

Compressive Inverse Light Transport

Xinqi Chu¹

<http://www.ifp.uiuc.edu/~chu36/>

Tian-Tsong Ng²

<http://www1.i2r.a-star.edu.sg/~ttng/>

Ramanpreet Pahwa¹

pahwa1@uiuc.edu

Tony Q.S. Quek²

<http://www1.i2r.a-star.edu.sg/~qsquek/>

Thomas S. Huang¹

huang@ifp.uiuc.edu

¹ Department of ECE

University of Illinois at

Urbana-Champaign

Urbana, IL, US

² Institute for Infocomm Research

A*STAR, Singapore

Abstract

This paper explores the possibility of acquiring inverse light transport directly. The current strategy of obtaining an inverse light transport matrix involves two steps: First, acquire the forward light transport matrix (f -LTM) and then calculate the inverse of the f -LTM. Both steps of the strategy requires considerable computational power. In addition to computational cost, the measurement error incurred at the first step inevitably propagates to or potentially gets amplified in the matrix inversion step.

In this paper, we propose a sensing strategy that acquires the inverse light transport matrix (i -LTM) directly, without reconstructing the f -LTM. Our direct strategy reduces both computational error and cost of acquiring i -LTM. For that, we propose a *compressive inverse theory*. Following the compressible property of i -LTM, a reconstruction condition for i -LTM is introduced. This new framework implies a trade-off between two factors: condition numbers of submatrices of f -LTM and the isometry constant of the illumination pattern. Our direct i -LTM reconstruction method is then demonstrated with a 2nd-bounce separation experiment on an M-shaped panel scene. Finally by quantitatively comparing our method with the existing two-stage approach, our method shows higher accuracy with lower complexity. The proofs of main theorem/lemma are contained in the supplementary material. The compressive inverse theory is general and potentially useful for wider application.

1 Introduction

A forward light transport simulates global illumination in a scene given direct lighting or corresponding light source emission. It embodies the forward rendering process, a cornerstone of computer graphics, which aggregates the effect of light bouncing in a scene. An inverse light transport reverses the forward process; it enables undoing of interreflections and separation of light bounces in a real scene [16, 20]. In all current work, an inverse light transport matrix (i -LTM) is obtained by inverting a forward light transport matrix (f -LTM). For a projector-camera setup [10, 16, 27], a f -LTM can easily exceed the size of $10^5 \times 10^5$.

Acquiring such a large f -LTM can take hours or days while inverting the f -LTM requires various forms of approximation which compromises on the accuracy of the inverse light transport [10, 16, 27].

In this work, we propose a way of computing i -LTM directly from the measurements without the need for a prior and explicit reconstruction of the f -LTM. We show that within the framework of *compressive inverse theory*, i -LTM can be obtained in a similar manner as the f -LTM without additional computational cost and aggregated error from matrix inversion. For simplicity, we assume a square matrix for light transport and defer the case of non-square matrices to future work.

2 Previous Work

As this work concerns with reconstructing i -LTM directly in the compressive sensing framework, the methodology is related to the prior works in inverse light transport, forward light transport acquisition, compressive sensing, and matrix inversion. We denote the f -LTM as \mathbf{T} henceforth.

2.1 Inverse light transport computation

Inverse global illumination was introduced by Yu *et al.* [28], but with the focus on estimating reflectance properties, rather than compensating input lighting patterns. Seitz *et al.* [20] proposed a theory of inverse light transport and demonstrated light bounce separation with inverse light transport.

Inverse light transport can be computed by inverting the forward light transport matrix [10, 12, 16, 27]. In all these works, the f -LTM is first reconstructed from the measurements before deriving the i -LTM from the f -LTM. In the conventional setting, the computation of i -LTM mainly involves solving the inverse problem

$$\mathbf{l}_{\text{out}} = \mathbf{T}\mathbf{l}_{\text{in}}, \quad (1)$$

where \mathbf{l}_{out} is the observed scene under global illumination and \mathbf{l}_{in} is the input light pattern. There are two approaches for solving the inverse light transport problem in Eq. 1. The first approach solves Eq. 1 as a system of linear equations [1]. Efficient methods of this approach are the Jacobi method and the Gauss-Seidel method that involves iterative vector-matrix multiplication. The second approach is to invert the matrix \mathbf{T} . Matrix inversion ideally corresponds to solving the matrix-matrix equation $\mathbf{T}\mathbf{T}^{-1} = \mathbf{I}$, hence the second approach has much higher computational complexity than the first approach. Once \mathbf{T}^{-1} is precomputed, given a new \mathbf{l}_{in} , \mathbf{l}_{out} can be easily obtained with a single-step vector-matrix multiplication. However, inverting a large-size f -LTM is computational and memory intensive, hence various forms of approximation were introduced [10, 12, 16, 27].

In contrast to this conventional procedure, our method enables inverse light transport to be computed directly from the measurements, without the need for explicitly reconstructing the f -LTM. By skipping the intermediate steps of f -LTM reconstruction and by avoiding the approximate light transport inversion, it is logical that our method will be more computationally efficient and accurate.

A recent work by O’Toole *et al.* [17] computes the output image of an inverse light transport in the optical domain, without explicitly computing the f -LTM or i -LTM. The work is in the same spirit as our work in reducing computational steps. However, pre-computing i -LTM is important when it comes to processing a video or a sequence of images. Without

an explicit i -LTM, the implicit matrix inversion process needs to be repeated for each frame of a video and hence incurs a high computational cost.

2.2 Forward light transport acquisition

A large body of work over the last decade in computer graphics and computer vision has dealt with acquiring f -LTM, indicating how a real scene responds to diffuse or focused light from projectors [9, 10, 18, 22, 26]. While we focus here on real scenes, the use of precomputed transport is increasingly common even for synthetic scenes, in applications like real-time relighting [15].

The brute-force acquisition method turns on the projector pixels one by one while the response of each projector pixel is recorded by a camera to form columns in a f -LTM. Alternatively, Hadamard light patterns can be used to achieve a better signal-to-noise ratio [19]. Sen et al. [22] proposed a multi-resolution and adaptive method to measure f -LTM. In [10], a deterministic stripe-scanning method was proposed to acquire \mathbf{T} where horizontal and vertical stripes scan through the scene in a sequential manner. For low-frequency diffuse light sources, a fast method in Wang et al. [26] was introduced to take advantage of the coherency in the rows and columns of f -LTM. Compressive sensing, which is closely related to our work, was also used to measure a f -LTM by exploiting its sparse or compressible property [18, 21].

2.3 Compressed sensing

Compressed sensing has opened a new frontier in sparse signal sampling and measurement. Candés *et al.* and Donoho [2, 4, 11] showed that sparse signal can be reconstructed with much fewer measurements than what is dictated by the Nyquist-Shannon sampling theory [13]. The canonical form of compressive sensing problem is given by

$$(P1) \quad \min \|\mathbf{x}\|_0 \text{ s.t. } \mathbf{y} = \mathbf{A}\mathbf{x}, \quad (2)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is an underdetermined measurement matrix with $n > m$, $\mathbf{x} \in \mathbb{R}^n$ is an unknown vector with sparse elements, and $\mathbf{y} \in \mathbb{R}^m$ is the measurement vector. The ℓ_0 norm in Eq. 2 measures the number of nonzero entries in \mathbf{x} . The optimization problem (P1) in Eq. 2 is unfortunately NP-complete. However, when the measurement matrix \mathbf{A} satisfies a condition known as the restricted isometry property (RIP) [4], the solution for problem (P1) can be obtained by replacing the ℓ_0 norm with the ℓ_1 norm,

$$(P2) \quad \min \|\mathbf{x}\|_1 \text{ s.t. } \mathbf{y} = \mathbf{A}\mathbf{x}, \quad (3)$$

which can be solved with a linear program in polynomial time [2, 11].

With compressed sensing, Peers *et al.* and Sen *et al.* [18, 21] proposed to exploit sparsity in both rows and columns of the f -LTM and compute the response of each pixel by projecting patterned illumination. They probe the light transport matrix by m illumination conditions $\mathbf{L} = [\mathbf{l}_0, \dots, \mathbf{l}_m]$ to obtain their corresponding observations $\mathbf{C} = [\mathbf{c}_0, \dots, \mathbf{c}_m]$, which is,

$$\mathbf{C} = \mathbf{T}\mathbf{L} \Leftrightarrow \mathbf{C}^T = \mathbf{L}^T\mathbf{T}^T \Leftrightarrow \mathbf{c}'_i = \mathbf{L}^T\mathbf{t}'_i, \quad (4)$$

where \mathbf{t}'_i is a column in \mathbf{T}^T that represents the reflectance function of the i -th pixel in the camera image and \mathbf{c}'_i is a column of \mathbf{C}^T . This formulation maps directly to a compressive sensing context where \mathbf{L}^T fulfills the role of the measurement ensemble \mathbf{A} , and the reflectance function \mathbf{t}'_i corresponds to the discrete signal \mathbf{x} . To exploit sparsity in both rows and columns of

	compressive sensing	compressive inverse
Encoding matrix	$\mathbf{L} = [\mathbf{l}_1, \dots, \mathbf{l}_m]$	$\mathbf{C}^T = \begin{bmatrix} \mathbf{c}_1^T \\ \vdots \\ \mathbf{c}_m^T \end{bmatrix}$
Codes	$\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_m]$	$\mathbf{L}^T = [\mathbf{L}(1), \dots, \mathbf{L}(N)]$
Relation	$\mathbf{C} = \mathbf{T}\mathbf{L}$	$\mathbf{C}^T (\mathbf{T}^{-1})^T = \mathbf{L}^T$
Reconstruction	$\min \ \hat{\mathbf{t}}_i\ _0$ s.t. $\mathbf{c}_i = \hat{\mathbf{t}}_i^T \mathbf{L}\mathbf{W}$	$\min \ \mathbf{h}_i\ _0$ s.t. $\mathbf{C}^T \mathbf{W}\mathbf{h}_i = \mathbf{L}_w(i)$

Table 1: A comparison with conventional compressed sensing framework.

\mathbf{T} , \mathbf{T} can be further represented as $\mathbf{T} = \mathbf{W}_1^T \hat{\mathbf{T}} \mathbf{W}_2$. The right transformation \mathbf{W}_2 operates on the rows (i.e., reflectance functions) of the transport matrix to exploit the coherency within the reflectance functions, while the left transformation \mathbf{W}_1^T operates on the columns of the transport matrix (i.e., photographs) to condense the inter-pixel information. The resulting doubly transformed f -LTM is even sparser, and thus is potentially inferable from fewer measurements at a higher accuracy.

3 Compressive Inverse

Noticing the computational cost and inaccuracy of the two-step strategy of i -LTM acquisition, in this section we develop the *compressive inverse* theory which enables a one-step acquisition of i -LTM. Eq. 4, relating illumination and observation by \mathbf{T} , can be rewritten as:

$$\mathbf{C}^T (\mathbf{T}^{-1})^T = \mathbf{L}^T \quad (5)$$

by left multiplying \mathbf{T}^{-1} and take transpose on both sides.

This equation is in the form of compressed sensing by considering columns of \mathbf{C} as a "sensing ensemble" and columns of \mathbf{L} as "measurements/observations" in terms of compressed sensing. Notice that the terminology is reversed if we consider the problem in a physical context where \mathbf{C} are the observations from the sensing patterns \mathbf{L} . In other words, mathematically in Eq. 6, we are probing the matrix \mathbf{T}^{-1} with \mathbf{C} where physically we are probing \mathbf{T} with \mathbf{L} . To avoid this confusion, we use a coding theory terminology in the seminal work [4] to dissociate physical meanings with the conventional terminology such as "sensing ensemble" or "measurements". We refer the matrix on the left of our target matrix as "encoding matrix" and the vectors comprising the matrix on the right hand side as "codes". A comparison of formulations of forward and inverse compressive sensing is in Table 1. As we shall prove this in section 3.1, $(\mathbf{T}^{-1})^T$ is compressible in both rows and columns with orthonormal basis \mathbf{W} (Eq. 9), substituting Eq. 9 into Eq. 6 and right multiply \mathbf{W} ,

$$\mathbf{C}^T \mathbf{W} (\widehat{\mathbf{T}^{-1}})^T = \mathbf{L}^T \mathbf{W}. \quad (6)$$

The optimization problem for reconstructing the i -LTM $(\mathbf{T}^{-1})^T$ then becomes:

$$(P3) \quad \min \|\mathbf{h}_i\|_0 \text{ s.t. } \mathbf{C}^T \mathbf{W}\mathbf{h}_i = \mathbf{L}_w(i), \quad i = 1, \dots, N \quad (7)$$

where \mathbf{h}_i denotes the i -th column of $(\widehat{\mathbf{T}^{-1}})^T$, i.e. $(\widehat{\mathbf{T}^{-1}})^T = [\mathbf{h}_0, \dots, \mathbf{h}_N]$, and $\mathbf{L}_w(i)$ denotes the i -th column of $\mathbf{L}^T \mathbf{W}$. By Eq. 9 in section 3.1, successfully solving (P3) implies the reconstruction of $(\mathbf{T}^{-1})^T$. Unlike the forward case where \mathbf{L} is in our direct control, i.e. \mathbf{L} can be designed as a Bernoulli ensemble which satisfies RIP, in the *compressive inverse* scenario, the encoding matrix \mathbf{C}^T is not something we can design. Nevertheless, the prior

knowledge that \mathbf{C} is obtained physically by the product of \mathbf{T} and \mathbf{L} is known to us though \mathbf{T} is unknown. Now the question becomes, can we still reconstruct \mathbf{h}_i ? If yes, what is the condition that guarantees the reconstruction? Fortunately, as shown in section 3.2, as long as any of the K -subspaces spanned by \mathbf{T} is not contained in the null space of \mathbf{L}^T , we can still reconstruct the inverse matrix $(\mathbf{T}^{-1})^T$ (Lemma 3.1), furthermore, if the K -submatrices of \mathbf{T} are well-conditioned, efficient algorithms exist for the reconstruction (Theorem 3.1). It is worthy to note that, the recent paper dealing with coherent dictionaries [5] cannot be used for our problem because the optimization yields the signal $(\mathbf{T}^T \mathbf{h}_i = \mathbf{e}_i$ in our case) directly. This is inapplicable to our case where the sparse coefficients are the sought quantities. Before we come to the main theorem for the i -LTM reconstruction, we investigate the following two conditions that must be satisfied before hand: \mathbf{T} must be *invertible* and can be *sparsely represented*.

3.1 Invertibility of f -LTM and compressibility of i -LTM

It is common that f -LTM \mathbf{T} under focused light sources such as light pixels of a projector are diagonally dominant¹ [1, 8, 23]. An example of such f -LTM is shown in Fig.1(a) with another one shown in Section 3 of the supplementary. As long as the dimension of the illumination pattern and observation are of the same dimension, \mathbf{T} is square and invertible due to applying Gershgorin's circle theorem to diagonally dominant matrices. Pseudo-inverse can be used for the non-square case but we focus on the square case in this work.

3.1.1 i -LTM has a sparse representation

Remark 3.1. *If a matrix has sparse representation in both rows and columns with some orthonormal basis, its inverse can be sparsified in rows and columns with the same basis.*

This claim can be argued in two steps. By exploiting sparsity in both rows and columns of \mathbf{T} , we have Eq. 8

$$\mathbf{T} = \mathbf{W}\widehat{\mathbf{T}}\mathbf{W}^T \quad (8)$$

where \mathbf{W} denotes an orthonormal basis and $\widehat{\mathbf{T}}$ is the sparsified matrix. Similarly, we denote $\widehat{\mathbf{T}}^{-1}$ as the matrix transformed by \mathbf{W} as in Eq. 9.

$$\mathbf{T}^{-1} = \mathbf{W}\widehat{\mathbf{T}}^{-1}\mathbf{W}^T. \quad (9)$$

A simple fact derived in [6], as stated in Eq. 10, is that we can obtain $\widehat{\mathbf{T}}^{-1}$ by taking the inverse of the sparse representation of \mathbf{T} , this is due to orthonormality of \mathbf{W} .

$$\widehat{\mathbf{T}}^{-1} = \widehat{\mathbf{T}}^{-1} \quad (10)$$

Now we argue that $\widehat{\mathbf{T}}^{-1}$ is sparse due to the following. Observe that $\widehat{\mathbf{T}}$ is also diagonally dominant by comparing the diagonal and off-diagonal elements. If the diagonal elements of $\widehat{\mathbf{T}}$ are nonzero (which is in general true for f -LTM of focused light sources), we define $\mathbf{E} := \mathcal{D}(\widehat{\mathbf{T}}) - \widehat{\mathbf{T}}$ and $\mathbf{F} := (\mathcal{D}(\widehat{\mathbf{T}}))^{-1}\mathbf{E}$, where $\mathcal{D}(\cdot)$ is a 2D operator that only retains the diagonal elements of a matrix, the following equality holds

$$\widehat{\mathbf{T}} = \mathcal{D}(\widehat{\mathbf{T}})(\mathbf{I} - \mathbf{F}) \quad (11)$$

¹Light transport matrix from low-frequency diffuse light sources may not be diagonally dominant [26]. However, such light transport matrix is in general highly compressible and low-rank. We will address the inverse of such type of light transport in a framework similar to pseudo-inverse in future work.




Figure 1: This figure shows an example of (a) f -LTM and its i -LTM (c) with their corresponding sparsefied matrices (b) and (d) by Haar wavelet basis.

from which we obtain the Neumann expansion

$$\hat{\mathbf{T}}^{-1} \approx \underbrace{(\mathbf{I} + \mathbf{F} + \mathbf{F}^2 + \dots + \mathbf{F}^p)}_{(\hat{\mathbf{T}}^{-1})^{(p)}} (\mathcal{D}(\hat{\mathbf{T}}))^{-1} \quad (12)$$

If $\hat{\mathbf{T}}$ is diagonally dominant and sparse, *i.e.*, the number of nonzeros of each row of $\hat{\mathbf{T}}$ is bounded by a small constant, \mathbf{F}^p is then sparse and the entries shall decrease in value as p increases, as stated in [24] and elaborated in [7, 14]. Hence, $(\hat{\mathbf{T}}^{-1})^{(p)}$ is also compressible. Convergence of the Neumann series for forward and inverse light transport can be found in [1, 16]. Now we concluded that $\hat{\mathbf{T}}^{-1}$ is sparse which means that $\widehat{\mathbf{T}}^{-1}$ must be sparse, and hence \mathbf{T}^{-1} can be sparsely represented but not as sparse as the sparse representation of \mathbf{T} . This conclusion is demonstrated in Fig.1.

3.2 Reconstruction condition for i -LTM

In this section we develop the conditions for reconstructing $\hat{\mathbf{T}}^{-1}$. Following the same developmental line in compressed sensing literature, first we establish a necessary and sufficient condition for the existence of a unique solution which can be obtained by NP-hard programs (with ℓ_0), followed by a theorem stating a condition under which an efficient algorithm exists for reconstruction (with ℓ_1). Same as in Table 1, we denote $\mathbf{L}_w(i)$ to be the i th column of the $m \times N$ matrix $\mathbf{L}^T \mathbf{W}$, and \mathbf{C} to be the encoding matrix. We also use \mathbf{h}_i to denote the i -th column of $(\widehat{\mathbf{T}}^{-1})^T$. Let \mathbb{S}_k be the set of all k -sparse vectors.

3.2.1 When does a solution exist?

Consider the problem (P3) in Eq. 7

$$(P3) \quad \min \|\mathbf{h}_i\|_0 \text{ s.t. } \mathbf{C}^T \mathbf{W} \mathbf{h}_i = \mathbf{L}_w(i), \quad i = 1, \dots, N$$

where \mathbf{C}^T is obtained physically by $\mathbf{C}^T = \mathbf{L}^T \mathbf{T}^T$ as in Eq. 4. For orthonormal matrix \mathbf{W} , the following lemma shows a necessary and sufficient condition for this problem to have a unique solution.

Lemma 3.1. *Given a set $J \subseteq \{1, 2, \dots, N\}$, define the matrix \mathbf{T}_J as the one formed from \mathbf{T} by using columns from the set J . Let $\mathbf{h}_i \in \mathbb{S}_k$ be the i th column of $(\widehat{\mathbf{T}}^{-1})^T$. There exist a unique solution $\hat{\mathbf{h}}_i$ to the minimization procedure P(3) such that $\hat{\mathbf{h}}_i = \mathbf{h}_i$, if and only if*

$$\ker(\mathbf{L}^T) \cap \mathbf{T}_J = \{0\}, \{\forall J | |J| \leq 2k\}$$

We can interpret this lemma as follows, for the minimizer $\hat{\mathbf{h}}_i$ of (P3) to be the i -th column of $(\widehat{\mathbf{T}}^{-1})^T$, a necessary and sufficient condition is that, any of the K -subspaces spanned by \mathbf{T}




Figure 2: This figure shows the trade-off between condition numbers of submatrices of \mathbf{T} and the isometry constant of the illumination pattern \mathbf{L}^T . The shaded region shows all the feasible combinations of $\kappa_{2k}^2(\mathbf{T})$ and $\gamma_{\mathbf{L}^T}(2k)$.

is not contained in the null space of \mathbf{L}^T . Intuitively, for the transformed sensing process to be reversible for reconstruction, the vectors in the K -subspaces of that transformation matrix must be one-to-one under the random mapping.

3.2.2 When does ℓ_1 reconstruction work?

In this section we develop a sufficient condition under which problem (P4) in Eq. 13 guarantees to have the same solution as problem (P3) thereby enabling efficient algorithms.

$$(P4) \quad \min \|\mathbf{h}_i\|_1 \text{ s.t. } \mathbf{C}^T \mathbf{W} \mathbf{h}_i = \mathbf{L}_w(i), \quad i = 1, \dots, N. \quad (13)$$

Our theorem below is based on a careful examination of submatrices consisting of an arbitrary collection of k columns. The symbols $\sigma_{\max}(\mathbf{T})$, $\sigma_{\min}(\mathbf{T})$, respectively denote the usual maximum, minimum nonzero singular values of a matrix \mathbf{T} . We use the superscript (k) to represent extremal values of the above spectral measures for its submatrices. For instance, $\sigma_{\max}^{(k)}(\mathbf{T}) = \sup\{\sigma_{\max}(\mathbf{T}_J), \forall |J| = k\}$, i.e., the largest singular value taken over all k -column submatrices of \mathbf{T} ; while $\sigma_{\min}^{(k)}(\mathbf{T}) = \inf\{\sigma_{\min}(\mathbf{T}_J), \forall |J| = k\}$. Before the statement of the theorem, we define the \mathbf{D} -adapted restricted isometry constant (\mathbf{D} -RIP) [5], which is an extension of a ubiquitous property (RIP) assumed in the analysis of compressed sensing. The dictionary is denoted as \mathbf{D} to comply with the notation in [5], but it represents \mathbf{T}^T in our case.

Definition 3.1. [5] For each integer $k = 1, 2, \dots$, define the \mathbf{D} -adapted isometry constant δ_k of a matrix \mathbf{L}^T as the smallest number such that

$$(1 - \delta_k) \|\mathbf{D} \mathbf{h}\|_2^2 \leq \|\mathbf{L}^T \mathbf{D} \mathbf{h}\|_2^2 \leq (1 + \delta_k) \|\mathbf{D} \mathbf{h}\|_2^2, \forall \mathbf{h} \in \mathbb{S}_k \quad (14)$$

with \mathbf{D} -RIP ratio

$$\gamma_{\mathbf{L}^T}(k) = \frac{1 + \delta_k}{1 - \delta_k} \quad (15)$$

Now we are ready to state the condition under which ℓ_1 reconstruction of i -LTM is valid,

Theorem 3.1. Let \mathbf{L}^T be any matrix of size $m \times N$ with \mathbf{D} -RIP ratio $\gamma_{\mathbf{L}^T}(2k)$. Let $\kappa_k(\mathbf{T}) = \sigma_{\max}^{(k)}(\mathbf{T}) / \sigma_{\min}^{(k)}(\mathbf{T})$, if

$$\kappa_{2k}^2(\mathbf{T}) \cdot \gamma_{\mathbf{L}^T}(2k) < \sqrt{2} + 1 \quad (16)$$

Then, ℓ_1 -minimization will exactly recover $\mathbf{h} \in \mathbb{S}_k$.

This theorem states that, provided that the submatrices of \mathbf{T} is reasonably well-conditioned, we can still ensure an exact recovery of a sparse \mathbf{h} . According to Theorem 3.1, we plot the feasible region of all possible $\kappa_{2k}^2(\mathbf{T})$ and $\gamma_{\mathbf{L}^T}(2k)$ in Fig. 2, which shows that, when \mathbf{L}^T has a smaller isometry constant δ_{2k} , *i.e.*, \mathbf{L}^T contains more random patterns, then there is larger room for $\kappa_{2k}^2(\mathbf{T})$ to vary while still ensures reconstruction. On the contrary, if the submatrices of \mathbf{T} are well-conditioned, *i.e.*, $\kappa_{2k}^2(\mathbf{T}) \rightarrow 1$, we can afford an illumination pattern \mathbf{L}^T with less rows (fewer measurements). Notice that, If \mathbf{T} is canonical or orthonormal, *i.e.* $\kappa_{2k}(\mathbf{T}) = 1$, straight away we have the condition $\delta_{2k} < \sqrt{2} - 1$ which is the state-of-the-art RIP constant derived in [3]. This also verifies that our bound in Eq. 16 is tight.

4 Experiments

In this section, we will validate our method of *i*-LTM estimation in terms of its computational efficiency and accuracy with real experiments on an M-shaped panel scene. We will compare our method with the conventional approach that consists of two separate steps: forward light transport acquisition followed by inversion. All the results in this section have a lighting resolution of 13×21 . Bernoulli patterns are used for the measurement matrix \mathbf{L} in our experiment. Haar wavelet is used as sparsifying basis \mathbf{W} for both rows and columns as in Eq. 9. ℓ_1 -Magic is the toolbox for solving the convex programming problem (P4).

2nd-bounce separation One consequence of the inverse light transport theory is that once the light transport has been acquired, we can quickly separate an image into the different bounces (direct, 1st indirect bounce, 2nd indirect bounce and so on). It follows from [1], that the k -th indirect bounce is

$$\mathbf{I}_{out}^{(k+1)} - \mathbf{I}_{out}^{(k)} = \mathbf{I}_d - \mathbf{F}\mathbf{T}^{-1}\mathbf{I}_{out}^{(k)}, \quad (17)$$

where \mathbf{F} is the first-bounce light transport which can be exactly computed for Lambertian scenes [20], as we do here, or approximated otherwise [16]. Thus, each successive run of our iterative inversion algorithm yields a bounce of light transport. Fig. 3 shows a didactic example demonstrating the second bounce light when lighting up different panels of the M-shaped scene, which consists of two concave V-wedges that induces light interreflection between the mutually facing panels. When panel 1 of the M-scene is lit, panel 2 receives significant second-bounce light bounced off from panel 1. As panel 1 is flat where any two points on the panel see each other, the second-bounce light does not light up panel 1 itself. Similar phenomena are observed when panel 2, 3 and 4 are lit. In this experiment, the *i*-LTM used in this experiment is acquired by 170 measurements.

Comparison on accuracy We compare the accuracy of our direct reconstruction method with the two-phase approach by calculating the relative error, which is the difference between the second-bounce light image (when all panels are lit) computed using the different *i*-LTM with respect to that from the reference *i*-LTM. The reference *i*-LTM we use is the exact inverse of the *f*-LTM acquired by the brute-force method that measures the response of each projector source element one by one. This comparison is shown in Fig. 4 for two cases with 120 and 150 random measurements respectively. On the left of the first row, with the *i*-LTM acquired by our direct method, a second bounce image of the M-shaped panel scene is shown under flood illumination. In the middle the image second bounce separated by the reference *i*-LTM is shown. On the right, a false color image map of the relative error between the two images is shown. The second row shows the corresponding results obtained through a two-stage approach: the *f*-LTM acquired with the compressive sensing method [18, 21] followed by an exact inversion. It is evident that our direct method has higher accuracy.

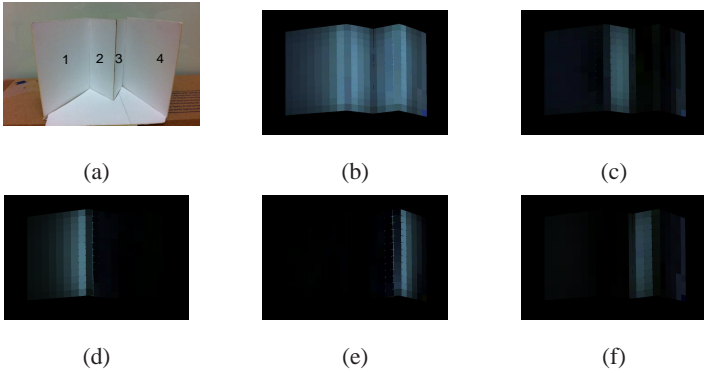


Figure 3: Second bounce images under floodlit at different panels on an M-shaped panel scene. Top row: (a)The M-shaped panel scene. (b)Second-bounce image under floodlit. (c) (d) (e) (f) are 2nd-bounce images under floodlit at panel 1,2,3,4 respectively.

	Comp. 2nd bounce	Ref. 2nd bounce	Error map
Comparison upon 120 measurements			
Direct:			
Two-stage:			 error:0.010
Comparison upon 150 measurements			
Direct:			 error:0.004
Two-stage:			 error:0.015

Figure 4: Quantitative comparison of the 2nd-bounces separated from the *i*-LTM acquired by our direct method and the two-phase method. The average error is computed by taking the sum of squared differences of the reconstruction and the reference images, divided by the number of pixels. The error at the corner is due to inadequate wavelet sparsification at the ending column of the matrix, see Fig.1(d).

Comparison on complexity The computational complexity for ℓ_1 minimization over all columns of i -LTM is $O(N^2m)$ on average. For the two-stage method, the lower bound for the computational complexity of the inversion of a matrix $N \times N$ is $\Omega(N^2 \log N)$ according to [25]. This is the overhead in addition to ℓ_1 minimization. Therefore two-phase method unnecessarily slows down the computation by a significant amount.

5 Conclusion and future work

In this paper, we propose a one-step sensing strategy to acquire the inverse light transport matrix (i -LTM) directly by developing the *compressive inverse* theory. This strategy reduces both error and computational cost of acquiring i -LTM. We first show that i -LTM can be represented sparsely followed by an inverse reconstruction condition for direct acquisition of i -LTM. The new framework implies a trade-off between two factors: condition numbers of submatrices of f -LTM and the isometry constant of the illumination pattern. i -LTM acquired by our method is then demonstrated with 2nd-bounce separation experiments on the M-shaped panel scene. Finally our one-step i -LTM reconstruction method outperforms the two-stage method with higher accuracy and lower complexity. The *compressive inverse* framework can also be extended to other inverse system identification problems. Furthermore, this work can be extended to a pseudo-inverse formulation to deal with the dimension mismatch of illuminations and observations.

References

- [1] J. Bai, M. Chandraker, T.-T. Ng, and R. Ramamoorthi. A dual theory of inverse and forward light transport. In *Proc. of European Conf. on Computer Vision*, 2010.
- [2] E. Candès. Compressive sampling. *International Congress of Mathematicians*, III:1433–1452, 2006.
- [3] E. Candès. The restricted isometry property and its implications for compressed sensing. *Compte Rendus de l'Academie des Sciences, Paris, Serie I*, 346:589–592, 2008.
- [4] E. Candès and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.
- [5] E. Candès, Y. Eldar, D. Needell, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 2010.
- [6] T. F. Chan, W. P. Tang, and W. L. Wan. Wavelet sparse approximate inverse preconditioners. *BIT Numerical Mathematics*, 37(3):644–660, 1997.
- [7] E. Chow. A priori sparsity patterns for parallel sparse approximate inverse preconditioners. *SIAM J. Sci. Comput.*, 21:1804–1822, 2000.
- [8] M. F. Cohen and J. R. Wallace. *Radiosity and Realistic Image Synthesis*. Morgan Kaufmann Publishers, 1993.
- [9] P. Debevec, T. Hawkins, C. Tchou, H. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Proc. of ACM SIGGRAPH*, pages 145–156, 2000.
- [10] Y. Ding, J. Xiao, K.-H. Tan, and J. Yu. Catadioptric projectors. In *Proc. of Computer Vision and Pattern Recognition*, 2009.

- [11] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, April 2006.
- [12] H. Habe, N. Saeki, and T. Matsuyama. Inter-reflection compensation for immersive projection display. In *Proc. of Computer Vision and Pattern Recognition*, 2007.
- [13] S. Haykin. *Communication systems*. Wiley, 2008.
- [14] T. Huckle. Approximate sparsity patterns for the inverse of a matrix and preconditioning. *Appl. Numer. Math.*, 30:291–303, 1999.
- [15] R. Ng, R. Ramamoorthi, and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics*, 22(3):376–381, 2003.
- [16] T.-T. Ng, R. S. Pahwa, J. Bai, Q.-S. Quek, and K.-H. Tan. Radiometric Compensation Using Stratified Inverses. In *Proc. of Int’l Conf. on Computer Vision*, 2009.
- [17] M. O’Toole and K. Kutulakos. Optical computing for fast light transport analysis. In *Proc. of ACM SIGGRAPH ASIA*, 2010.
- [18] P. Peers, D. Mahajan, B. Lamond, A. Ghosh, W. Matusik, R. Ramamoorthi, and P. Debevec. Compressive light transport sensing. *ACM Transactions on Graphics*, 28(1), 2009.
- [19] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. Multiplexing for optimal lighting. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, pages 1339–1354, 2007.
- [20] S. Seitz, Y. Matsushita, and K. Kutulakos. A theory of inverse light transport. In *Proc. of Int’l Conf. on Computer Vision*, pages 1440–1447, 2005.
- [21] P. Sen and S. Darabi. Compressive Dual Photography. *EUROGRAPHICS 09*, 28(2):609–618, 2009.
- [22] P. Sen, B. Chen, G. Garg, S. Marschner, M. Horowitz, M. Levoy, and H. Lensch. Dual Photography. *ACM Transactions on Graphics*, 24(3):745–755, 2005.
- [23] Q. O. Snell and J. L. Gustafson. Parallel hierarchical global illumination. In *IEEE Int High Performance Distributed Computing Symp*, pages 12–19, 1997.
- [24] J. Tang and Y. Saad. A probing method for computing the diagonal of the matrix inverse. Technical Report 42, University of Minnesota, 2010.
- [25] A. Tveit. On the complexity of matrix inversion. Mathematical Note, IDI, NTNU, Trondheim, Norway, 2003.
- [26] J. Wang, Y. Dong, X. Tong, Z. Lin, and B. Guo. Kernel Nyström method for light transport. *ACM Transaction on Graphics*, 28(3), 2009.
- [27] G. Wetzstein and O. Bimber. Radiometric Compensation through Inverse Light Transport. In *Pacific conference on computer graphics and applications*, pages 391–399, 2007.
- [28] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proc. of ACM SIGGRAPH*, pages 215–224, 1999.