

# A Dual Theory of Inverse and Forward Light Transport

Jiamin Bai<sup>1</sup>, Manmohan Chandraker<sup>1</sup>, Tian-Tsong Ng<sup>2</sup>, and Ravi Ramamoorthi<sup>1</sup>

<sup>1</sup> University of California, Berkeley

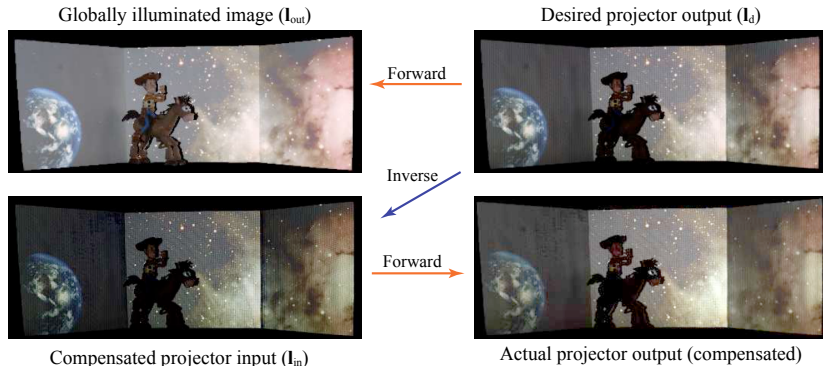
<sup>2</sup> Institute for Infocomm Research, Singapore

**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, we show 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 transport matrix, which is impractical for realistic resolutions. A natural consequence of our theoretical framework is a suite of fast computational algorithms for light transport inversion – analogous 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.

## 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 [9]. In contrast, most computer vision algorithms are forced to simply ignore interreflections, where one would ideally like to invert the rendering equation to undo their effects. Recently, Seitz et al. [18] 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 already 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.



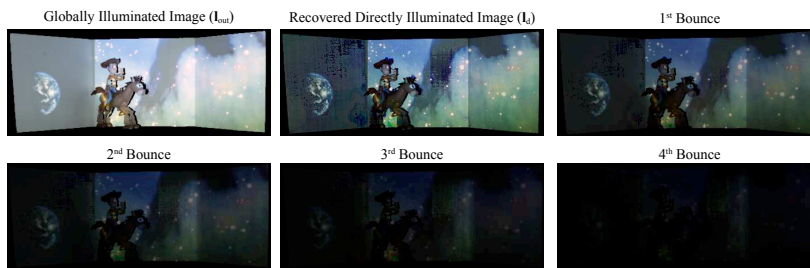
**Fig. 1.** Application of inverse light transport for projector compensation in a real scene. **Top:** The desired projector output (right) leads to significant interreflections when displayed (left). **Bottom:** Our theory determines the pattern (left) whose projection is close to the desired (right). Our fast iterative method involves only matrix-vector multiplications, with each iteration taking only 0.03 sec. For a transport matrix of size  $10^5 \times 10^5$ , the full image  $I_{in}$  is computed after several iterations in 2-3 secs.

Specifically, forward rendering readily admits to a Neumann series solution. We derive a similar series for the inverse solution and show formally that just as each term of the forward Neumann series adds bounces of light transport, each term of the inverse series zeroes out the corresponding bounce (but unlike in the forward case, also affects higher-order bounces). However, convergence of the inverse series is oscillatory. While the forward series convergence condition corresponds to energy conservation, in the inverse case the condition is more complex—a sufficient condition is that the albedo of surfaces is below 0.5, so that the net global illumination is still less than the direct lighting component.

Recent techniques for acquiring the light transport of real scenes [11, 16] have facilitated relighting applications in computer graphics, equivalent to matrix-vector multiplication. While light transport *inversion* enables new applications like illumination estimation, separating bounces of global illumination [18] and projector radiometric compensation [21], the high resolution of real transport data ( $10^5 \times 10^5$  or higher) often makes standard matrix inversion impractical.

Inspired by efficient algorithmic approaches such as finite element radiosity [4] and Monte Carlo methods [9, 20] 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 full radiometric compensation of interreflections while projecting complex scenes (Fig. 1), as well as separation of individual local and global illumination components or bounces (Fig. 2).

To summarize, the main contributions of this paper are:



**Fig. 2.** Separation of bounces of interreflection using our iterative light transport inversion technique, that runs in 3 seconds on a  $131K \times 131K$  light transport matrix.

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

## 2 Previous Work

Our work builds most closely on Seitz et al. [18], who introduce the problem of inverse light transport. This paper elucidates novel theoretical connections to the forward problem and proposes new algorithms that are far more efficient (hence, practical on high resolution data) than the direct matrix inversion of [18]. Nayar et al. [13] present a fast direct and global separation where the entire scene is lit by a light source. In contrast, we acquire the full light transport, but can then separate each bounce of light and consider general illumination conditions.

Our approach is distinct from inverse rendering methods [10, 17] that acquire lighting and reflectance, as well as the inverse global illumination method of [22] for BRDF estimation, all of which assume known scene geometry. In contrast, we observe only the light transport matrix—both geometry and reflectance are unknowns in this work. We focus on the case where scene elements are illuminated individually by a single projector, with a camera recording the output [15, 16, 18]. Extensions to incident (and reflected) light fields [7, 11, 19] are encompassed by the theory, but not yet considered in our practical applications.

One important application is projector radiometric compensation, where we seek to project a desired image, while compensating for global illumination effects on the display surface. Many recent efforts considered non-uniform scene reflectance, but not interreflections [5, 6, 14]. Clusters of camera-projector pixels are formed in [21], with a brute-force transport inversion within clusters, but inter-cluster interactions are ignored. Iterative inverse methods for diffuse scenes are proposed in [3], and a series expansion for inverse light transport, denoted as

$\mathbf{l}_{\text{in}}$	Incident lighting or projected pattern	$\mathbf{l}_{\text{d}}$	Direct light from sources
$\mathbf{l}_{\text{g}}$	Global light from interreflections	$\mathbf{l}_{\text{out}}$	Outgoing light ( $\mathbf{l}_{\text{out}} = \mathbf{l}_{\text{d}} + \mathbf{l}_{\text{g}}$ )
$\mathbf{S}$	Forward transport operator	$\mathbf{S}^{-1}$	Inverse transport operator
$\mathbf{R}$	Interreflections operator ( $\mathbf{R} = \mathbf{S} - \mathbf{I}$ )	$\mathbf{K}$	Local reflection operator
$\mathbf{G}$	Geometric operator	$\mathbf{A}$	Net global transport, $\mathbf{A} = \mathbf{K}\mathbf{G}$
$\mathbf{F}$	First bounce from projector	$\mathbf{T}$	Observed light transport, $\mathbf{T} = \mathbf{S}\mathbf{F}$
$N$	Transport resolution (matrix size $N^2$ )	$p$	$\ \mathbf{K}\ $ , related to max albedo ( $p < 1$ )

**Fig. 3.** Table of the main notation used in the paper.

the stratified inverse, is derived by Ng et al. [15]. We show that this series is a natural analog to the forward Neumann series. Our dual formulation enables us to go much further, by showing that the inverse series subtracts physical bounces of light, analyzing convergence conditions and providing fast algorithms that relate to radiosity, wavelet and Monte Carlo methods in forward rendering.

### 3 Preliminaries

Owing 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:

$$\mathbf{l}_{\text{out}} = \mathbf{S}\mathbf{l}_{\text{d}}, \quad (1)$$

where  $\mathbf{l}_{\text{out}}$  is the outgoing “global” light, and  $\mathbf{l}_{\text{d}}$  is the direct lighting on surfaces due to external sources. In continuous form,  $\mathbf{l}_{\text{out}}$  and  $\mathbf{l}_{\text{d}}$  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,  $\mathbf{l}_{\text{out}}$  and  $\mathbf{l}_{\text{d}}$  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 computations, we do not see the light source directly, but rather its effect on the scene, which we denote as the direct component,  $\mathbf{l}_{\text{d}}$ . The inverse light transport problem considered here is simply

$$\mathbf{l}_{\text{d}} = \mathbf{S}^{-1}\mathbf{l}_{\text{out}}, \quad (2)$$

where we seek to invert the operator  $\mathbf{S}^{-1}$ , 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,

$$\mathbf{l}_{\text{d}} = \mathbf{F}\mathbf{l}_{\text{in}} \quad \mathbf{l}_{\text{out}} = \mathbf{T}\mathbf{l}_{\text{in}} = \mathbf{S}\mathbf{F}\mathbf{l}_{\text{in}}, \quad (3)$$

where  $\mathbf{l}_{\text{in}}$  is the incident pattern projected (or distant light sources turned on), and  $\mathbf{F}$  is a “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  $\mathbf{S}$ ,

	Forward	Inverse
Problem	$\mathbf{l}_{\text{out}} = \mathbf{S}\mathbf{l}_{\text{d}}$	$\mathbf{l}_{\text{d}} = \mathbf{S}^{-1}\mathbf{l}_{\text{out}}$
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	$\mathbf{l}_{\text{out}}^{(k)} = \mathbf{l}_{\text{d}} + \mathbf{A}\mathbf{l}_{\text{out}}^{(k-1)}$	$\mathbf{l}_{\text{d}}^{(k)} = \mathbf{l}_{\text{out}} - \mathbf{R}\mathbf{l}_{\text{d}}^{(k-1)}$
Monte Carlo	$\sum \mathbf{A}_{i_0 i_1} \mathbf{A}_{i_1 i_2} \dots \mathbf{l}_{\text{d}}(i_k)$	$\sum (-1)^k \mathbf{R}_{i_0 i_1} \mathbf{R}_{i_1 i_2} \dots \mathbf{l}_{\text{out}}(i_k)$

**Fig. 4.** Duality of forward and inverse light transport, indicating analogous relations for some key properties. (Monte Carlo equations abbreviated; full forms in text.)

using  $\mathbf{S} = \mathbf{T}\mathbf{F}^{-1}$ . Moreover, applications like radiometric compensation actually seek to recover  $\mathbf{l}_{\text{in}}$  (rather than  $\mathbf{l}_{\text{d}}$  in (2)) given by  $\mathbf{l}_{\text{in}} = \mathbf{T}^{-1}\mathbf{l}_{\text{out}} = \mathbf{F}^{-1}\mathbf{l}_{\text{d}}$ .

Since we focus on global illumination  $\mathbf{S}$ , we will consider setups where  $\mathbf{S}$  is easy to obtain from  $\mathbf{T}$ , i.e., 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 (where  $\mathbf{F}$  is low rank for diffuse surfaces [17]). After geometric calibration, we can use the same parameterization for projection and camera images [18].  $\mathbf{F}$  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.

## 4 Dual Forward and Inverse Light Transport

In this section, we show that the structure of the rendering equation exposes a strong duality between forward and inverse light transport. We derive analogous inverse and forward Neumann series expansions, and interpret them in terms of physical bounces of light. Key theoretical results are summarized in Fig. 4.

In the operator notation of [1], the rendering equation is written as

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

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.<sup>3</sup> It naturally follows that

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

<sup>3</sup> This formulation is valid for any opaque BRDF when considering the full light field. While the theory is fully general, our experiments will consider projection to a single view, which introduces practical limitations, as discussed in Sec. 7.

This well known result shows that the forward problem formally involves a matrix or operator inversion. Also note that if the scene geometry and reflectance (and hence  $\mathbf{A}$ ) are known, we simply have  $\mathbf{S}^{-1} = \mathbf{I} - \mathbf{A}$ , as noted by [18, 12]. We focus here on cases where we only measure  $\mathbf{S}$ , but do not know or compute  $\mathbf{A}$ .

We can separate  $\mathbf{l}_{\text{out}}$  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:

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

The very similar or dual forms of (5) and (7) is a key insight in this paper, and allows direct leveraging of many forward rendering theories and algorithms for inverse rendering and inverse light transport algorithms in computer vision.

**Neumann Forward and Inverse Series:** The forward equations (4) and (5) have well-known series expansions corresponding physically to light bounces,

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

We can also relate the global illumination operator  $\mathbf{R}$  to this expansion,

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

Mathematically, the dual formulation in (7) has a series analogous to (8),

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

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).

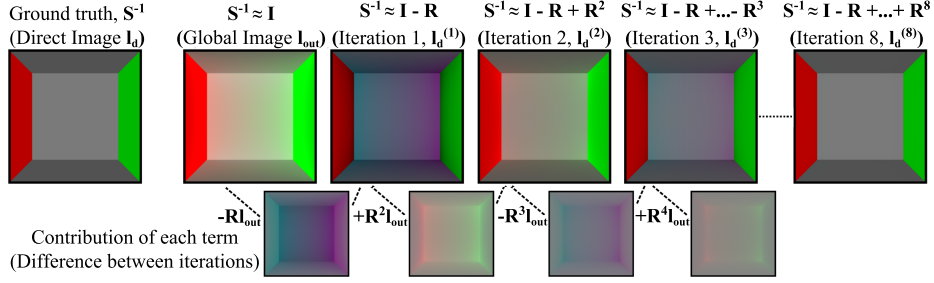
With suitable algebraic manipulations, one may note that (10) explains the stratified inverses of Ng *et al.* [15] and relates it to the rendering equation.<sup>4</sup>

**Interpretation as Physical Bounces of Light:** Consider an approximation up to order  $n$ , that we denote as  $\mathbf{S}_n$  or  $\mathbf{S}_n^{-1}$ . In the forward case, it is clear that

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

where the first  $n$  physical bounces of light are represented (each term adds the next bounce), and the error is from neglecting bounces  $n + 1$  onwards.

<sup>4</sup> 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  $\mathbf{T}^{-1} = \mathbf{F}^{-1}\mathbf{S}^{-1}$  enables one to derive the results in [15].



**Fig. 5. Top:** From left to right, we add more terms of the inverse series, going from the simulated global illumination output  $\mathbf{I}_{\text{out}}$  to the “direct lighting” result  $\mathbf{I}_d$  (shown leftmost). These terms also correspond to the iterations introduced in Sec. 6. **Bottom:** Contributions of individual terms (neutral grey is 0).

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 [2], we derive a surprising result: each term of the inverse series *cancels or zeros out the corresponding bounce of light transport*, analogous to the forward case. Formally, we show that  $\mathbf{S}_0^{-1} = \mathbf{I}$ ,  $\mathbf{S}_1^{-1} = \mathbf{I} - \mathbf{A} - \mathbf{A}^2 - \dots$ , and for  $n > 1$ ,

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

Now, consider the case when  $m \leq n$ . In this case, the second summation has a limit of  $m > 1$ , and the coefficient of  $\mathbf{A}^m$  becomes  $\sum_{k=1}^m (-1)^k \binom{m-1}{k-1}$ , which is the binomial expansion of  $(1+x)^{m-1}$  with  $x = -1$ , thus, identically 0.

This implies a key result, that the  $\mathbf{A}^m$  terms vanish for  $2 \leq m \leq n$ ,

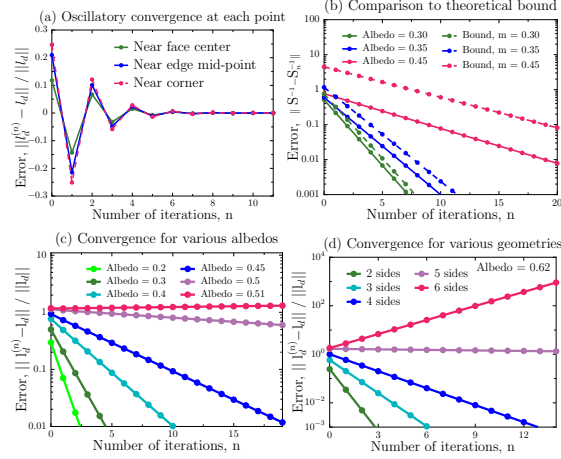
$$\boxed{\mathbf{S}_n^{-1} = \mathbf{I} - \mathbf{A} + O(\mathbf{A}^{n+1}) \quad \mathbf{S}_n^{-1} - \mathbf{S}^{-1} = O(\mathbf{A}^{n+1})} \quad (13)$$

analogous to the forward series in (11). An exact expression can be derived as  $\mathbf{S}_n^{-1} = \mathbf{I} - \mathbf{A} + (-1)^n \sum_{m=n+1}^{\infty} \binom{m-2}{n-1} \mathbf{A}^m$ . Note the oscillatory series behavior from the  $(-1)^n$ . Finally, since  $\mathbf{S} = (\mathbf{I} - \mathbf{A})^{-1}$ ,

$$\mathbf{S}_n^{-1} \mathbf{S} = [(\mathbf{I} - \mathbf{A}) + O(\mathbf{A}^{n+1})] [\mathbf{I} - \mathbf{A}]^{-1} = \mathbf{I} + O(\mathbf{A}^{n+1}). \quad (14)$$

In other words, the  $n$  term series  $\mathbf{S}_n^{-1}$  annihilates the first  $n$  physical bounces of light (each term in the series zeroes the corresponding interreflection bounce), leaving only bounces  $n+1$  and higher. However, 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. An exact result is

$$\mathbf{S}_n^{-1} \mathbf{S} = \mathbf{I} + (-1)^n \sum_{m=n+1}^{\infty} \binom{m-1}{n} \mathbf{A}^m. \quad (15)$$



**Fig. 6.** Error analysis of inverse series. (a): Convergence at different points (center, edge, corner). (b): Comparison of error to theoretical bound for different albedos showing good agreement. (c): Convergence for different albedos (is faster for lower albedos up to 0.5). (d): As expected, an albedo of 0.62 diverges for a closed box (6 sides), shows slow convergence for a 5-sided box and rapid convergence for more open environments.

## 5 Convergence and Error Analysis

For the forward case, Arvo et al. [1] prove several results, briefly summarized here. For a closed enclosure,  $\|\mathbf{G}\| = 1$  (less for open scenes). From energy conservation, excluding perfect reflectors,  $\|\mathbf{K}\| \leq p < 1$ , where  $p$  relates to the surface albedo (for non-diffuse materials, it is the maximum over all incident directions of the fraction of total energy reflected).<sup>5</sup> Since  $\|\mathbf{A}\| \leq \|\mathbf{K}\| \|\mathbf{G}\|$ , it follows that  $\|\mathbf{A}\| \leq p < 1$ , so the forward series always converges.

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

$$\|\mathbf{R}\| \leq \|\mathbf{A}\| + \|\mathbf{A}^2\| + \dots \leq p + p^2 + \dots = \frac{p}{1-p}. \quad (16)$$

If  $p < \frac{1}{2}$ , 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 non-diffuse 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 to bound 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 ( $\mathbf{S}_n$  or  $\mathbf{S}_n^{-1}$ ) for forward and inverse series can be bounded as

$$\|\mathbf{S} - \mathbf{S}_n\| \leq \sum_{k=n+1}^{\infty} \|\mathbf{A}^k\| \leq \sum_{k=n+1}^{\infty} p^k = \frac{p^{n+1}}{1-p}. \quad (17)$$

$$\|\mathbf{S}^{-1} - \mathbf{S}_n^{-1}\| \leq \sum_{k=n+1}^{\infty} \|\mathbf{R}^k\| \leq \sum_{k=n+1}^{\infty} \left(\frac{p}{1-p}\right)^k = \frac{p^{n+1}}{(1-p)^n(1-2p)}. \quad (18)$$

<sup>5</sup> These relations hold in any  $L^p$  norm, since by reciprocity,  $\|\mathbf{K}\|_1 = \|\mathbf{K}\|_{\infty} = p$ , and  $\|\cdot\|_q \leq \max(\|\cdot\|_1, \|\cdot\|_{\infty})$ .





**Fig. 7.** Validation of the theory for shadowed and non-Lambertian scenes. Our iterative method recovers  $\mathbf{l}_d$  in 10 iterations for the shadowed scene and 20 for the glossy one.

**Numerical Simulations:** For simplicity, we consider a synthetic diffuse box (closed, so  $\|\mathbf{G}\|=1$ ), without shadows but with interreflections. Fig. 5 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, while forward global illumination in  $\mathbf{l}_{out}$  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  $\mathbf{l}_d$  has no color bleeding, as desired.

In Fig. 6, we analyze errors and convergence. Fig. 6a indicates similar oscillatory convergence behavior near corners, edges and face centers. Fig. 6b shows excellent agreement, up to a constant factor, between error for the whole  $\mathbf{S}^{-1}$  operator and the theoretical bound in (18). Fig. 6c 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. 6d shows the variation of convergence with geometry (that is,  $\|\mathbf{G}\|$ ). For an albedo of 0.62, close to the theoretical limit for a 5-sided box, we observe very slow convergence for a 5-sided box, divergence for a 6-sided box and rapid convergence for more open geometries.

Finally, Fig. 7 shows a scene with occlusions and glossy surfaces. Similar behaviors hold as above, with convergence of the inverse series to direct lighting.


## 6 Exploiting Duality for Fast Light Transport Inversion

We now introduce efficient algorithms for high-resolution light transport inversion, exploring duals to iterative finite element radiosity, wavelet accelerations and Monte Carlo methods.

**Finite Element Methods:** Forward rendering rarely computes the series in (8) to explicitly determine  $\mathbf{S}$ , due to the high cost of matrix-matrix multiplications on high-resolution scenes. Instead, finite element and radiosity methods [4] try to solve  $\mathbf{l}_{out} = \mathbf{l}_d + \mathbf{A}\mathbf{l}_{out}$ , which corresponds directly to (4), iteratively,

$$\mathbf{l}_{out}^{(k)} = \mathbf{l}_d + \mathbf{A}\mathbf{l}_{out}^{(k-1)}. \quad (19)$$

This iteration is numerically stable, and requires only the matrix-vector multiplication for  $\mathbf{A}\mathbf{l}_{out}$ . The superscript stands for the step  $k$ , and  $\mathbf{l}_{out}^{(0)} = \mathbf{l}_d$ . Note that  $n$



Method	Resolution of transport matrix ( $N$ )				
	80	320	1280	5120	10240
Series	1.0	47.0	2.3e3	9.3e4	7.5e5
Iterative	0.1	0.5	10.3	162.5	679.0
Wavelet	1.2	1.3	1.6	2.3	4.0

**Fig. 8.** Timings for series, iterative finite element, and wavelet accelerated methods (using Daubechies<sub>4</sub> wavelets).  $N$  is the transport resolution (matrix size is  $N^2$ ). We normalize timings so that 1.0 corresponds to  $5.57 \times 10^{-4}$  seconds, with experiments in Matlab on an Intel i7 machine. All methods are run until 1% error.

steps simply compute the effect of the first  $n$  terms of the series in (8). For inverse light transport, one can derive a similar relation, starting from  $\mathbf{l}_d = \mathbf{l}_{\text{out}} - \mathbf{R}\mathbf{l}_d$ , that follows from (6). The iterative solution naturally follows, dual to (19),

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

with  $\mathbf{l}_d^{(0)} = \mathbf{l}_{\text{out}}$ . Again, the first  $n$  steps correspond to the first  $n$  terms in (10). Note the negative sign on  $\mathbf{R}$  that determines the oscillatory nature of the series.

**Matrix Iteration:** In cases where we seek to precompute  $\mathbf{S}^{-1}$ , there is also a corresponding full matrix iteration. The dual forward and inverse relations are

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

with  $\mathbf{S}_0 = \mathbf{S}_0^{-1} = \mathbf{I}$ . These equations provide a numerically stable iteration.

**Wavelet 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 it up. This is analogous to wavelet radiosity and light transport in forward rendering [8].

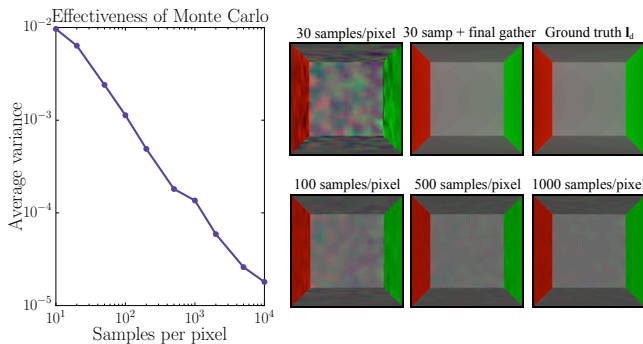
**Monte Carlo Methods:** For the matrix  $\mathbf{A}$ , [9] considers all index permutations

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

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.

Analogously, the inverse series in (10) has a similar form,

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



**Fig. 9. Left:** Variance in Monte Carlo methods. **Right:** Top row shows final gather inversion with only 30 samples. Bottom row shows effects of increasing samples with pure Monte Carlo. Transport resolution is  $N = 5120$ .

where the oscillatory behavior requires the additional  $(-1)^k$  factor. A direct Monte Carlo algorithm uses a number of samples, drawing the indices  $i_1, i_2, \dots, i_k$  at random for each. The expectation of these samples gives the desired result. One may also use fewer samples for the iteration, but compute the final step with a direct matrix-vector multiplication, akin to *final gather* in forward rendering.

**Numerical Simulations:** As timing baseline, we use matrix-matrix multiplications to directly compute the series in (10) (explicit matrix inversion is intractable for high resolutions). In Fig. 8, 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  $O(N^2)$ , with a speedup of three orders of magnitude for large sizes. Wavelet acceleration leads to linear  $O(NW)$  performance, where the number of wavelets  $W$  is relatively insensitive to  $N$ .

Fig. 9 shows the expected inverse relation between variance and number of samples for the Monte Carlo method. The images in the top row show the power of final gather—Monte Carlo with 30 samples is noisy as expected, but is nearly smoothed out using one direct iteration. The bottom row shows that, as expected, pure Monte Carlo converges as the number of samples increases.

## 7 Experiments with Real Data

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 is performed to ensure linearity of the corresponding signals [2]. We assemble 8 images at various exposures into a high dynamic range image.

We present two applications of our iterative light transport inversion — projector radiometric compensation and separating the bounces of light transport. As mentioned in Sec. 3, choosing  $\mathbf{F}$  as the diagonal of  $\mathbf{T}$  is accurate for radiometric compensation, even in non-Lambertian scenes. For our single projector-camera setup, that is only an approximation for applications like bounce separation in specular scenes. However, higher bounces rapidly become diffuse in practice and

our experiments show robust results even for non-diffuse scenes. We refer the reader to [2] for a more complete discussion.

**Projector Radiometric Compensation:** The ubiquitous use of projectors may necessitate inverting photometric distortions and interreflection effects to simulate any desired appearance in non-flat, non-Lambertian spaces. In terms of our theory, given a desired appearance  $\mathbf{l}_{\text{out}}$ , we seek to invert the light transport to find  $\mathbf{l}_{\text{d}} = \mathbf{S}^{-1}\mathbf{l}_{\text{out}}$ . As discussed in Sec. 3, we 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 secs. While such high resolutions may be infeasible for a straightforward matrix inversion, based on the patterns in Fig. 8, the stratified inverses method of [15] will require 1 – 2 orders of magnitude more time. Also, in contrast to the method of [21], our algorithms are physically motivated and not contingent on any tunable parameters.

**Separating Bounces:** One consequence of our theory is that once the light transport has been acquired, we can quickly separate an image into the different bounces (direct, 1st bounce indirect, 2nd bounce indirect and so on). It follows from (19), noting that  $\mathbf{S}^{-1} = \mathbf{I} - \mathbf{A}$ , that the  $k$ -th indirect bounce is


$$\mathbf{l}_{\text{out}}^{(k+1)} - \mathbf{l}_{\text{out}}^{(k)} = \mathbf{l}_{\text{d}} - \mathbf{S}^{-1}\mathbf{l}_{\text{out}}^{(k)}. \quad (24)$$

Thus, each successive run of our iterative inversion algorithm yields a bounce of light transport. Fig. 10 shows a didactic example demonstrating the accuracy of the bounce separation. 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 perfectly between the two walls, as expected. Fig. 11 demonstrates the same with a non-Lambertian occluder present in the scene. We observe that the specular highlight is limited only to the direct component and absent from the indirect bounces, which is also expected.

This application is the same as [18], 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. 11, 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. 2 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. 8, a brute force inversion will require nearly 150 hours. In contrast, we require only 33ms per iteration in our (unoptimized) Matlab implementation, for a total of about 3 sec to separate each bounce. Note that the faster method of [13] yields only the top row of Fig. 10 for a *particular* lighting configuration, while we can separate all the bounces for *any* lighting, albeit at the expense of a more laborious acquisition.



**Fig. 10.** Separation of individual bounces. 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.



**Fig. 11.** 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.

We share some restrictions with other projector-camera systems, such as shutter speeds limited by projector refresh rates, color bleeding and non-linear color mixing ratios. For radiometric compensation, the projector cannot display negative values, which may lead to clipping artifacts in dark regions.

## 8 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 perform radiometric compensation in dynamic scenes.

**Acknowledgments:** This work is funded by ONR YIP grant N00014-10-1-0032, ONR PECASE grant N00014-09-1-0741, a National Science Scholarship from A\*STAR Graduate Academy of Singapore, as well as generous support from Adobe, NVIDIA, Intel and Pixar. We thank Joo Hwee Lim and Zhiyong Huang for kind support at I<sup>2</sup>R and anonymous reviewers for useful comments.

## References

1. J. Arvo, K. Torrance, and B. Smits. A framework for the analysis of error in global illumination algorithms. In *SIGGRAPH 94*, pages 75–84, 1994.
2. Jiamin Bai, Manmohan Chandraker, Tian-Tsong Ng, and Ravi Ramamoorthi. A dual theory of inverse and forward light transport. Technical Report UCB/EECS-2010-101, UC Berkeley, Jun 2010.
3. O. Bimber, A. Grundhoefer, T. Zeidler, D. Danch, and P. Kapakos. Compensating indirect scattering for immersive and semi-immersive projection displays. In *IEEE Virtual Reality*, pages 151–158, 2006.
4. M. Cohen and J. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press, 1993.
5. Y. Ding, J. Xiao, K.-H. Tan, and J. Yu. Catadioptric projectors. In *CVPR*, 2009.
6. K. Fujii, M. Grossberg, and S. Nayar. A Projector-camera System with Real-time Photometric Adaptation for Dynamic Environments. In *CVPR 05*, pages 814–821, 2005.
7. G. Garg, E. Talvala, M. Levoy, and H. Lensch. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *EGSR 06*, pages 251–262, 2006.
8. S. Gortler, P. Schröder, M. Cohen, and P. Hanrahan. Wavelet radiosity. In *SIGGRAPH 93*, pages 221–230, 1993.
9. J. Kajiya. The rendering equation. In *SIGGRAPH 86*, pages 143–150, 1986.
10. S. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell Univ., 1998.
11. V. Masselus, P. Peers, P. Dutre, and Y. Willems. Relighting with 4D incident light fields. *ACM Trans. on Graphics (SIGGRAPH 03)*, 22(3):613–620, 2003.
12. Y. Mukaigawa, T. Kakinuma, and Y. Ohta. Analytical compensation of inter-reflection for pattern projection. In *ACM VRST*, pages 265–268, 2006.
13. 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. on Graphics (SIGGRAPH 06)*, 25(3):935–944, 2006.
14. S. Nayar, H. Peri, M. Grossberg, and P. Belhumeur. A Projection System with Radiometric Compensation for Screen Imperfections. In *IEEE PROCAMS*, 2003.
15. T.-T. Ng, R. S. Pahwa, J. Bai, Q.-S. Quek, and K.-H. Tan. Radiometric Compensation Using Stratified Inverses. In *ICCV*, 2009.
16. 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. on Graphics (SIGGRAPH 06)*, 25(3):746–753, 2006.
17. R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH 01*, pages 117–128, 2001.
18. S. Seitz, Y. Matsushita, and K. Kutulakos. A theory of inverse light transport. In *ICCV 05*, pages 1440–1447, 2005.
19. P. Sen, B. Chen, G. Garg, S. Marschner, M. Horowitz, M. Levoy, and H. Lensch. Dual Photography. *ACM Trans. on Graphics (SIGGRAPH)*, 24(3):745–755, 2005.
20. E. Veach. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1998.
21. G. Wetzstein and O. Bimber. Radiometric Compensation through Inverse Light Transport. In *Pacific Conf. on Comp. Graphics and Appl.*, pages 391–399, 2007.
22. Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *SIGGRAPH 99*, pages 215–224, 1999.