in another lighting condition $B_{lm}^{1,2}$ and apply the lighting transfer function L'_{lm} obtained from *another* object. In practice, we prefer using the BRDF transfer function (equation 35), since γ_l is more robust to noise. This is because, γ_l is independent of m . Hence for a given l , we can average over different values of m , thus reducing the noise in the coefficients. In contrast, the lighting transfer functions are more sensitive to noise. Certain frequency modes in the source image might be suppressed, leading to division by zero in the lighting transfer function. Hence, the image (e), obtained using lighting transfer function L'_{lm}




Fig. 10. **Top Row:** Shows two different approaches for image relighting. We can either use BRDF Transfer function (a) or Lighting Transfer function (b). All the spheres are synthetic, with Lighting 1 being St. Peters environment map and Lighting 2 Grace Cathedral. We have used Phong BRDF, with Phong exponent $s = 500$ for object 1 and $s = 100$ for object 2. **Bottom Row:** Comparison of spheres generated using two approaches with actual sphere. Sphere (e) generated using Lighting Transfer has artifacts (note the ringing) whereas the sphere (c) generated using BRDF Transfer matches closely with the actual sphere (d).

has artifacts, whereas the one (c), obtained by using BRDF transfer function γ_l is consistent with actual image (d) due to robustness of γ_l to noise.

The idea of estimating the fourth light/BRDF image, given the other three, has some conceptual similarity to learning image analogies [HJO⁺01]. However, we are considering a convolution of lighting and BRDF, while image analogies try to synthesize images by rearranging input pixels, irrespective of the physics, and cannot achieve the desired result in general. Since *none* of the lightings or BRDFs are known, it would also be very difficult to render $B_{lm}^{2,2}$ with alternative physics-based inverse rendering methods.

VII. IMPLICATIONS AND DISCUSSION

We now briefly discuss some of the broader implications of our theory. First, we extend the two BRDF/two lighting case to multiple lighting conditions and BRDFs. Then, we discuss spatial domain setups and identities analogous to our frequency domain analysis. Finally, we

show how many previous spatial domain algorithms and invariants can be considered special cases, extensions or variants of this general class of identities.

A. Multiple lighting conditions and BRDFs

Let us consider p lighting conditions and q BRDFs, instead of assuming $p = q = 2$, with superscripts $i \leq p$ and $j \leq q$, so that

$$B_{lm}^{i,j} = A_l^j L_{lm}^i \implies \mathbf{B}_{lm} = \mathbf{L}_{lm} \mathbf{A}_l^T, \quad (36)$$

where in the last part, for a given spherical harmonic index (l, m) , we regard \mathbf{B}_{lm} as an $p \times q$ matrix obtained by multiplying column vectors \mathbf{L}_{lm} ($p \times 1$), corresponding to the lighting conditions, and the transpose of \mathbf{A}_l ($q \times 1$), corresponding to the BRDFs.

Equation 36 makes it clear that there is a *rank 1 constraint* on the $p \times q$ matrix \mathbf{B}_{lm} . Section VI-C has considered the special case $p = q = 2$, corresponding to a 2×2 matrix, where the rank 1 constraint leads to a single basic identity (equation 32). In fact, equation 32 simply states that the determinant of the singular 2×2 matrix \mathbf{B}_{lm} is zero.

B. Spatial Domain Analog

Equation 36 expresses the image of a homogeneous glossy material in the *frequency domain* as a *product* of lighting and BRDF. Analogously, a difficult to analyze frequency domain convolution corresponds to a simple spatial domain product. For example, the image of a textured Lambertian surface in the *spatial domain* is a *product* of albedo ρ_k and irradiance E_k , where k denotes the pixel.

$$B_k^{i,j} = \rho_k^j E_k^i \implies \mathbf{B}_k = \mathbf{E}_k \rho_k^T. \quad (37)$$

Equation 37 has the same product form as the basic convolution equation ($B_{lm} = A_l L_{lm}$). Hence an identity similar to equation 32 holds in the angular domain for textured Lambertian objects.

$$B_{\text{diffuse}}^{1,1}(\theta, \phi) B_{\text{diffuse}}^{2,2}(\theta, \phi) = B_{\text{diffuse}}^{1,2}(\theta, \phi) B_{\text{diffuse}}^{2,1}(\theta, \phi) \quad (38)$$

The BRDF transfer function $\gamma(\theta, \phi)$ is just the ratio of diffuse albedos and is constant for homogeneous objects.

These identities enable spatial domain techniques for re-rendering the diffuse component (which in our case has constant albedo since the material is homogeneous), while still using

the frequency domain for the specular component. In order to separate the diffuse and specular components from the images, we observe that in a parameterization by surface normals, B_{lm} will have essentially all of its diffuse energy for $l \leq 2$, while the specular energy falls away much more slowly [RH02], and therefore mostly resides in $l > 2$. So we assume that

$$B_{\text{diffuse}}(\theta, \phi) \approx \sum_{l=0}^2 \sum_{m=-l}^l B_{lm} Y_{lm}(\theta, \phi) \quad (39)$$

But a single image gives information only for a hemisphere of surface normals, so we cannot directly calculate B_{lm} for the normal parameterization. Spherical harmonics do not form a linearly independent basis for the hemisphere. We pose the diffuse computation as a fitting problem where we want to find $B_{lm}, l \leq 2$ that best fits the hemisphere. We solve a system of equations $AX = B$ corresponding to equation 39, where A is an $N \times 9$ matrix of Y_{lm} computed at N sample points on the hemisphere, X is a 9×1 matrix of the corresponding 9 B_{lm} coefficients and B is an $N \times 1$ matrix of irradiance at sample points. The specular component can then be handled as discussed in the previous section and the diffuse component can be computed using equation 38. The diffuse computation is more stable in the angular domain than in the spherical harmonics domain. This method is used in all our rendering examples. As expected, our practical results work less well for the extremes when the specular intensity is very small relative to the diffuse component (in the limit, a purely Lambertian surface) or vice versa (a purely specular object).

C. Analogies with Previous Spatial Domain Results

While the exact form of, and rank 1 constraint on, equation 37 is not common in previous work, many earlier spatial domain invariants and algorithms can be seen as using special cases and extensions thereof. We briefly discuss some prominent results in our framework, also describing our analogous frequency domain results. In this way, we provide a unified view of many spatial and frequency domain identities, that we believe confers significant insight.

Reflectance ratios [NB96] are widely used for recognition. The main observation is that at adjacent pixels, the irradiance is essentially the same, so that the ratio of image intensities corresponds to the ratio of albedos. Using superscripts for the different pixels as usual (we do not need multiple super- or any subscripts in this case), we have $B^2/B^1 = \rho^2/\rho^1$. The analogous frequency domain result is equation 26, corresponding to the two BRDFs, same lighting case.

In both cases, by dividing the image intensities or spherical harmonic coefficients, we obtain a result *independent of the illumination*.

Similarly, a simple version of the recent BRDF-invariant stereo work of Davis et al. [DYW05] can be seen as the two lighting, same BRDF case. For fixed view and point source lighting, a variant of equation 37 still holds, where we interpret ρ_k^j as the (spatially varying) BRDF for pixel k and fixed view, and E_k^i as the (spatially varying) light intensity at pixel k . If the light intensity changes (for the same pixel/BRDF), we have $B^2/B^1 = E^2/E^1$. The frequency domain analog is equation 29. In both cases, we have *eliminated the BRDF* by dividing image intensities or spherical harmonic coefficients.

Narasimhan et al. [NRN03] also assume point source lighting to derive photometric invariants in the spatial domain—note that our frequency domain framework, by contrast, easily handles general complex lighting. Narasimhan et al. [NRN03] consider a variant of equation 37 with a summation of multiple terms (such as diffuse plus specular). For each term, ρ encodes a material property such as the diffuse albedo, while E encodes the illumination intensity and geometric attributes (such as a cosine term for diffuse or a cosine lobe for specular). Their work can be seen as effectively deriving a rank constraint on \mathbf{B} , corresponding to the number of terms summed. For diffuse objects, this is a rank 1 constraint, analogous to that in the frequency domain for equation 36. For diffuse plus specular, this is a rank 2 constraint. They then effectively use the rank constraint to form appropriate determinants that eliminate either material or geometry/lighting attributes, as in our frequency domain work. Jin et al. [JSY03] employ a similar rank 2 constraint for multi-view stereo with both Lambertian and specular reflectance.

Finally, we note that while there are many analogies between previous spatial domain identities and those we derive in the spherical/angular frequency domain, *some of our frequency domain results have no simple spatial domain analog*. For example, the concept of angular radial symmetry does not transfer to the spatial domain, and there is no known spatial analog of the identities in equations 18, 19, 22, 24, and 27.

VIII. EXPERIMENTAL VALIDATION AND RESULTS

We now present some experiments to validate the theory, and show potential applications. We start with diffuse plus specular spheres in figure 11, since they correspond most closely with

our theory. We then describe results with a complex cat geometry (figures 1, 12 and 13). All of these results show that the theory can be applied in practice with real data, where objects are not perfectly homogeneous, there is noise in measurement and calibration, and specular reflectance is not perfectly radially symmetric.

Experimental Setup

We ordered spheres from <http://www.mcmaster.com>. The cat model was obtained at a local craft sale. All objects were painted to have various specular finishes and diffuse undercoats. While homogeneous overall, small geometric and photometric imperfections on the objects were visible at pixel scale and contributed “reflection noise” to the input images. To control lighting, we projected patterns onto two walls in the corner of a room. We placed a Canon EOS 10D camera in the corner and photographed the objects at a distance of 2-3m from the corner (see top left of figure 11). This setup has the advantage of more detailed frontal reflections, which are less compressed than those at grazing angles. However, frontal lighting also gives us little information at grazing angles, where the BRDF might violate the assumption of radial symmetry due to Fresnel effects; we hope to address this limitation in future experiments. To measure the lighting, we photographed a mirror sphere. To measure BRDFs (only for deconvolution), we imaged a sphere under a point source close to the camera, determining A_l by simply reading off the profile of the highlight, and K_d by fitting to the diffuse intensity. For all experiments, we assembled high-dynamic range images.

Glossy Spheres

Figure 11 shows the two lighting, two materials case. The top right shows a relighting application. We assume (b1) is unknown, and we want to synthesize it from the other 3 lighting/BRDF images (a1,a2,b2). We also do the same for rendering (b2) assuming we know (a1,a2,b1). The results are visually quite accurate, and in fact reduce much of the noise in the input. Quantitatively, the L_1 norm of the errors for (b1) and (b2) are 9.5% and 6.5% respectively. In the bottom row, we tamper (b2) by using image processing to squash the highlight slightly. With the naked eye, it is difficult to detect that image (c) is not consistent with lighting 2 or the other spheres. However, all three identities discussed in the previous section correctly detect the tampering.




Fig. 11. **Top Left:** Experimental setup. **Top Middle:** Two lightings (shown only for reference) and images of two glossy (diffuse plus specular) spheres in that lighting. **Top Right:** We can accurately render (b1), given (a1,a2,b2), and render (b2), given (a1,a2,b1). **Bottom:** We tamper (b2) to generate (c) by squashing the specular highlights slightly in photoshop. While plausible to the naked eye, all three identities in section VI clearly indicate the tampering (red graphs).

Complex Geometry

For complex (mostly convex) known geometry, we can map object points to points on the sphere with the same surface normal, and then operate on the resulting spherical image. Deconvolution is shown in figure 12. We used a sphere painted with the same material as the cat to acquire both the cat geometry, using example-based photometric stereo [HS05] for the normals, and the BRDF (needed only for deconvolution). Errors (unrelated to our algorithm) in the estimated geometry lead to some noise in the mapping to the sphere. Our deconvolution method for lighting estimation substantially sharpens the reflections, while removing much of the input noise. Moreover, our results are consistent with taking the actual lighting and convolving it with the product of the BRDF and Wiener spherical harmonic filters.

The cat can also be used directly as an object for relighting/rendering and consistency checking. An example of rendering is shown in figure 1. The L_1 norm of the error is somewhat higher than




Fig. 12. Deconvolution on a real cat image. **Left:** Geometry estimation, using example-based photometric stereo (we take a number of images with the cat and example sphere; the sphere is also used to find the BRDF). **Middle:** Input image under unknown lighting, and mapping to a sphere using the surface normals. **Right:** Closeups, showing the original sphere map, and our deconvolved lighting estimate on top. This considerably sharpens the original, while removing noise, and resembles the BRDF*Wiener filter applied to the actual lighting (bottom row).




Fig. 13. Image consistency checking for cat (labels are consistent with figure 1). The tampered image (c) is obtained by splicing the top half (b1) under lighting 1 and the bottom half (b2) under lighting 2. Image (c) looks quite plausible, but the splicing is clearly detected by our identities.

in figure 11, at 12%, primarily because this is a much more challenging example. We are using the BRDF transfer function from a much lower-frequency material to a higher-frequency one—the blue sphere has a much broader specular lobe than the green cat. Moreover, inaccuracies in the normal estimation (not part of our algorithm) lead to some visible contouring in the results. Nevertheless, we see that the results are visually plausible. Note that, even though our theory requires the full range of normals in the image in order to calculate the spherical harmonics transform, in practice it works well even when the estimated normals are noisy or some of the normals are missing.

Figure 13 illustrates photomontage image tampering, in which the top half under lighting 1

(b1 in figure 1) is spliced with the bottom half under lighting 2 (b2 in figure 1). While the image (c) looks plausible in itself, the identities for both single and multiple images clearly detect the tampering.

For image consistency checking, our identities require explicit knowledge of 3D geometry, as well as a homogeneous object. However, increasingly good methods exist to acquire this geometry from a single image, both without and with minimal user assistance [ZPSS02], [OCDD01]. Moreover, diffuse and specular components can be separated automatically with modern techniques to handle textured objects [MZB06]. Finally, many of the applications of image consistency checking are focused on changes and inconsistencies to the appearance of known objects like human faces, where it is easy to find generic 3D models—or with the increasing popularity of 3D scanners, even an actual 3D model and appearance model of the subject. Famous recent examples of digital forgeries or touchups are the darkening of the OJ Simpson photograph on the Time magazine cover¹¹, and a recent forged image of John Kerry and Jane Fonda appearing together¹²(their faces were composited with inconsistent lighting). Our theoretical framework provides a solid foundation for applying practical image consistency checks to determine consistency of lighting and shading in these scenarios.

IX. CONCLUSIONS AND FUTURE WORK

In this paper, we have introduced a new theoretical framework for using spherical convolution and deconvolution in inverse rendering, BRDF/lighting transfer and image consistency checking. The main contribution is the set of new frequency domain invariants, which represent fundamental identities following from the convolution theorem. These identities often eliminate the lighting and/or BRDF, enabling a new class of inverse rendering algorithms that can relight or change materials by using BRDF/lighting transfer functions, without explicit illumination or BRDF estimation. In the future, similar ideas may be applied to other problems, such as BRDF-invariant stereo and photometric stereo, or lighting-insensitive recognition. The theoretical framework also makes a contribution to the relatively new area of image consistency checking, describing a suite of frequency domain identities to detect tampering and other undesirable image processing

¹¹http://www.authentichistory.com/diversity/african/images/1994_OJ_Simpson_Time_Magazine.html

¹²<http://www.snopes.com/photos/politics/kerry2.asp>

operations. We have also presented a new unified view of spatial and frequency domain identities and rank constraints, that can give insight for developing future algorithms in either, or even a combination of both domains.

In the future, from a theoretical perspective, we want to develop a framework for operating on local subsets of the entire image, corresponding to small portions of the full sphere of directions. From a practical perspective, we want to better understand the sensitivity of our identities—initial tests indicate they are fairly robust, but more work needs to be done. We wish to apply our algorithms in more complex cases like faces where the geometry is not known accurately, and where objects may not be perfectly convex. We would also like to handle textured objects by automatic diffuse/specular separation methods [MZB06]. We believe that the theory may also lead to the construction of better light probes where we can replace the mirror sphere by a sphere of general material and hence bypass the serious issues like dynamic range associated with current light probes.

In summary, we see this paper as introducing the basic theory, that can lead to much future theoretical and practical work in inverse rendering and image consistency checking.

ACKNOWLEDGEMENTS

We thank Sameer Agarwal for many helpful discussions. This work was supported in part by an ONR Young Investigator Award (Mathematical Models of Illumination and Reflectance for Image Understanding and Machine Vision), NSF grant # 0430258 (CyberTrust - Restore the Trustworthiness of Digital Photographs: Blind Detection of Digital Photograph Tampering), as well as NSF grants # 0098005, # 0305322 and # 0446916, the Washington Research Foundation, Microsoft Research, and the University of Washington Animation Research Labs.

REFERENCES

- [BJ01] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. In *CVPR 01*, pages II-374–II-381, 2001.
- [BJ03] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, 2003.
- [BV99] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH 1999*, pages 187–194, 1999.
- [DYW05] J. Davis, R. Yang, and L. Wang. BRDF invariant stereo using light transport constancy. In *ICCV 05*, pages 436–443, 2005.
- [GW03] R. Gonzalez and R. Woods. *Digital Image Processing, Second Edition*. Pearson Education, 2003.
- [HJO⁺01] A. Hertzmann, C. Jacobs, N. Oliver, B. Curless, and D. Salesin. Image analogies. In *SIGGRAPH 01*, pages 327–340, 2001.
- [HS05] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *PAMI*, 27(8):1254–1264, 2005.

- [JF05] M. Johnson and H. Farid. Exposing digital forgeries by detecting inconsistencies in lighting. In *ACM Multimedia and Security Workshop*, pages 1–10, 2005.
- [JSY03] H. Jin, S. Soatto, and A. Yezzi. Multi-view stereo beyond lambert. In *CVPR 03*, pages 171–178, 2003.
- [LWTS05] Z. Lin, R. Wang, X. Tang, and H. Shum. Detecting doctored images using camera response normality and consistency. In *CVPR 05*, pages 1087–1092, 2005.
- [MG97] S. Marschner and D. Greenberg. Inverse lighting for photography. In *Color Imaging Conf.*, pages 262–265, 1997.
- [MRC06] D. Mahajan, R. Ramamoorthi, and B. Curless. A theory of spherical harmonic identities for brdf/lighting transfer and image consistency. In *ECCV 06*, pages IV: 41–55, 2006.
- [MWLT00] S. Marschner, S. Westin, E. Lafortune, and K. Torrance. Image-Based BRDF measurement. *Applied Optics*, 39(16):2592–2600, 2000.
- [MZB06] S. Mallick, T. Zickler, and P. Belhumeur. Specularity removal in images and videos: A PDE approach. In *ECCV 06*, pages I: 550–563, 2006.
- [NB96] S. Nayar and R. Bolle. Reflectance based object recognition. *IJCV*, 17(3):219–240, 1996.
- [NCS04] T. Ng, S. Chang, and Q. Sun. Blind detection of photomontage using higher order statistics. In *IEEE International Symposium on Circuits and Systems*, pages V–688–V–691, 2004.
- [NRN03] S. Narasimhan, V. Ramesh, and S. Nayar. A class of photometric invariants: Separating material from shape and illumination. In *ICCV 03*, pages 1387–1394, 2003.
- [OCDD01] B. Oh, M. Chen, J. Dorsey, and F. Durand. Image-based modeling and photo editing. In *SIGGRAPH 01*, pages 433–442, 2001.
- [RH01] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, pages 117–128, 2001.
- [RH02] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. *ACM Transactions on Graphics (SIGGRAPH 2002)*, 21(3):517–526, 2002.
- [Rus98] S. Rusinkiewicz. A new change of variables for efficient BRDF representation. In *Eurographics Rendering Workshop*, 1998.
- [SFB03] D. Simakov, D. Frolova, and R. Basri. Dense shape reconstruction of a moving object under arbitrary, unknown lighting. In *ICCV 03*, pages 1202–1209, 2003.
- [SKS02] P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics (SIGGRAPH 2002)*, 21(3):527–536, 2002.
- [SSI99] I. Sato, Y. Sato, and K. Ikeuchi. Illumination distribution from brightness in shadows: adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. In *ICCV 99*, pages 875 – 882, 1999.
- [Wie42] N. Wiener. Extrapolation, interpolation and smoothing of stationary time series. In *the MIT Press*, 1942.
- [WLH03] Z. Wen, Z. Liu, and Thomas S. Huang. Face relighting with radiance environment maps. In *CVPR 2003*, pages II: 158–165, 2003.
- [ZPSS02] Li Zhang, G. Phocion, J. Samson, and S. Seitz. Single view modeling of free-form scenes. *Journal of Visualization and Computer Animation*, 13(4):225–235, 2002.
- [ZWS05] L. Zhang, S. Weng, and D. Samaras. Face synthesis and recognition from a single image under arbitrary unknown lighting using a spherical harmonic basis morphable model. In *CVPR 2005*, pages 209–216, 2005.