

A Theory of Spherical Harmonic Identities for BRDF/Lighting Transfer and Image Consistency

Dhruv Mahajan¹, Ravi Ramamoorthi¹, and Brian Curless²

¹ Columbia University
{dhruv, ravir}@cs.columbia.edu
² University of Washington
curless@cs.washington.edu

Abstract. We develop new mathematical results based on the spherical harmonic convolution framework for reflection from a curved surface. We derive novel identities, which are the angular frequency domain analogs to common spatial domain invariants such as reflectance ratios. They 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. Denote the spherical harmonic coefficients by $B_{lm}^{light,material}$, where the subscripts refer to the spherical harmonic indices, and the superscripts to the lighting (1 or 2) and object or material (again 1 or 2). We derive a basic identity, $B_{lm}^{1,1} B_{lm}^{2,2} = B_{lm}^{1,2} B_{lm}^{2,1}$, independent of the specific lighting configurations or BRDFs. 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.

1 Introduction

Recent work by Basri and Jacobs [2], and Ramamoorthi and Hanrahan [15] has shown that the appearance of a curved surface can be described as a spherical convolution of the (distant) illumination and BRDF. This result often enables computer vision algorithms, previously restricted to point sources without attached shadows, to work in general complex lighting. Many recent articles have explored theoretical and practical applications for Lambertian surfaces (e.g., [1, 16]). However, there has been relatively little work in vision on using the convolution formulae for general glossy objects.

The main goal of this paper is to derive new formulae and identities for direct frequency domain spherical (de)convolution. As with most previous work in the area, we assume curved homogeneous objects (single BRDF) of known shape lit by complex distant illumination, and neglect cast shadows and interreflections. As explained in section 2, we also assume the BRDF is radially symmetric, which is a good approximation for most specular reflectance. A first example of our framework (section 3.1) is illumination estimation 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.

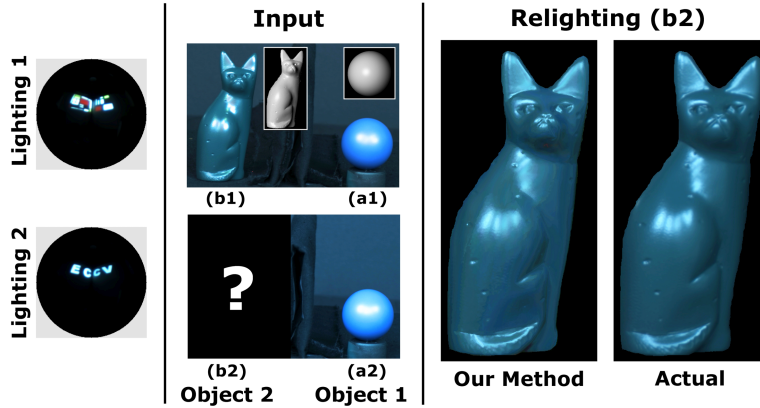


Fig. 1. One application of our framework. We are given real photographs of two objects of known geometry (shown in inset; note that both objects can be arbitrary, and one of them is a sphere here only for convenience). The two objects have different (and unknown) diffuse and specular material properties. 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.

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 [15], but cannot be robustly applied, since BRDF coefficients become small at high frequencies. Our first contribution is the adaptation of Wiener filtering [4, 11] from image processing to develop robust deconvolution filters (figures 2 and 5).

Deconvolution can be considered a particular case, involving a single image of an object with known reflectance in unknown lighting. In this paper, we study single (section 3) and multiple (section 4) images under single and multiple lighting conditions. *Our main contribution is the derivation of a number of novel frequency domain identities.* For example, one important case we consider (section 4.3) is that of two different glossy¹ materials in two different lighting environments. 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 lightings or BRDFs.

We show (section 4.4) that this and a related class of identities can be considered the frequency domain analog of fundamental spatial domain invariants, such as reflectance ratios (Nayar and Bolle [13]) or photometric invariants (Narasimhan et al. [12]).

¹ Parts of the theory (in sections 3.2 and 4) 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.

This paper is motivated by two important practical applications. The first is inverse rendering [10, 15]. Besides estimation of lighting and BRDFs, we also develop more general algorithms, which directly relight and change material properties by transferring the BRDF or lighting from another object or illumination. For example, our identity above enables us to render the fourth light/BRDF image (say $B_{lm}^{2,2}$), 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.

The second, newer application, is to verify image consistency and detect tampering (Johnson and Farid [7], Lin et al. [8]). The widespread availability of image processing tools enables creation of “forgeries” such as splicing images together (one example is shown in figure 6). Most previous work has focused on checking consistency at a signal or pixel level, such as the camera response [8], or wavelet coefficients (Ng et al. [14]). However, (in)consistencies of lighting and shading also provide valuable clues. This paper takes an important first step in laying theoretical foundations for this new area.

2 Background

We now briefly introduce the spherical convolution and signal-processing framework [2, 15] 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 this 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:

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

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 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, good approximations [15] 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 frequency domain, the reflected light B is given by a simple product formula or spherical convolution (see [2, 15] for the derivation and an analysis of convolution),

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

where for convenience, we define A_l and the normalization constant A_l as

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

The remainder of this paper derives new identities and formulae from equation 4, $B_{lm} = A_l L_{lm}$. Practical spherical harmonic computations are possible from only a *single image*, since a single view of a sufficiently curved object (assuming a distant viewer) sees all reflected directions. Most glossy BRDFs (such as Torrance-Sparrow) are approximately radially symmetric, especially for non-grazing angles of reflection [15]. Most of the theory in this paper also carries over to general isotropic materials, if we consider the entire light field, corresponding to multiple views [9].

3 Theoretical Analysis: Single Image of One Object

We develop the theory for the simplest case of a single image of one object in this section, with multiple objects and lighting conditions discussed later in section 4.

3.1 Known BRDF: Deconvolution to Estimate Lighting

We start by considering a known BRDF, where we want to determine the lighting. 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, we use the dual lighting estimation algorithm [15], which finds the specular B_{lm} consistent with the diffuse component. As per equation 4, B_{lm} will be a blurred version of the lighting, filtered by the glossy BRDF.

From equation 4 in the spherical harmonic domain, we derive

$$L_{lm} = \frac{B_{lm}}{A_l} = A_l^{-1} B_{lm}, \quad (6)$$

where the last identity makes explicit that we are convolving with a new radially symmetric kernel A_l^{-1} , which can be called the inverse, sharpening or deconvolution filter.

Unfortunately, it is difficult to apply equation 6 directly, since A_l in the denominator will become small for high frequencies, or alternatively the inverse filter A_l^{-1} will become very large. This will lead to substantial amplification of noise at high frequencies.

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 [4, 11] for this purpose. Specifically, we define a new inverse filter,

$$A_l^* = \frac{1}{A_l} \left(\frac{|A_l|^2}{|A_l|^2 + K} \right) = \frac{A_l}{|A_l|^2 + K} \quad L_{lm} = A_l^* B_{lm}, \quad (7)$$




Fig. 2. Deconvolution filter in the angular (a), and frequency (b) domains. The combined effect (multiplication in the frequency domain) of Phong and deconvolution filters is in (c).

where K is a small user-controlled constant. 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 2 shows a deconvolution filter A_l^* in the angular (a) and frequency (b) domains. The convolution of this filter with the original Phong filter (blue graph in c) lets through most frequencies without attenuation, while filtering out the very high frequencies. Figure 5 shows an application of deconvolution with real images and objects.

3.2 Known Lighting: Identity Obtained by Eliminating the BRDF

We now consider the converse case, where the lighting is known, but the BRDF is not. We will *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

$$A_l = \frac{B_{lm}}{L_{lm}}. \quad (8)$$

This expression could be used to solve for BRDF coefficients.² 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

$$\frac{B_{li}}{L_{li}} = \frac{B_{lj}}{L_{lj}} \quad (9)$$

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

$$B_{li}L_{lj} - B_{lj}L_{li} = 0. \quad (10)$$

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 (spherical) domain. However, they are related by a 1D radially symmetric BRDF, leading to a 1D redundancy, that can be used for consistency checking in equation 10.

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

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

² Since natural lighting usually includes higher frequencies than the BRDF, we can apply equation 8 directly without regularization, and do not need to explicitly discuss deconvolution.

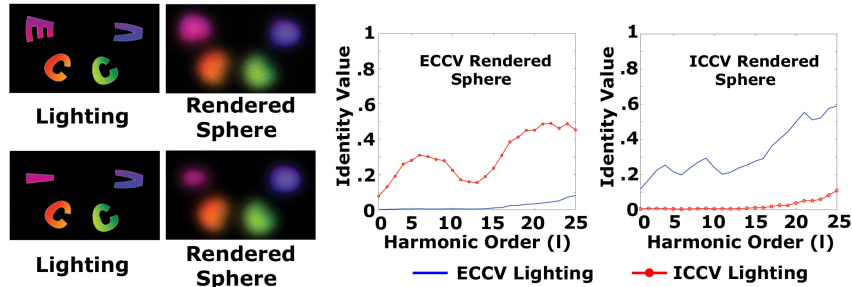


Fig. 3. Left: The synthetic images used. These correspond to closeups of specular spheres rendered with “ECCV” and “ICCV” lighting. To the naked eye, the two images look very similar. **Middle and Right:** The graphs show that our identity can clearly distinguish consistent image/lighting pairs (lower line) from those where lighting and image are inconsistent (upper line).

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 3 applies our theory to synthetic data of an ideal Phong BRDF, with noise added. We show closeups of (a rectangular region in) 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 10 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 [7], 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 (though we still require known geometry, and assume homogeneous materials.)

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.³ 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 4 and 6.

3.3 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.

³ Our 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 may 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.

Common Parameterization: While both diffuse and specular components are radially symmetric, they are so in different parameterizations (normal vs reflected direction). Hence, we first express the diffuse irradiance in the reflected parameterization,

$$B_{lm} = K_d D_{lm} + A_l^{\text{spec}} L_{lm}, \quad (11)$$

where D_{lm} are the spherical harmonic coefficients of the irradiance written in the reflected parameterization⁴. They depend only on the lighting, which is assumed known, and are small [9] for $|m| > 2$. 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).

Determining A_l and Image Consistency: For example, we can now eliminate K_d ,

$$\frac{B_{li} - A_l L_{li}}{D_{li}} = \frac{B_{lj} - A_l L_{lj}}{D_{lj}} \implies A_l = \frac{B_{li} D_{lj} - B_{lj} D_{li}}{L_{li} D_{lj} - L_{lj} D_{li}}. \quad (12)$$

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 12 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 .

Since the formula above is true for all i, j , we get an identity for image consistency. An application to detect splicing for a real object is shown in the left graph of figure 6.

$$\frac{B_{li} D_{lj} - B_{lj} D_{li}}{L_{li} D_{lj} - L_{lj} D_{li}} = \frac{B_{lm} D_{ln} - B_{ln} D_{lm}}{L_{lm} D_{ln} - L_{ln} D_{lm}}. \quad (13)$$

4 Theoretical Analysis: Two Materials and/or Lighting Conditions

Section 3 analyzed the single object, single image case. In this section⁵, 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. Finally, we discuss some broader implications.

4.1 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,

⁴ D_{lm} depend linearly on lighting coefficients L_{lm} 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. The coefficients T_{lmn} can be determined analytically or numerically [9].

⁵ This 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 4.4. Section 5 provides experimental validation with objects containing both diffuse and specular components.

$$B_{lm}^1 = A_l^1 L_{lm} \quad B_{lm}^2 = A_l^2 L_{lm}. \quad (14)$$

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

$$\frac{B_{lm}^2}{B_{lm}^1} = \frac{A_l^2}{A_l^1} = \gamma_l. \quad (15)$$

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, and can be used in any (unknown) natural illumination. Also note that γ_l is *independent of m* , so we can average over all m , which makes it robust to noise—in our experiments, we have not needed explicit regularization for the frequencies of interest. Moreover, we do not need to know or estimate the individual BRDFs.

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 15 should hold for all m , so

$$\frac{B_{li}^2}{B_{li}^1} = \frac{B_{lj}^2}{B_{lj}^1} \implies B_{li}^2 B_{lj}^1 - B_{li}^1 B_{lj}^2 = 0. \quad (16)$$

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 inconsistent lighting. As with equation 10, the images B are 2D functions, but related by a 1D radially symmetric BRDF, leading to redundancies in the spherical harmonic coefficients.

Finally, note that the single image identity (equation 10) is a special case of equation 16, where one of the objects is simply a mirror sphere (so, for instance, $B^1 = L$).

4.2 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,

$$B_{lm}^1 = A_l L_{lm}^1 \quad B_{lm}^2 = A_l L_{lm}^2. \quad (17)$$

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

$$\frac{B_{lm}^2}{B_{lm}^1} = \frac{L_{lm}^2}{L_{lm}^1} = L'_{lm}. \quad (18)$$

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. 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 18 can be expanded in the form of an identity,

$$B_{lm}^2 L_{lm}^1 - B_{lm}^1 L_{lm}^2 = 0. \quad (19)$$

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.

4.3 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,

$$B_{lm}^{1,1} B_{lm}^{2,2} = B_{lm}^{1,2} B_{lm}^{2,1} = A_l^1 A_l^2 L_{lm}^1 L_{lm}^2, \quad (20)$$

or for the purposes of consistency checking,

$$\boxed{B_{lm}^{1,1} B_{lm}^{2,2} - B_{lm}^{1,2} B_{lm}^{2,1} = 0.} \quad (21)$$

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 19.

Equation 20 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 20, we have

$$B_{lm}^{2,2} = \frac{B_{lm}^{1,2} B_{lm}^{2,1}}{B_{lm}^{1,1}} \quad (22)$$

$$= B_{lm}^{1,2} \left(\frac{B_{lm}^{2,1}}{B_{lm}^{1,1}} \right) = B_{lm}^{1,2} L'_{lm} \quad (23)$$

$$= B_{lm}^{2,1} \left(\frac{B_{lm}^{1,2}}{B_{lm}^{1,1}} \right) = B_{lm}^{2,1} \gamma_l. \quad (24)$$

This makes it clear that we can visualize the process of creating $B_{lm}^{2,2}$ in two different ways. One approach is to 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. Alternatively, we 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 1. In practice, we prefer using the BRDF transfer function (equation 24), since γ_l is more robust to noise. However, the equations above make clear that both interpretations are equivalent, following naturally from equation 20. Since *none* of the lightings or BRDFs are known, it would be very difficult to render $B_{lm}^{2,2}$ with alternative physics-based inverse rendering methods.

4.4 Implications and Discussion

We now briefly discuss some of the broader implications of our theory, and previous spatial domain identities and invariants analogous to our frequency domain results.

Multiple lighting conditions and BRDFs: Let us consider r lighting conditions and s BRDFs, instead of assuming $r = s = 2$, with superscripts $i \leq r$ and $j \leq s$, so that

$$B_{lm}^{i,j} = A_l^j L_{lm}^i \implies \mathbf{B}_{lm} = \mathbf{L}_{lm} \mathbf{A}_l^T, \quad (25)$$

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

Equation 25 makes it clear that there is a *rank 1 constraint* on the $r \times s$ matrix \mathbf{B}_{lm} . Section 4.3 has considered the special case $r = s = 2$, corresponding to a 2×2 matrix, where the rank 1 constraint leads to a single basic identity (equation 21). In fact, equation 21 simply states that the determinant of the singular 2×2 matrix \mathbf{B}_{lm} is zero.

Spatial Domain Analog: Equation 25 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.

$$B_k^{i,j} = \rho_k^j E_k^i \implies \mathbf{B}_k = \mathbf{E}_k \rho_k^T. \quad (26)$$

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.⁶ We use the *spatial domain analogs* of equations 20 and 24 to compute the diffuse component of $B^{2,2}$.

⁶ We first separate diffuse and specular components by observing 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, and therefore mostly resides in $l > 2$. Since we have only the front facing normals in a single image, we use fitting to remove the low-frequency terms with $l \leq 2$, and thereby separate diffuse and specular. As expected, this works 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).

In this case, the spatial analog of the BRDF transfer function γ in equation 24 simply corresponds to the ratio of the (in our case uniform) diffuse albedos of the two objects.

Analogies with Previous Spatial Domain Results: While the exact form of, and rank 1 constraint on, equation 26 is not common in previous work, many earlier spatial domain invariants and algorithms can be seen as using special cases and extensions.

Reflectance ratios [13] 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 15, corresponding to the two BRDFs, same lighting case. In both cases, by dividing the image intensities (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. [3] can be seen as the two lighting, same BRDF case. For fixed view and point source lighting, a variant of equation 26 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 18. In both cases, we have *eliminated the BRDF* by dividing image intensities or spherical harmonic coefficients.

Narasimhan et al. [12] also assume point source lighting to derive spatial domain photometric invariants. By contrast, our frequency domain framework handles general complex lighting (but is limited to homogeneous objects with known shape, and a global spherical harmonic analysis as opposed to local pixel operations.) Narasimhan et al. [12] consider a variant of equation 26 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 illumination and geometric attributes. Their work can be seen as effectively deriving a rank constraint on \mathbf{B} , corresponding to the number of terms. For diffuse objects, this is a rank 1 constraint, analogous to that in the frequency domain for equation 25. For diffuse plus specular, this is a rank 2 constraint. They then effectively use the rank constraint to form determinants that eliminate either material or geometry/lighting attributes, as in our frequency domain work. Jin et al. [6] use a similar rank 2 constraint for multi-view stereo with both Lambertian and specular reflectance.

Finally, we note that 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 equations 10, 13, and 16.

5 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 4, since they correspond most closely with our theory. We then describe results with a complex cat geometry (figures 1, 5 and 6). 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.

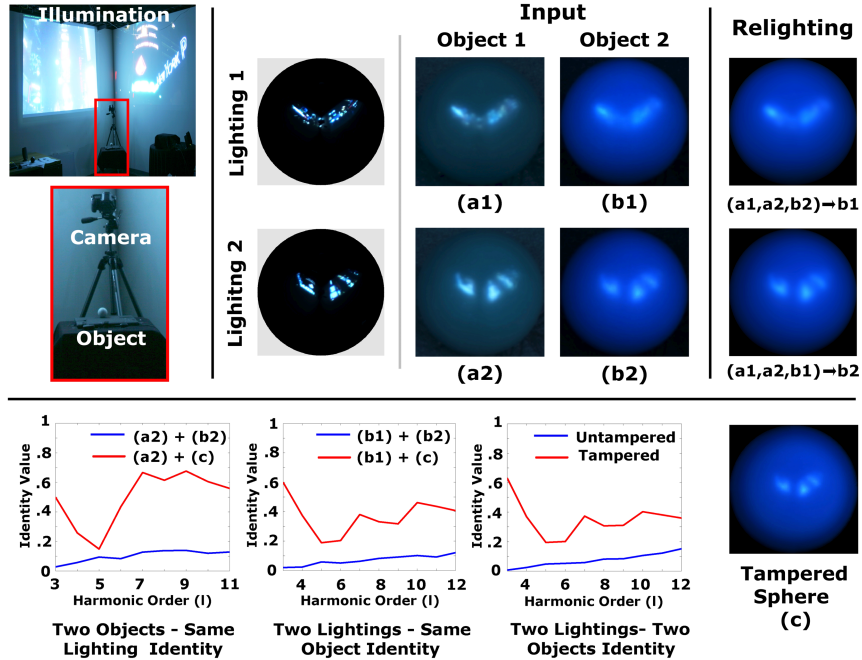


Fig. 4. Top Left: Experimental setup. **Top Middle:** Two lightings (shown only for reference) and images of two glossy (diffuse plus specular) spheres in that lighting. **Top Right:** We can accurately render (b1), given (a1,a2,b2), and render (b2), given (a1,a2,b1). **Bottom:** We tamper (b2) to generate (c) by squashing the specular highlights slightly in photoshop. While plausible to the naked eye, all three identities in section 4 clearly indicate the tampering (red graphs).

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 4). 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 4 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




Fig. 5. 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).

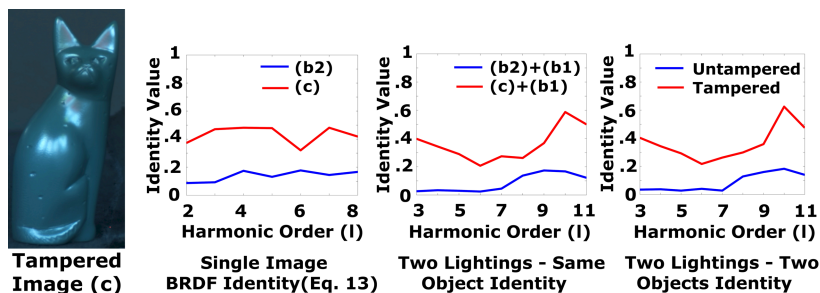


Fig. 6. Image consistency checking for cat (labels are consistent with figure 1). The tampered image (c) is obtained by splicing the top half (b1) under lighting 1 and the bottom half (b2) under lighting 2. Image (c) looks quite plausible, but the splicing is clearly detected by our identities.

(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.

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 5. We used a sphere painted with the same material as the cat to acquire both the cat geometry, using example-based photometric stereo [5] 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.

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. Note that this is a very challenging example, since 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. Nevertheless, we see that the results are visually plausible. Figure 6 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, the identities for both single and multiple images clearly detect tampering.

6 Conclusions and Future Work

In this paper, we have derived a set of novel frequency space identities. 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. In the future, similar ideas may be applied to other problems, such as BRDF-invariant stereo or lighting-insensitive recognition. We would also like to extend our identities and algorithms to work with textured objects, local patches, and unknown or approximate geometry. Beyond inverse rendering, our 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. Finally, we have presented a new unified view of spatial and frequency domain identities and rank constraints.

Acknowledgements. We thank Sameer Agarwal for many helpful discussions. This work was supported by 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

1. R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. In *CVPR 01*, pages II-374–II-381, 2001.
2. R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, 2003.
3. J. Davis, R. Yang, and L. Wang. BRDF invariant stereo using light transport constancy. In *ICCV 05*, 2005.
4. R. Gonzalez and R. Woods. *Digital Image Processing*. Pearson Education, 2003.
5. A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *PAMI*, 27(8):1254–1264, 2005.
6. H. Jin, S. Soatto, and A. Yezzi. Multi-view stereo beyond lambert. In *CVPR 03*, pages 171–178, 2003.
7. M. Johnson and H. Farid. Exposing digital forgeries by detecting inconsistencies in lighting. In *ACM Multimedia and Security Workshop*, 2005.
8. Z. Lin, R. Wang, X. Tang, and H. Shum. Detecting doctored images using camera response normality and consistency. In *CVPR 05*, 2005.
9. D. Mahajan. A theory of spherical harmonic identities for inverse rendering, BRDF/lighting transfer and image consistency checking. Master’s thesis, Columbia, 2006.
10. S. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell, 1998.

11. Wiener N. Extrapolation, interpolation and smoothing of stationary time series. In *the MIT Press*, 1942.
12. 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.
13. S. Nayar and R. Bolle. Reflectance based object recognition. *IJCV*, 17(3):219–240, 1996.
14. T. Ng, S. Chang, and Q. Sun. Blind detection of photomontage using higher order statistics. In *IEEE International Symposium on Circuits and Systems*, 2004.
15. R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, pages 117–128, 2001.
16. 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.