

On the Duality of Forward and Inverse Light Transport

Manmohan Chandraker, *Member, IEEE*,
 Jiamin Bai, *Student Member, IEEE*,
 Tian-Tsong Ng, *Member, IEEE*, and
 Ravi Ramamoorthi, *Member, IEEE*

Abstract—Inverse light transport seeks to undo global illumination effects, such as interreflections, that pervade images of most scenes. This paper presents the theoretical and computational foundations for inverse light transport as a dual of forward rendering. Mathematically, this duality is established through the existence of underlying Neumann series expansions. Physically, it can be shown that each term of our inverse series cancels an interreflection bounce, just as the forward series adds them. While the convergence properties of the forward series are well known, we show that the oscillatory convergence of the inverse series leads to more interesting conditions on material reflectance. Conceptually, the inverse problem requires the inversion of a large light transport matrix, which is impractical for realistic resolutions using standard techniques. A natural consequence of our theoretical framework is a suite of fast computational algorithms for light transport inversion—analogue to finite element radiosity, Monte Carlo and wavelet-based methods in forward rendering—that rely at most on matrix-vector multiplications. We demonstrate two practical applications, namely, separation of individual bounces of the light transport and fast projector radiometric compensation, to display images free of global illumination artifacts in real-world environments.

Index Terms—Please provide.

1 INTRODUCTION

GLOBAL illumination effects are key visual features of real-world scenes. Simulation of these effects in forward rendering has been extensively studied in computer graphics, with a theoretical foundation based on the rendering equation [1]. In contrast, most computer vision algorithms are forced to simply ignore interreflections where one would ideally like to undo their effects. Recently, Seitz et al. [2] formalized this as the problem of inverse light transport. However, little is known about the theory and algorithms for efficient light transport inversion in practical scenes.

This paper lays the mathematical and computational foundations of inverse light transport by exposing a strong duality to the mature framework of forward light transport. Intuitively, the duality arises because solving the (forward) rendering equation itself involves an operator or matrix inverse. Exploiting this duality allows us to leverage many theoretical results and algorithms from forward global illumination for the inverse problem in computer vision.

Specifically, forward rendering readily admits to a Neumann series solution. We derive a similar series for the inverse problem and show that just as each term of the forward series adds bounces of light, each term of the inverse series zeroes out the corresponding bounce (but, unlike in the forward case, also affects higher order bounces, leading to oscillatory convergence). The forward

series convergence corresponds to energy conservation, but the inverse condition is more complex—a sufficient condition is that the albedo of surfaces is below 0.5 so that the net global illumination is less than the direct lighting component.

Recent techniques for acquiring the light transport of real scenes [3], [4] have facilitated relighting applications in graphics, equivalent to matrix-vector multiplication. The light transport matrices in this paper are of full rank, corresponding to projector-camera setups (like [2], [5], but unlike [6]). While transport inversion enables new applications like illumination estimation, separating bounces of global illumination [2], projector radiometric compensation [7], and shape recovery [8], the high resolution of real transports ($10^5 \times 10^5$ or higher) makes standard matrix inversion impractical.

Inspired by efficient solutions such as finite element radiosity [9] and Monte Carlo methods [1], [10] for the forward problem, we propose fast algorithms for canceling interreflections, which require only matrix-vector multiplications (as opposed to a full matrix inversion). We demonstrate practical applications of these algorithms, such as radiometric compensation of interreflections while projecting complex scenes (Fig. 1), as well as separation of local and global illumination components or individual bounces (Fig. 2).

To summarize, this paper makes several important theoretical, algorithmic, and practical contributions.

- A theoretical framework that provides novel insights into light transport inversion by posing it as a dual to forward rendering.
- Efficient algorithms for high-resolution transport inversion, with rigorous convergence analysis.
- Demonstration of practical applications such as bounce separation and radiometric compensation in complex, non-Lambertian scenes.

2 PREVIOUS WORK

This paper builds upon and extends the theoretical framework proposed in [11].

Inverse rendering. Lighting and reflectance properties are acquired in [12], [13], and [14] assuming known scene geometry. Like [2], we observe only the light transport matrix—both geometry and reflectance are neither known, nor explicitly estimated. In contrast to all of the above, our algorithms can handle high resolutions for both lighting and geometry.

Forward rendering. This work draws upon operator formulations and error analysis for global illumination [15], Monte Carlo algorithms [10], and finite element radiosity methods [9]. Many iterative radiosity techniques are also closely related to numerical linear algebra methods [16], [17] for solving systems of linear equations, such as Jacobi and Gauss-Seidel iterations. Our framework enables similar relations to be drawn for inverse rendering. Similarly, our Monte Carlo method bears similarities to forward path tracing [1], as well as von Neumann and Ulam's original Monte Carlo matrix inversion method [18].

Computational light transport inversion. Clusters of camera-projector pixels are formed in [7], doing a brute-force transport inversion within clusters, but not considering intercluster interactions. This method is aimed at efficiency, but without clear error control. Iterative inverse methods for diffuse scenes are proposed in [19], which effectively solves (20) of this paper, but assuming a known form factor matrix. Recently, Ng et al. [5] developed a series expansion for inverse light transport, referred to as stratified inverses. We show that this series is a natural analog to the forward Neumann series. Our dual formulation enables us to go much further, clarifying the nature of convergence conditions. Most importantly, we derive new computational analogs to iterative finite element radiosity, as well as Monte Carlo methods. The resulting algorithms are significantly faster since they involve only matrix-vector multiplications (rather than matrix-matrix ones).

- M. Chandraker, J. Bai, and R. Ramamoorthi are with the Electrical Engineering and Computer Science Department, University of California, Berkeley, CA 94720. E-mail: {manukc, bjiamin, ravir}@eecs.berkeley.edu.
- T.-T. Ng is with the Institute for Infocomm Research, 1 Fusionopolis Way, Singapore 138632. E-mail: ttng@i2r.a-star.edu.sg.

Manuscript received 13 Dec. 2010; revised 8 Apr. 2011; accepted 2 May 2011; published online 7 June 2011.

Recommended for acceptance by C.-K. Tang.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2010-12-0947.

Digital Object Identifier no. 10.1109/TPAMI.2011.124.

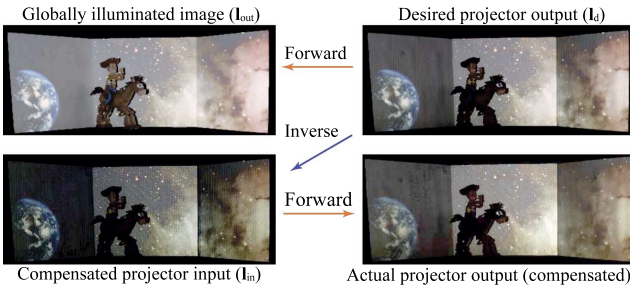


Fig. 1. Example application of inverse light transport for projector compensation in a real scene. Top: The desired projector output (right) leads to significant interreflections when displayed (left). Bottom: Our theory determines the pattern (left) whose projection is close to the desired (right). In effect, we have gone from global to local illumination, undoing interreflections. Our fast iterative method involves only matrix-vector multiplications and a Matlab implementation computes \mathbf{I}_{in} for a transport matrix of size $131K \times 131K$ in 2-3 sec.

Subsequent to this work, Krylov subspace methods have been implemented in the optical domain [20]. Acquisition and inverse iterations are coupled within the same loop in such active methods, which requires a more complex photometric setup. They require fewer images to invert a *specific* lighting configuration—in contrast, we acquire the entire transport matrix, but any input image can be inverted offline.

Practical applications. Usual approaches to radiometric compensation [21], [22] do not consider interreflections. Techniques for making one object look like another are also related [23]. Our main practical contribution is a suite of fast algorithms for light transport inversion. Another application is rapid direct and global separation for unstructured lighting, related to [24]. Our work requires prior acquisition of the light transport, but is valid for the entire light field and under multiple illuminants, while Nayar et al. [24] work for a single image, assuming a single source. Moreover, we can separate individual bounces of global illumination like [2] and can do so with much higher resolution transport matrices at interactive rates.

3 PRELIMINARIES

Due to the linearity of light transport, the image formation process is governed by a linear operator \mathbf{S} that encodes the effects of global illumination:

$$\mathbf{I}_{out} = \mathbf{S}\mathbf{I}_d, \quad (1)$$

where \mathbf{I}_{out} is the outgoing “global” light and \mathbf{I}_d is the direct lighting on surfaces due to external sources. In continuous form, \mathbf{I}_{out} and \mathbf{I}_d are functions (of spatial location and outgoing direction), while \mathbf{S} is a linear operator that accounts for global illumination. When discretized for practical applications, \mathbf{I}_{out} and \mathbf{I}_d are vectors, while \mathbf{S} is the interreflection matrix. Note that (1) depends only on linearity, and holds for the light field, as well as a single camera view (image).

Unlike forward global illumination, we do not see the light source directly, but rather its effect on the scene, which we




Fig. 2. Didactic real-data example illustrating separation of individual bounces of global illumination. The scene is a white concave dihedral, with flat green projection on the left half. Top row: Input image and separated direct and net global components. Bottom row: Recovered indirect bounces. Note that successive bounces illuminate alternating walls of the dihedral, as expected.

denote as the direct component, \mathbf{I}_d . The inverse light transport problem is simply

$$\mathbf{I}_d = \mathbf{S}^{-1}\mathbf{I}_{out}, \quad (2)$$

where we seek to invert the operator \mathbf{S} , undoing the effects of interreflections.

Practical issues. In practice, it is rare that \mathbf{S} is measured directly. Instead, a projector or illumination source lights the scene:

$$\mathbf{I}_d = \mathbf{F}\mathbf{I}_{in} \quad \mathbf{I}_{out} = \mathbf{T}\mathbf{I}_{in} = \mathbf{S}\mathbf{F}\mathbf{I}_{in}, \quad (3)$$

where \mathbf{I}_{in} is the incident projection and \mathbf{F} is the “first-bounce” matrix or operator. The actual acquired light transport is $\mathbf{T} = \mathbf{S}\mathbf{F}$. The above expression holds for any light transport acquisition system.

The remainder of the theoretical development focuses on analyzing and computing \mathbf{S}^{-1} . Eventual practical applications do need to convert from \mathbf{T} to \mathbf{S} , using $\mathbf{S} = \mathbf{T}\mathbf{F}^{-1}$. Moreover, applications like radiometric compensation actually seek to recover \mathbf{I}_{in} (rather than \mathbf{I}_d in (2)) given by $\mathbf{I}_{in} = \mathbf{T}^{-1}\mathbf{I}_{out} = \mathbf{F}^{-1}\mathbf{I}_d$.

Since we focus on global illumination \mathbf{S} , we consider setups where \mathbf{S} is easy to obtain from \mathbf{T} , that is, where \mathbf{F} is simple and at least approximately invertible. Therefore, we consider projector-based acquisition that illuminates a single spatial location, rather than light sources that illuminate the whole object (\mathbf{F} is low rank for diffuse surfaces [13]). After geometric calibration, we can use the same parameterization for projection and camera images [2]. \mathbf{F} is then a diagonal matrix, with \mathbf{F}^{-1} being trivial to compute.

Note that \mathbf{F} need not correspond to the actual first bounce for an accurate light transport inversion. In numerical terms, choosing $\mathbf{F} = \text{diag}(\mathbf{T})$ amounts to Jacobi preconditioning, which is convergent if \mathbf{T} is diagonally dominant (also see Section 8.2).

4 DUAL FORWARD-INVERSE TRANSPORT

In this section, we show that the rendering equation underlines a strong mathematical and physical duality between forward and inverse light transport. Key theoretical results are summarized in Fig. 3.

	Forward	Inverse
Problem	$\mathbf{I}_{out} = \mathbf{S}\mathbf{I}_d$	$\mathbf{I}_d = \mathbf{S}^{-1}\mathbf{I}_{out}$
Duality	$\mathbf{S} = (\mathbf{I} - \mathbf{A})^{-1}$	$\mathbf{S}^{-1} = (\mathbf{I} + \mathbf{R})^{-1}$
Series	$\mathbf{S} = \mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \dots$	$\mathbf{S}^{-1} = \mathbf{I} - \mathbf{R} + \mathbf{R}^2 - \dots$
Bounces	$\mathbf{S}_n = \sum_{k=0}^n \mathbf{A}^k = \mathbf{S} + O(\mathbf{A}^{n+1})$	$\mathbf{S}_n^{-1} = \mathbf{S}^{-1} + O(\mathbf{A}^{n+1})$
Iteration	$\mathbf{I}_{out}^{(k)} = \mathbf{I}_d + \mathbf{A}\mathbf{I}_{out}^{(k-1)}$	$\mathbf{I}_d^{(k)} = \mathbf{I}_{out} - \mathbf{R}\mathbf{I}_d^{(k-1)}$
Monte Carlo	$\mathbf{I}_d(i_0) + \sum_k \sum_i \mathbf{A}_{i_0 i_1} \mathbf{A}_{i_1 i_2} \dots \mathbf{I}_d(i_k)$	$\mathbf{I}_{out}(i_0) + \sum_k (-1)^k \sum_i \mathbf{R}_{i_0 i_1} \mathbf{R}_{i_1 i_2} \dots \mathbf{I}_{out}(i_k)$

Fig. 3. Duality of forward and inverse light transport, indicating analogous relations for some key properties. (Monte Carlo equations abbreviated, full forms in the text.)

In the notation of [15], the rendering equation is

$$\mathbf{l}_{\text{out}} = \mathbf{l}_d + \mathbf{K}\mathbf{G}\mathbf{l}_{\text{out}} \Rightarrow \mathbf{l}_{\text{out}} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{l}_d, \quad (4)$$

where \mathbf{K} considers the local reflection at a surface, governed by the BRDF, \mathbf{G} is a geometric operator that transports outgoing to incident radiance, and $\mathbf{A} = \mathbf{K}\mathbf{G}$ corresponds to one physical bounce of light. Note that this formulation is valid for any opaque BRDF when considering the full light field. While our theory is fully general, our experiments will consider single view projections, which introduce practical limitations, as discussed in Section 8. It follows that

$$\mathbf{S} = (\mathbf{I} - \mathbf{A})^{-1}. \quad (5)$$

This well-known result shows that the forward problem formally involves a matrix or operator inversion. Also, as noted in [2] and [25], if the scene geometry and reflectance (and hence \mathbf{A}) are known, we simply have $\mathbf{S}^{-1} = \mathbf{I} - \mathbf{A}$. We focus here on cases where we only measure \mathbf{S} , but do not know or compute \mathbf{A} .

We can separate \mathbf{l}_{out} into direct \mathbf{l}_d and indirect or global \mathbf{l}_g components:

$$\mathbf{l}_{\text{out}} = \mathbf{l}_d + \mathbf{l}_g = \mathbf{l}_d + \mathbf{R}\mathbf{l}_d \quad \mathbf{l}_{\text{out}} = (\mathbf{I} + \mathbf{R})\mathbf{l}_d, \quad (6)$$

where $\mathbf{R} = \mathbf{S} - \mathbf{I}$ is a linear operator that accounts only for global illumination. We are now ready to present an expression for inverse light transport:

$$\mathbf{S}^{-1} = (\mathbf{I} + \mathbf{R})^{-1}. \quad (7)$$

The very similar or dual forms of (5) and (7) are a key insight in this paper, allowing us to directly leverage many forward rendering theories and algorithms for inverse rendering in computer vision.

4.1 Neumann Forward and Inverse Series

The forward equations (4) and (5) have series expansions corresponding physically to light bounces:

$$\mathbf{S} = \mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \mathbf{A}^3 + \dots \quad (8)$$

Relating the global illumination operator \mathbf{R} to (8):

$$\mathbf{R} = \mathbf{S} - \mathbf{I} = \mathbf{A} + \mathbf{A}^2 + \mathbf{A}^3 + \dots, \quad (9)$$

the dual series analogous to (8) is

$$\mathbf{S}^{-1} = \mathbf{I} - \mathbf{R} + \mathbf{R}^2 - \mathbf{R}^3 + \dots \quad (10)$$

Note that the positive sign of \mathbf{R} implies the series is oscillatory. Intuitively, from (6), $\mathbf{l}_d = \mathbf{l}_{\text{out}} - \mathbf{R}\mathbf{l}_d$. Since the unknown \mathbf{l}_d appears on the right-hand side, a first approximation as $\mathbf{l}_d \approx \mathbf{l}_{\text{out}}$ calculates $\mathbf{l}_d \approx \mathbf{l}_{\text{out}} - \mathbf{R}\mathbf{l}_{\text{out}}$. This overcompensation is corrected by higher order terms, leading to the alternating signs in (10).

With suitable algebraic manipulations, one may note that (10) explains the stratified inverses of Ng et al. [5] and relates it to the rendering equation.¹

4.2 Interpretation as Physical Bounces of Light

The results of this section were originally derived in [26], where we point the reader for complete details. Here, we present a brief summary.

Consider an approximation up to order n that we denote as \mathbf{S}_n or \mathbf{S}_n^{-1} . In the forward case, this physically corresponds to considering the first n bounces:

1. In particular, note that $\mathbf{R} = \mathbf{S} - \mathbf{I}$, which is $\mathbf{T}\mathbf{F}^{-1} - \mathbf{I}$. A final binomial expansion in $\mathbf{T}\mathbf{F}^{-1}$ and using $\mathbf{T}^{-1} = \mathbf{F}^{-1}\mathbf{S}^{-1}$ enable one to derive the results in [5].

$$\mathbf{S}_n = \sum_{k=0}^n \mathbf{A}^k, \quad \mathbf{S}_n - \mathbf{S} = O(\mathbf{A}^{n+1}). \quad (11)$$

A physical interpretation for the inverse series seems non-intuitive since (10) is expressed in terms of \mathbf{R} that includes all global illumination terms. Nevertheless, in [26], we derive a surprising result: Each term of the inverse series *cancels or zeroes out the corresponding bounce of light transport*. We show

$$\mathbf{S}_n^{-1} = \mathbf{I} - \mathbf{A} + \sum_{j=2}^{\infty} \left[\sum_{k=1}^{\min(j,n)} (-1)^k \binom{j-1}{k-1} \right] \mathbf{A}^j, \quad (12)$$

which, using $\mathbf{S}^{-1} = \mathbf{I} - \mathbf{A}$, leads to a key result that the \mathbf{A}^j terms vanish for $2 \leq j \leq n$:

$$\mathbf{S}_n^{-1} = \mathbf{I} - \mathbf{A} + O(\mathbf{A}^{n+1}), \quad \mathbf{S}_n^{-1} - \mathbf{S}^{-1} = O(\mathbf{A}^{n+1}). \quad (13)$$

Note that, as opposed to the forward series where the higher bounces are simply 0 until they are added in, the values for the higher bounces in the inverse series oscillate until they are zeroed—this is related to the oscillatory convergence of the inverse series.

5 CONVERGENCE AND ERROR ANALYSIS

For the forward case, we note the results of [15]. For a closed enclosure, $\|\mathbf{G}\| = 1$ (less for open scenes). By energy conservation, excluding perfect reflectors, $\|\mathbf{K}\| \leq m < 1$, where m relates to surface albedo (for nondiffuse materials, it is the maximum over all incident directions of the fraction of total energy reflected).² Since $\|\mathbf{A}\| \leq \|\mathbf{K}\| \|\mathbf{G}\|$, it follows that $\|\mathbf{A}\| \leq m < 1$, so the forward series always converges.

For the inverse series in (10), a bound from (9) is

$$\|\mathbf{R}\| \leq \|\mathbf{A}\| + \|\mathbf{A}^2\| + \dots \leq m + m^2 + \dots = \frac{m}{1-m}. \quad (14)$$

If $m < \frac{1}{2}$, we obtain $\|\mathbf{R}\| < 1$, which is sufficient for convergence (though *not* necessary). Intuitively, if the diffuse albedo (or maximum fraction of energy reflected for any incident direction for nondiffuse materials) is less than 1/2, the norm of the total global illumination operator \mathbf{R} is less than that of the direct lighting operator \mathbf{I} . In matrix terms, $\mathbf{S} = \mathbf{I} + \mathbf{R}$ is diagonally dominant. Since the inverse series is oscillatory, we require bounding the full global illumination, rather than just each bounce separately, as in the forward case.

Error analysis. The error introduced in an n -term expansion for forward and inverse series (denoted \mathbf{S}_n or \mathbf{S}_n^{-1}) can be easily bounded. In the forward case,

$$\|\mathbf{S} - \mathbf{S}_n\| \leq \sum_{k=n+1}^{\infty} \|\mathbf{A}^k\| \leq \sum_{k=n+1}^{\infty} m^k = \frac{m^{n+1}}{1-m}. \quad (15)$$

Similarly, for the inverse series,

$$\begin{aligned} \|\mathbf{S}^{-1} - \mathbf{S}_n^{-1}\| &\leq \sum_{k=n+1}^{\infty} \|\mathbf{R}^k\| \leq \sum_{k=n+1}^{\infty} \left(\frac{m}{1-m} \right)^k \\ &= \frac{m^{n+1}}{(1-m)^n(1-2m)}. \end{aligned} \quad (16)$$

Numerical simulations. For simplicity, we consider a synthetic diffuse box (closed, so $\|\mathbf{G}\| = 1$), without shadows but with interreflections. Fig. 4 assumes that \mathbf{l}_d is constant on each surface, which have different albedos. From left to right, addition of more terms from (10) causes oscillations between over and under-compensating interreflections, till convergence to \mathbf{l}_d . Interestingly,

2. These relations hold in any L^p norm since, from reciprocity, $\|\mathbf{K}\|_1 = \|\mathbf{K}\|_{\infty} = p$ and $\|\cdot\|_p \leq \max(\|\cdot\|_1, \|\cdot\|_{\infty})$.




Fig. 4. Top: From left to right, we add more terms of the inverse series, going from the simulated global illumination \mathbf{I}_{out} to the “direct lighting” \mathbf{I}_d (leftmost). These terms also correspond to the iterations in Section 6. Bottom: Contributions of individual terms (neutral gray is 0).

while forward global illumination in \mathbf{I}_{out} results in predictable red and green color bleeding, odd terms of the inverse series give rise to cyan and magenta colors. The final inverse light transport solution for \mathbf{I}_d has no color bleeding, as desired.

In Fig. 5, we analyze errors and convergence. Fig. 5a indicates similar oscillatory convergence behavior near corners, edges, and face centers. Fig. 5b shows excellent agreement, up to a constant factor, between error for the whole \mathbf{S}^{-1} operator and the theoretical bound in (16). Fig. 5c illustrates the inverse relation of convergence rate and albedo. Even albedos near the theoretical limit (like 0.45) converge in a few iterations; those very close to 0.5 converge slowly and those greater than 0.51 diverge. Fig. 5d shows the variation of convergence with geometry (that is, $\|\mathbf{G}\|$). For an albedo of 0.62, close to the theoretical limit for a five-sided box, we observe very slow convergence for a five-sided box, divergence for a six-sided box, and rapid convergence for more open geometries.

Finally, Fig. 6 shows a scene with occlusions and glossy surfaces. Similar behaviors hold as above, with convergence of the inverse series to direct lighting.

6 FAST ITERATIVE COMPUTATION

In this section, we introduce the first of our algorithmic contributions—a fast iterative method to compute inverse light transport, using only matrix-vector multiplications. This method is

dual to iterative forward rendering methods like finite element radiosity. We also explore analogous wavelet accelerations.

For forward rendering, one rarely computes the series in (8) to explicitly determine \mathbf{S} . This is mainly because of the high cost of matrix-matrix multiplications on high-resolution scenes. Instead, finite element and radiosity methods [9] try to solve

$$\mathbf{I}_{\text{out}} = \mathbf{I}_d + \mathbf{A}\mathbf{I}_{\text{out}} \quad (17)$$

iteratively, which corresponds to (4). This iteration is numerically stable and requires only the matrix-vector multiplication for $\mathbf{A}\mathbf{I}_{\text{out}}$. Each step computes

$$\mathbf{I}_{\text{out}}^{(k)} = \mathbf{I}_d + \mathbf{A}\mathbf{I}_{\text{out}}^{(k-1)}, \quad (18)$$

where the superscript stands for step k , with $\mathbf{I}_{\text{out}}^{(0)} = \mathbf{I}_d$. Importantly, note that n steps correspond simply to computing the effect of the first n terms of (8).

Using (6), one can similarly derive the analog to (17) for inverse rendering:

$$\mathbf{I}_d = \mathbf{I}_{\text{out}} - \mathbf{R}\mathbf{I}_d. \quad (19)$$

With $\mathbf{I}_d^{(0)} = \mathbf{I}_{\text{out}}$, the iterative solution dual to (18) is

$$\mathbf{I}_d^{(k)} = \mathbf{I}_{\text{out}} - \mathbf{R}\mathbf{I}_d^{(k-1)}, \quad (20)$$

where again the first n steps correspond to the first n terms in (10). Note $-\mathbf{R}$ (compare to $+\mathbf{A}$ in (18)), corresponding to the oscillatory nature of the series.

As written, (20) and (18) correspond to the Jacobi iteration for solving systems of linear equations [16]. If \mathbf{I}_d are updated in place (instead of at the end of a step), this is the Gauss-Seidel method. Both techniques are popular in forward radiosity. By framing inverse light transport as dual to forward transport, we could also leverage other computational methods in the future, such as Southwell iteration, successive overrelaxation, and conjugate gradient solutions.

Matrix iteration. To precompute \mathbf{S} or \mathbf{S}^{-1} , there are corresponding iterations for the full matrix:

$$\mathbf{S}_k = \mathbf{I} + \mathbf{A}\mathbf{S}_{k-1}, \quad \mathbf{S}_k^{-1} = \mathbf{I} - \mathbf{R}\mathbf{S}_{k-1}^{-1}, \quad (21)$$

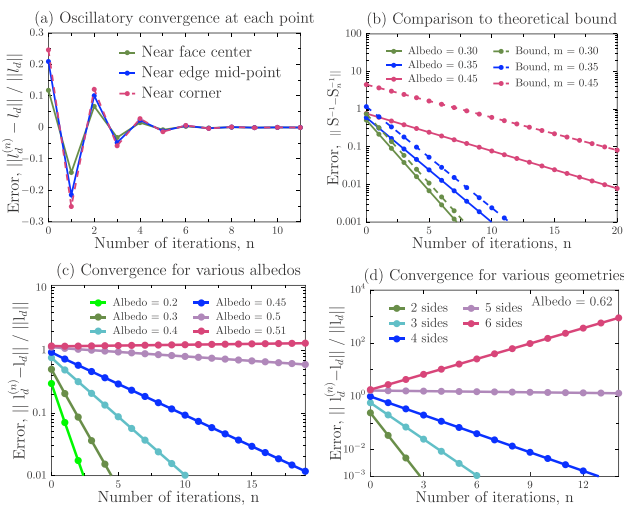


Fig. 5. Error analysis for convergence of inverse series. (a) Similar convergence behavior at different points (center, edge, and corner). (b) Comparison of error to theoretical bound for different albedos showing good agreement. (c) Convergence for different albedos. As predicted by theory, convergence is faster for lower albedos, up to the limit of 0.5. An albedo of 0.51 leads to divergence. (d) An albedo of 0.62 diverges for a closed box (six sides) and shows very slow convergence, as expected for a five-sided box, but rapid convergence for more open environments (fewer sides, smaller $\|\mathbf{G}\|$).




Fig. 6. Validation of the theory for shadowed and non-Lambertian scenes. Our iterative method recovers \mathbf{I}_d in 10 iterations for the shadowed scene and 20 for the glossy one.




Fig. 7. Timings for series, iterative finite element, and wavelet accelerated methods (using Daubechies4 wavelets). N is the transport resolution (matrix size is N^2). We normalize timings so that 1.0 corresponds to 5.57×10^{-4} sec, with experiments in Matlab on an Intel i7 machine. All methods are run until 1 percent error.

with $S_0 = \mathbf{I}$ and $S_0^{-1} = \mathbf{I}$. These iterations are not significantly more efficient than a direct factorization for (8) and (10) and, thus, rarely used. But they do provide an elegant and numerically stable iterative scheme.

Wavelet and hierarchical methods. The matrix-vector multiplication $\mathbf{R}\mathbf{l}_d$ in (20) is the time-consuming step. We can wavelet-transform and approximate the vector \mathbf{l}_d , as well as the rows of \mathbf{R} , to speed up the matrix-vector multiply. This is analogous to wavelet radiosity and light transport in forward rendering [27], [28]. Other hierarchical approaches, analogous to [29], can also be explored.

Numerical simulations. As baseline, we use matrix-matrix multiplications for the series in (10) (explicit matrix inversion is intractable for high resolutions). In Fig. 7, we compare to iterations in (20) and wavelet accelerations. For transport resolution N , the series method scales as $O(N^3)$ and rapidly becomes impractical. The iterative method uses only matrix-vector multiplications and is much faster $O(N^2)$, with a speedup of three orders of magnitude for large sizes. Wavelet acceleration leads to linear $O(NW)$ performance, where the number of wavelets W in each row is relatively insensitive to N . The benefits are more noticeable at higher resolutions, where wavelet sparsity outweighs the transform overhead—wavelets provide significant savings at real-world resolutions.

7 MONTE CARLO ALGORITHMS

Besides finite element methods like radiosity, forward rendering has developed a suite of Monte Carlo techniques [1]. Treating \mathbf{A} as a matrix, we need to consider all permutations of indices:

$$\mathbf{l}_{\text{out}}(i_0) = \mathbf{l}_d(i_0) + \sum_{k=1}^{\infty} \sum_{i_1, i_2, \dots, i_k} \mathbf{A}_{i_0 i_1} \mathbf{A}_{i_1 i_2} \cdots \mathbf{A}_{i_{k-1} i_k} \mathbf{l}_d(i_k), \quad (22)$$

where the first summation is over all terms k in the series or all path lengths in a path tracing context. The different indices correspond to all matrix sums, or equivalently all paths, where each i_j chooses a particular point on the path. In Monte Carlo path tracing, essentially the above form is implicitly used, but the \mathbf{A} matrix is not usually computed explicitly, and elements of it are generated on the fly.

The inverse series in (10) has an analogous form

$$\mathbf{l}_d(i_0) = \mathbf{l}_{\text{out}}(i_0) + \sum_{k=1}^{\infty} (-1)^k \sum_{i_1, i_2, \dots, i_k} \mathbf{R}_{i_0 i_1} \mathbf{R}_{i_1 i_2} \cdots \mathbf{R}_{i_{k-1} i_k} \mathbf{l}_{\text{out}}(i_k), \quad (23)$$

where the oscillatory behavior requires the additional $(-1)^k$ factor. A direct Monte Carlo algorithm is to use a number of samples, for each of which the indices i_1, i_2, \dots, i_k are drawn at random. The expectation of these samples gives the desired result. Our implementation makes a number of optimizations, corresponding to analogous techniques in forward rendering.

First, for each sample, we choose a path length k . We assign probabilities to different path lengths in proportion to their expected contribution, which decays with k . From the convergence and error analysis in (16), we use the normalized probability

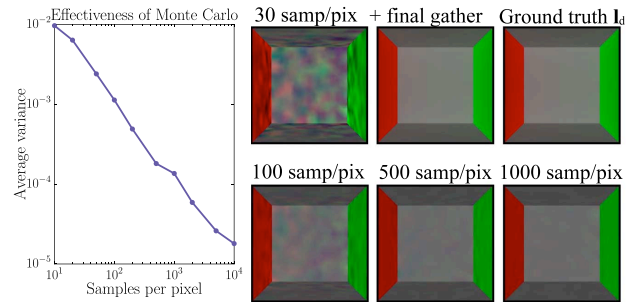


Fig. 8. Left: Graph of variance in Monte Carlo methods, which varies inversely with the number of samples per pixel, as expected. Right: In the top row, we show that only 30 samples per pixel (that in itself is extremely noisy) are adequate to produce good results using final gather. In the bottom row, as expected, Monte Carlo becomes more accurate with more samples. The transport resolution in these experiments is $N = 5,120$.

$$p(k) = \binom{m}{1-m} \frac{k(1-2m)}{m}, \quad (24)$$

with m being an estimate of the average albedo of the scene. We next choose indices i_1, i_2, \dots, i_k . These can be chosen randomly or we can use importance sampling on each row of the matrix \mathbf{R} :

$$p(i_j | i_{j-1}) = \frac{\mathbf{R}_{i_{j-1} i_j}}{\sum_{i_p} \mathbf{R}_{i_{j-1} i_p}} = \frac{\mathbf{R}_{i_{j-1} i_j}}{|\mathbf{R}_{i_{j-1}}|}, \quad (25)$$

where we normalize by the sum of elements over the full row and the last step simply denotes the row sum as $|\mathbf{R}_{i_{j-1}}| = \sum_{i_p} \mathbf{R}_{i_{j-1} i_p}$. These row sums correspond to a generalized analog of albedos or BRDFs along the path, suitably weighted. Note that importance sampling in forward rendering is usually based on some partial information like lighting or BRDF, but we have the luxury of the full \mathbf{R} matrix to importance sample.³ This greatly simplifies the final expressions.

Finally, the net image is just the expected value over all paths/samples. For each path, we must divide the value f by the probability (in this case, f is simply the appropriate term in the summation on the right-hand side of (23)). Since f involves expressions of the form $\mathbf{R}_{i_{j-1} i_j}$, they cancel with the above probabilities, as they should for good importance sampling:

$$\frac{f(k; i_1, i_2, \dots, i_k)}{p(k; i_1, i_2, \dots, i_k)} = \frac{(-1)^k}{p(k)} |\mathbf{R}_{i_0}| |\mathbf{R}_{i_1}| \cdots |\mathbf{R}_{i_{k-1}}| \mathbf{l}_{\text{out}}(i_k), \quad (26)$$

where we must average (take the expected value) over all samples to obtain \mathbf{l}_d and also add the initial term $\mathbf{l}_{\text{out}}(i_0)$, as per (23). Note that the above simplified form is valid only if we importance sample properly when choosing the next index along a path.

Hybrid methods. Besides the above pure Monte Carlo path tracing analog, we can also explore hybrids of iterative and Monte Carlo techniques. For example, we can speed up the matrix-vector multiplication $\mathbf{R}\mathbf{l}_d$ in (20) by Monte Carlo sampling only some of the columns for each row, using the importance sampling scheme above. We also explore an analogy to final gather in forward rendering where we use fewer samples for the iteration, but then compute the final step with a direct matrix-vector multiplication. For many applications, this may be sufficiently accurate and faster than a full multistep iterative approach. We point the reader to [30] for a detailed exposition on hybrid methods, while noting that our duality theory ensures that the benefits observed there for forward rendering translate directly to the inverse problem.

Numerical simulations. Fig. 8 first demonstrates that variance varies inversely with the number of samples per pixel as expected.

3. Note that building the probability tables for importance sampling requires a preprocess for each row of the matrix. This is done once, after acquisition and before any specific \mathbf{l}_{out} is chosen.




Fig. 9. Approximate bounce separation with occlusions and specularities. Top row: Input image and separated direct and net global components. Bottom row: Recovered indirect bounces. Note that successive bounces illuminate alternating walls and the specular highlight is present only in the direct component.

The top row shows the power of final gather—Monte Carlo with 30 samples is noisy, but is smoothed out almost completely using one direct iteration (the final gather). In the bottom row, we see that, as expected, pure Monte Carlo converges as the number of samples is increased.⁴

8 EXPERIMENTS WITH REAL DATA

In this section, we illustrate practical applications of our iterative inverse light transport algorithm. The accuracy of our algorithms is established by a few didactic examples, while their computational utility is demonstrated on high-resolution transport matrices. Where appropriate, we discuss the limitations imposed by our choice of experimental conditions.

8.1 Acquisition Details

Our acquisition setup consists of a Dell 4310WX projector and a Canon EOS 5D Mark II camera. An accurate, one-time, radiometric calibration of the projector and camera response curves is performed to ensure linearity of the corresponding signals. While prior work has obtained transport matrices at comparable resolutions [31], it has mainly been for applications akin to relighting. In contrast, our inverse applications require greater fidelity of the transport matrix. Thus, a judicious consideration of a signal to noise ratio is necessary to capture as many of the weaker interreflection bounces as possible while discarding sensor noise. To faithfully capture the energy of the transport matrix, up to eight images at various exposures are assembled into a high dynamic range image. The projector’s black offset is computed at the highest exposure to average out high-frequency fluctuations. For higher resolution scenes, a hierarchical subdivision scheme is used to simultaneously acquire portions of the transport matrix which are not in mutual conflict [31].

Limitations. Our setup shares the restrictions of other projector-camera systems, such as shutter speeds limited by projector refresh rates, color bleeding, and nonlinear color mixing ratios. Clipping artifacts may arise in applications due to the inability of projectors to display negative values. Note that we require a prior acquisition of \mathbf{T} , for which methods like [31] may become expensive under strong interreflections.

8.2 Projector Radiometric Compensation

The ubiquitous use of projectors may necessitate inverting photometric distortions and interreflection effects to simulate any desired appearance in nonflat, non-Lambertian spaces. In terms of our theory, given a desired appearance \mathbf{l}_{out} , we seek to invert the light transport to find $\mathbf{l}_{\text{d}} = \mathbf{S}^{-1}\mathbf{l}_{\text{out}}$. As discussed in Section 3, we

4. Our Matlab implementation is not optimized for the sampling process, making a direct timing comparison to finite elements difficult. Hence, we simply report on the number of samples per pixel.

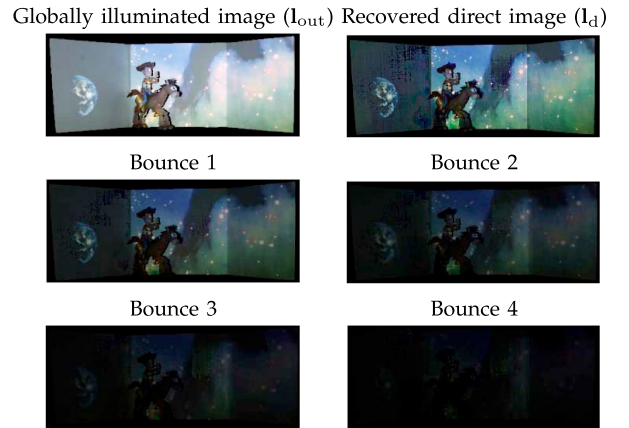


Fig. 10. Approximate bounce separation on a complex, high-resolution scene using our iterative light transport inversion. Each bounce is obtained in just 3 sec. for a $131K \times 131K$ light transport.

must account for the first bounce \mathbf{F} from the projector, and actually compute $\mathbf{l}_{\text{in}} = \mathbf{T}^{-1}\mathbf{l}_{\text{out}}$.

Fig. 1 shows results for radiometric compensation to project a desired image onto a scene with non-Lambertian materials, occlusions, and interreflections. Clearly, the desired appearance is closely matched. The size of the transport matrix is $131K \times 131K$, for which our iterative algorithm performs radiometric compensation in only about 3 sec. While such high resolutions may be infeasible for a straightforward matrix inversion, based on the patterns in Fig. 7, the stratified inverses method in [5] will require one to two orders of magnitude more time. Also, in contrast to the method in [7], our algorithms are physically motivated and not contingent on any tunable parameters.

Note on generality. Even when $\mathbf{F} = \text{diag}(\mathbf{T})$ is not the actual “first-bounce” operator, it is a valid Jacobi preconditioner, which leads to a convergent inversion for any light transport as long as global effects do not dominate. Thus, our radiometric compensation is accurate for any opaque BRDF, including non-Lambertian ones.

8.3 Separating Bounces

One consequence of our theory is that once the light transport has been acquired, we can quickly separate an image into its constituent interreflection bounces. It follows from (18), noting that $\mathbf{S}^{-1} = \mathbf{I} - \mathbf{A}$, that the k th indirect bounce is

$$\mathbf{l}_{\text{out}}^{(k+1)} - \mathbf{l}_{\text{out}}^{(k)} = \mathbf{l}_{\text{d}} - \mathbf{S}^{-1}\mathbf{l}_{\text{out}}^{(k)}. \quad (27)$$

Thus, each successive run of our iterative inversion algorithm yields a bounce of light transport. Fig. 2 demonstrates this on a didactic example. The scene consists of a white dihedral with green light projected on the left half. Note that successive bounces of indirect illumination in the bottom row alternate between the two walls, as expected. Fig. 9 demonstrates the same with a non-Lambertian occluder present in the scene. We observe that the specular highlight is limited to the direct component and absent from the indirect bounces, which is also expected.

This application is the same as [2], but our algorithms are far more efficient. For instance, our iterative method recovers the direct component as well as each bounce of indirect illumination in 0.09 sec for the $4K \times 4K$ transport matrix in Fig. 9, while straightforward matrix inversion requires 4.6 sec. More importantly, our methods can efficiently operate on much higher resolution scenes that direct inversion cannot handle—for instance, Fig. 10 demonstrates bounce separation in a $131K \times 131K$ transport matrix. While an uncompressed matrix of that size cannot even be loaded in RAM, extrapolating from Fig. 7, a brute force inversion will require nearly 150 hours. In contrast, we require only 33 msec per iteration in our (unoptimized) Matlab implementation, for a

total of about 3 sec to separate each bounce. Note that the faster method in [24] yields only the top row of Fig. 2 for a *particular* lighting configuration, while we can separate all of the bounces for *any* lighting, albeit at the expense of a more laborious acquisition.

Approximations in single projector-camera setup. While our theory is valid for all opaque BRDFs when considering the full light field, the use of a single projector and camera in our applications makes the bounce separation only approximately correct. A single projector-camera setup necessitates additional operators to project the output light field to the image (\mathbf{P}) and raise the projector input to the full light field (\mathbf{Q}). The light transport now becomes

$$\mathbf{T} = \mathbf{P}\mathbf{S}\mathbf{F}\mathbf{Q} = \mathbf{P}(\mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \dots)\mathbf{F}\mathbf{Q} \quad (28)$$

$$= \mathbf{P}\mathbf{F}\mathbf{Q} + \mathbf{P}\mathbf{A}\mathbf{F}\mathbf{Q} + \mathbf{P}\mathbf{A}^2\mathbf{F}\mathbf{Q} + \dots \quad (29)$$

The direct lighting component in the observed image is $\mathbf{P}\mathbf{F}\mathbf{Q}$. Higher bounces for the full light field are generated by a simple operator action \mathbf{A} . However, that is not true here unless \mathbf{P} and \mathbf{A} commute multiplicatively (for which it is unlikely that any physical meaning exists). Thus, our physical interpretations in terms of bounces of light in Section 4.2 are only valid for the full light field. Also, this makes the theory inexact for the bounce separation application with a single projector-camera in non-Lambertian scenes.

But note that the bounce separation is exact for Lambertian scenes even in the single projector-camera case. Since the camera direction is immaterial for Lambertian scenes, one need not consider \mathbf{P} for a radiometric analysis (or even \mathbf{Q} , if one ignores visibility and shadowing issues, as in [2]). In practice, specular effects rapidly decay with bounces of interreflection (that is, higher bounces are increasingly diffuse), so our bounce separation experiments in Figs. 9 and 10 are still robust to moderate amount of gloss.

9 CONCLUSIONS AND FUTURE WORK

The main contribution of this paper is a formulation of inverse light transport in computer vision, as a dual to the theory of forward rendering in computer graphics. This lends new insights for canceling interreflections in complex scenes, as well as fast computational methods for doing so. Our efficient algorithms, analogous to finite element radiosity and Monte Carlo path tracing in forward rendering, can handle transport resolutions far higher than previous methods.

From a theoretical perspective, we have just scratched the surface of analogies between forward and inverse methods. It is our hope that the framework of this paper forms the basis for discovering further insights into the structure of light transport and developing methods that couple fast acquisition and iterative inversion to handle dynamic scenes.

ACKNOWLEDGMENTS

This work was funded by US Office of Naval Research (ONR) YIP grant N00014-10-1-0032, ONR PEACASE grant N00014-09-1-0741, a National Science Scholarship from the A*STAR Graduate Academy of Singapore, as well as generous support from Adobe, NVIDIA, Intel, and Pixar. The authors would like to thank Joo Hwee Lim and Zhiyong Huang for kind support at I²R and the anonymous reviewers of this paper and [11] for useful comments.

REFERENCES

[1] J. Kajiya, "The Rendering Equation," *Proc. ACM SIGGRAPH*, pp. 143-150, 1986.

[2] S. Seitz, Y. Matsushita, and K. Kutulakos, "A Theory of Inverse Light Transport," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 1440-1447, 2005.

[3] P. Peers, K. Berge, W. Matusik, R. Ramamoorthi, J. Lawrence, S. Rusinkiewicz, and P. Dutre, "A Compact Factored Representation of Heterogeneous Subsurface Scattering," *ACM Trans. Graphics*, vol. 25, no. 3, pp. 746-753, 2006.

[4] V. Masselus, P. Peers, P. Dutre, and Y. Willems, "Relighting with 4D Incident Light Fields," *ACM Trans. Graphics*, vol. 22, no. 3, pp. 613-620, 2003.

[5] T.-T. Ng, R.S. Pahwa, J. Bai, Q.-S. Quek, and K.-H. Tan, "Radiometric Compensation Using Stratified Inverses," *Proc. IEEE Int'l Conf. Computer Vision*, 2009.

[6] P. Debevec, T. Hawkins, C. Tchou, H. Duiker, W. Sarokin, and M. Sagar, "Acquiring the Reflectance Field of a Human Face," *Proc. ACM SIGGRAPH*, pp. 145-156, 2000.

[7] G. Wetzstein and O. Bimber, "Radiometric Compensation through Inverse Light Transport," *Proc. Pacific Conf. Computer Graphics and Applications*, pp. 391-399, 2007.

[8] S. Liu, T.-T. Ng, and Y. Matsushita, "Shape from Second Bounce of Light Transport," *Proc. European Conf. Computer Vision*, 2010.

[9] M. Cohen and J. Wallace, *Radiosity and Realistic Image Synthesis*. Academic Press, 1993.

[10] E. Veach, "Robust Monte Carlo Methods for Light Transport Simulation," PhD dissertation, Stanford Univ., 1998.

[11] J. Bai, M. Chandraker, T.-T. Ng, and R. Ramamoorthi, "A Dual Theory of Inverse and Forward Light Transport," *Proc. European Conf. Computer Vision*, 2010.

[12] S. Marschner, "Inverse Rendering for Computer Graphics," PhD dissertation, Cornell Univ., 1998.

[13] R. Ramamoorthi and P. Hanrahan, "A Signal-Processing Framework for Inverse Rendering," *Proc. ACM SIGGRAPH*, pp. 117-128, 2001.

[14] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs," *Proc. ACM SIGGRAPH*, pp. 215-224, 1999.

[15] J. Arvo, K. Torrance, and B. Smits, "A Framework for the Analysis of Error in Global Illumination Algorithms," *Proc. ACM SIGGRAPH*, pp. 75-84, 1994.

[16] G. Golub and C. van Loan, *Matrix Computations*. John Hopkins Univ. Press, 1996.

[17] J. Demmel, *Applied Numerical Linear Algebra*. SIAM, 1997.

[18] G. Forsythe and R. Leibler, "Matrix Inversion by a Monte Carlo Method," *Math. Tables and Other Aids to Computation*, vol. 4, no. 31, pp. 127-129, July 1950.

[19] O. Bimber, A. Grundhoefler, T. Zeidler, D. Danch, and P. Kapakos, "Compensating Indirect Scattering for Immersive and Semi-Immersive Projection Displays," *Proc. IEEE Virtual Reality*, pp. 151-158, 2006.

[20] M. O'Toole and K. Kutulakos, "Optical Computing for Fast Light Transport Analysis," *Proc. SIGGRAPH Asia 2010 Papers*, 2010.

[21] K. Fujii, M. Grossberg, and S. Nayar, "A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2005.

[22] Y. Ding, J. Xiao, K.-H. Tan, and J. Yu, "Catadioptric Projectors," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.

[23] R. Raskar, G. Welch, K. Low, and D. Bandyopadhyay, "Shader Lamps," *Proc. EuroGraphics Workshop Rendering*, 2001.

[24] S. Nayar, G. Krishnan, M. Grossberg, and R. Raskar, "Fast Separation of Direct and Global Components of a Scene Using High Frequency Illumination," *ACM Trans. Graphics*, vol. 25, no. 3, pp. 935-944, 2006.

[25] Y. Mukaigawa, T. Kakinuma, and Y. Ohta, "Analytical Compensation of Inter-Reflection for Pattern Projection," *Proc. ACM Symp. Virtual Reality Software and Technology (VRST)*, pp. 265-268, 2006.

[26] J. Bai, M. Chandraker, T.-T. Ng, and R. Ramamoorthi, "A Dual Theory of Inverse and Forward Light Transport," Technical Report UCB/EECS-2010-101, Univ. of California Berkeley, 2010.

[27] S. Gortler, P. Schröder, M. Cohen, and P. Hanrahan, "Wavelet Radiosity," *Proc. ACM SIGGRAPH*, pp. 221-230, 1993.

[28] R. Ng, R. Ramamoorthi, and P. Hanrahan, "All-Frequency Shadows Using Non-Linear Wavelet Lighting Approximation," *ACM Trans. Graphics*, vol. 22, no. 3, pp. 376-381, 2003.

[29] P. Hanrahan, D. Salzman, and L. Uppерle, "A Rapid Hierarchical Radiosity Algorithm," *Proc. ACM SIGGRAPH*, pp. 197-206, 1991.

[30] P.H. Christensen, E.J. Stollnitz, D.H. Salesin, and T.D. DeRose, "Global Illumination of Glossy Environments Using Wavelets and Importance," *ACM Trans. Graphics*, vol. 15, pp. 37-71, 1996.

[31] P. Sen, B. Chen, G. Garg, S. Marschner, M. Horowitz, M. Levoy, and H. Lensch, "Dual Photography," *ACM Trans. Graphics*, vol. 24, no. 3, pp. 745-755, 2005.