

Spectral Theory of Light Transport Operators

CYRIL SOLER, INRIA - Grenoble University, France
KARTIC SUBR, University of Edinburgh, UK

Light Transport Operators (LTOs) represent a fundamental concept in computer graphics, modeling single bounces of light within a virtual environment as linear operators on infinite dimensional spaces. While the LTOs play a crucial role in rendering, prior studies have primarily focused on spectral analyses of the light field rather than the operators themselves. This paper presents a rigorous investigation into the spectral properties of the LTOs. Due to their non-compact nature, traditional spectral analysis techniques face challenges in this setting. However, many practical rendering methods effectively employ compact approximations, suggesting that non-compactness is not an absolute barrier. We show the relevance of such approximations and establish various path integral formulations of their spectrum. These findings enhance the theoretical understanding of light transport and offer new perspectives for improving rendering efficiency and accuracy.

CCS Concepts: • **Computing methodologies** → **Rendering**; • **Mathematics of computing** → *Functional analysis; Integral equations.*

Additional Key Words and Phrases: Light Transport Operator, Compactness, Fredholm Equations

ACM Reference Format:

Cyril Soler and Kartic Subr. 2025. Spectral Theory of Light Transport Operators. *ACM Trans. Graph.* 1, 1 (November 2025), 18 pages. https://doi.org/10.475/123_4

1 Introduction

Light Transport Operators (LTOs) are a mathematical construct in computer graphics, that formalizes light propagation within a virtual environment as linear operators acting on a given light distribution. These operators encapsulate one bounce of the input light distribution capturing its intricate interactions with geometric structure and material reflectance properties in the environment. During rendering, the light energy at each pixel of a virtual camera is estimated by evaluating the sum of an infinite series of increasing powers (multiple bounces) of the operators applied to the distribution of emitted light.

Despite the foundational role of transport operators in rendering, most prior studies have concentrated on spectral analyses of the light field [Durand et al. 2005; Mahajan et al. 2007] rather than the operator that governs its evolution. Consequently, the spectral properties of the LTOs remain underexplored. In this work, we use the term *spectral* strictly to refer to harmonic analysis of the LTOs

and not to wavelength-dependent phenomena. Furthermore, our analysis is restricted to monochromatic light (in the absence of fluorescence, light transport operators are defined independently across wavelengths). In this paper, we define a *scene* to mean a combination of geometry and their material properties—light sources are excluded since they are considered part of the input space on which the LTOs operate.

The LTOs inherently act in infinite-dimensional function spaces, posing substantial challenges for both theoretical investigation and computational implementation. Traditional approaches often involve discretization to make the problem tractable. Recent findings by Soler *et al.* [2022] reveal that LTOs are generally non-compact, leading to non-uniform convergence with respect to the number of elements of the discretization. Non-compactness complicates the application of classical spectral analysis results from finite-dimensional settings. Nevertheless, practical techniques such as Galerkin approximations [Baranoski et al. 1997] and precomputed radiance transport [Sloan et al. 2002] circumvent these issues by approximating the LTO with compact operators, suggesting that non-compactness is not an absolute limitation. This observation motivates deeper inquiry into the relationships between finite-dimensional approximations and the infinite-dimensional operators.

Our theoretical framework aims to bridge these gaps, connecting the eigenspectrum of the LTOs to measures of specific light path sets within a scene. By establishing these connections, we contribute to both the theoretical understanding of light transport and the development of practical methodologies for spectral analysis. We envision that our findings may have future direct implications for improving the efficiency and accuracy of rendering algorithms. In this paper, we present a detailed analysis of the spectral properties of the infinite-dimensional LTOs, with the following key contributions:

- A novel and rigorous mathematical characterization of the eigenspectrum of the LTOs;
- an investigation into how compact approximations of light transport operators can be used to approximate their spectral properties; and
- identification of connections between the eigenvalues of LTO's and path space integration, setting foundations for fresh perspectives on using eigenanalysis in rendering.

This paper presents an analysis of the light transport operator rather than of the light distribution. The operator is defined specifically by structure within a scene and so the specific choice of emitters within the scene is largely irrelevant. Occasionally, we exemplify our discussion with specific scenes that we chose to illustrate (only showing geometry), rather than render, to avoid misleading readers.

We separate the treatment of related work (Section 2) and the mathematical background needed (Section 3) to explain our work. Before performing a spectral analysis of operators that are non-compact in general, we explain how they permit compact approximations which lend themselves to analysis (Section 4). The core

Authors' Contact Information: Cyril Soler, INRIA - Grenoble University, Montbonnot, France, cyril.soler@inria.fr; Kartic Subr, University of Edinburgh, Edinburgh, UK, k.subr@ed.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM 1557-7368/2025/11-ART
https://doi.org/10.475/123_4

results surrounding spectral analysis of LTO's are presented as theorems (Section 5). We further explore how these spectra relate to different aspects of light transport (Section 6). Next, we briefly describe practical methods for computing the eigenvalues of the LTOs, providing insights into the computational challenges and opportunities in spectral analysis (Section 7). We conclude with a discussion of the implications of our findings and outline future research directions (Section 8).

2 Previous work

This paper pertains to the spectral analysis of linear operators over infinite dimensional spaces. Although the operators that we analyze have a long history in computer graphics, there has been little theoretical analysis of spectral properties. Yet, several practical approaches have exploited their *low-dimensional* approximations.

Light transport operators and their spectrum. We analyze operators that are minorly adapted from standard formulations [Arvo 1996; Veach 1997]. These formulations have provided insight to various applications such as inverting transport equations to compensate for global illumination [Ng et al. 2012], unification of forward and inverse rendering [Bai et al. 2010], analysis of compactness [Soler et al. 2022], etc. Their spectrum however has received little attention.

Baranoski [1997] examined the spectrum of the radiosity matrix in the discrete Lambertian case. Ashdown [2001] reformulated the radiosity transport matrix as an electrical network and studied its eigenvectors via analysis of the conductance matrix of the resulting network. Machida [2014] showed that, in an infinite and isotropic medium, the singular eigenfunctions of the radiative transport equations can be independently solved along each dimension.

Spectral analysis of linear operators. Although this is a vast area of study within mathematics, light transport operators exhibit challenging characteristics—such as having asymmetric unbounded and discontinuous kernels—restricting the applicability of standard results from the theory of linear operators. Despite this, we draw from a range of theoretical tools and concepts including trace-class operators [Brislaw 1988], non self-adjoint operators [Gohberg and Krein 1978], perturbation theory for linear operators [Kato 1995] and Fredholm equations [Zemyan 2012]. The connections between the eigenvalues of an operator and its discretized counterpart can be found in the work of Boffi [2010], Sun [2016], and Chatelin [2011]. We provide specific citations to the relevant sections of these works in our derivations.

Dimensionality of light transport matrices. Discrete light transport operators have been observed empirically to be conducive to dimensionality reduction. Precomputed Radiance Transfer methods [Nowrouzezahrai et al. 2007; Sloan et al. 2002, 2005; Wang et al. 2007] and modular radiance transfer [Loos et al. 2011] utilized this property to compress transport matrices. Wang [2009] exploited this property to reconstruct the matrix from a limited set of images. The solution space of light distributions obtained via low-dimensional approximations of the transport matrix is limited by the low rank of the matrix. Some methods attempt to adjust spatial sampling to optimize for this [Belcour et al. 2022; Huang and Ramamoorthi 2010] or use low rank approximations to parts of the transport

matrix [Garg et al. 2006]. Dimensionality reduction has also been applied to simplify computation in the ‘many-light’ setting with several point lights [Hašan et al. 2007; Ou and Pellacini 2011]. Mahajan *et al.* [2007] presented a theory of locally low dimensional light transport as the number of principal components required to represent the output of the transport matrix under varying conditions. This is analogous to the problem of estimating the number of eigenvalues of the light transport matrix necessary for local approximations. Lessig [2010] relates the problem of finding the effective dimension of local bandlimited and radially symmetric light transport to the spatio-spectral concentration problem. While all the above works operate on discretizations of the operator as matrices, the non-compactness of these operators [Soler et al. 2022] precludes finite dimensional uniform approximations whose spectra are not representative of the original operator's spectrum. In this paper, we analyze the spectrum of the light transport operator without discretization.

3 Background

3.1 Definitions: Linear operators and their properties

We study linear operators over Hilbert spaces. We denote spaces, operators and space elements using upper case caligraphic, upper case bold, and roman letters respectively.

Norm and convergence of linear operators. Let $\mathcal{H}$ be a Hilbert space. The norm of a linear operator $A : \mathcal{H} \rightarrow \mathcal{H}$ is defined as $\|A\| = \sup_{\|f\|=1} \|Af\|$. The remaining definitions in this section apply to operators for which the supremum exists, or *bounded linear operators*.

A sequence of linear operators $\{A_n\}$ is said to be *uniformly converging* to an operator A when

$$\lim_{n \rightarrow \infty} \|A - A_n\| = 0. \quad (1)$$

A sequence of linear operators $\{A_n\}$ is said to be *strongly converging* to an operator A when

$$\forall f \in \mathcal{H} \quad \lim_{n \rightarrow \infty} \|Af - A_n f\| = 0. \quad (2)$$

Uniform convergence implies strong convergence [Chatelin 1981].

Symmetrizable operators. The adjoint of operator $A : \mathcal{H} \rightarrow \mathcal{H}$ is the unique operator A^* for which

$$\forall (f, g) \in \mathcal{H}^2 \quad \langle Af, g \rangle = \langle f, A^*g \rangle.$$

If $A^* = A$, then A is *self-adjoint* and enjoys special properties (See Lax [Lax 2002], chap.28). If there exists a non negative self-adjoint operator S such that SA is self-adjoint, then A is known as *symmetrizable*. In this situation, A is self-adjoint for the dot product

$$\langle f, g \rangle_S = \langle f, Sg \rangle. \quad (3)$$

Symmetrizable operators in separable Hilbert spaces share many properties with self-adjoint operators.

The definition of $\langle \cdot, \cdot \rangle_S$ in Equation 3 requires S to be bounded. Alternatively, if A can be symmetrized by a positive self-adjoint operator S' with a bounded inverse on $\mathcal{H}' \subset \mathcal{H}$ (See Silbertein [1962] part II, Th. 9.1) then A is a self-adjoint operator on $\mathcal{H}'$ with dot product $\langle f, g \rangle_{S'} = \langle \sqrt{S'}f, \sqrt{S'}g \rangle$, so it still shares the properties of symmetrizable operators in this subspace.




Fig. 1. An illustrated overview of the central research question tackled in this paper. (a) The radiance distribution in a scene L (magenta) at equilibrium satisfies the rendering equation $L = E + TL$, where T is a linear operator that transports radiance by one bounce and L_e is the emitted radiance distribution (yellow). (b) The operator T is fully specified by the geometry of the scene along with the reflectance distributions associated with all surfaces in the scene. It is independent of the choice of emitted distribution L_e . (c) If the scene (and hence L) were discretized then T could be represented as a matrix and its spectrum (e.g. eigenvalues and eigenvectors) could be analyzed empirically. In this paper, we present mathematical results about the spectrum of the infinite-dimensional operator T without discretization.

Compact operators. A compact operator maps any bounded sequence in the domain into a sequence containing a converging subsequence. The family of compact operators itself forms a closed set, meaning that the limit of a uniformly converging sequence of compact operators is always a compact operator ([Gohberg and Krein 1978] p67). Practically speaking, compact operators behave like “infinite dimensional matrices”. Light transport operators are generally not compact [Soler et al. 2022].

Trace-class and Hilbert-Schmidt operators. An orthogonal basis $\{e_i\}_{i \geq 0}$ of $\mathcal{H}$ that generalizes to infinite-dimensional spaces is known as a Schauder orthogonal basis: every element in $\mathcal{H}$ can be expressed as a—possibly infinite—linear combination of e_i . Such a basis can be used to define the *trace* and the *Hilbert-Schmidt norm*

$$\text{Tr}(\mathbf{A}) = \sum_{i=0}^{\infty} \langle \mathbf{A}e_i, e_i \rangle \quad \text{and} \quad \|\mathbf{A}\|_{\text{HS}} = \sum_{i=0}^{\infty} \|\mathbf{A}e_i\|^2 \quad (4)$$

respectively. In finite-dimensional spaces, these correspond to the trace of the corresponding matrix and its Frobenius norm. *Trace-class* operators are those for which $\text{Tr}(\mathbf{A})$ converges (See Gohberg [1978] §8)¹ and *Hilbert-Schmidt* operators are operators for which the Hilbert-Schmidt norm converges. All trace-class operators are Hilbert-Schmidt and all Hilbert-Schmidt operators are compact (See Simon [Barry 2000], sec.5). The product of two Hilbert-Schmidt operators is trace-class (following Gohberg [1978] Cor. 4.1 and §4.3).

Integral operator. If $\mathcal{H}$ is a space of integrable functions over a domain S , an operator $\mathbf{A} : \mathcal{H} \rightarrow \mathcal{H}$ is an integral operator if

$$\forall f \in \mathcal{H} \quad (\mathbf{A}f)(\mathbf{x}) = \int_S \kappa(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) d\mathbf{y},$$

where $\kappa : S \times S \rightarrow \mathbb{C}$, is known as the integration kernel. A necessary and sufficient condition for an integral operator to be Hilbert-Schmidt is that $\int_S \int_S \kappa(\mathbf{x}, \mathbf{y})^2 d\mathbf{x} d\mathbf{y} < \infty$ (Gohberg [2012] Chp.IX), which is automatically verified when κ is bounded over its domain. The trace of trace-class integral operators can be computed (Barry [2000] Th. 3.9)

$$\text{Tr}(\mathbf{A}) = \int_S \kappa(\mathbf{x}, \mathbf{x}) d\mathbf{x}. \quad (5)$$

¹Another definition of the trace involving the polar decomposition of $\mathbf{A}$ is sometimes found.

Convergence of this integral does not, by itself, imply that $\mathbf{A}$ is trace-class².

3.2 Spectra of operators

In finite-dimensional spaces, linear operators are matrices. The spectrum of a $n \times n$ matrix $\mathbf{M}$ is the set of exactly n eigenvalues λ_i (accounting for multiplicity) such that $\mathbf{M} - \lambda \mathbf{I}$ is singular, and corresponds to the roots of the characteristic polynomial of $\mathbf{M}$. In infinite-dimensional spaces these definitions are slightly more complex:

Spectrum, resolvent set and resolvent operator. The spectrum of an operator $\mathbf{A}$ which we denote $\sigma(\mathbf{A})$ is the set of complex scalars λ for which $\mathbf{A} - \lambda \mathbf{I}$ is not bijective and therefore cannot be inverted. The *resolvent set* $\rho(\mathbf{A}) = \mathbb{C} \setminus \sigma(\mathbf{A})$ is the complement of the spectrum, for which the resolvent operator $R(\mathbf{A}; \lambda) = (\lambda \mathbf{I} - \mathbf{A})^{-1}$ exists. When $\lambda \in \rho(\mathbf{A})$, the resolvent operator is a bounded bijection over $\mathcal{H}$ (See Jeribi [2021] Sec 1.2.2).

Point spectrum, continuous spectrum and residual spectrum. The non-bijectivity of $\mathbf{A} - \lambda \mathbf{I}$ for $\lambda \in \sigma(\mathbf{A})$ reflects different situations: it might be non injective, or injective but not surjective in which case the range of the operator may be dense or not in $\mathcal{H}$. This defines the following partition of the spectrum:

- σ_p : the *point-spectrum* is the set of $\lambda \in \mathbb{C}$ such that $\mathbf{A} - \lambda \mathbf{I}$ is not injective;
- σ_r : the *residual spectrum* is the set of $\lambda \in \mathbb{C}$ such that $\mathbf{A} - \lambda \mathbf{I}$ is injective, but does not have a dense range;
- σ_c : the *continuous spectrum* is the set of $\lambda \in \mathbb{C}$ such that $\mathbf{A} - \lambda \mathbf{I}$ is injective with a dense range, but not bounded below.

The *eigenvalues* of $\mathbf{A}$ are the complex numbers λ such that there exists $\Lambda \in \mathcal{H}$ for which $\mathbf{A}\Lambda = \lambda\Lambda$, which corresponds to the elements of $\sigma_p(\mathbf{A})$. As such, we’re interested in this paper in the *point-spectrum* of light transport operators.

Intuition about operator spectra. Although the point-spectrum of an operator is a generalization of the eigenvalues of a matrix

²These are generally proved when κ is continuous and can be extended to domains where κ is continuous almost everywhere

to infinite-dimensional spaces, the general properties of the point-spectrum of operators cannot be extrapolated from those of matrices [Mortad 2022]. Some classical examples of operators and their point-spectrum include the Fourier transform, which point spectrum is the finite set $\{1, i, -1, -i\}$, and the unidirectional right-shift operator S_r (resp. left-shift operator S_l) over square-summable sequences in $\mathbb{C}$, for which $\sigma_p(S_r) = \emptyset$ whereas $\sigma_p(S_l)$ is the open unit disc [D’Aniello and Maiuriello 2022]. In general, σ_p may be either empty, finite, countably infinite, or even uncountable. Even “good looking” integral operators over $L^2[0, 1]$ may have any compact subset of $[0, 1]$ as a point-spectrum [Kalisch 1972]. Like matrices however, the span of eigenfunctions (including those corresponding to a zero eigenvalue) is not guaranteed to be the entire space. Except when standard theorems apply, proofs for every particular operator need a careful derivation.

One advantage of studying compact operators is that their spectrum boils down to a countable point-spectrum, with zero as the only possible accumulation point (Simon [2000] Th. 1.1). As such, they behave like “infinite-dimensional matrices”. This also explains why they are the closure of the set of finite dimensional operators, in operator norm. In this paper, we analyse the spectral properties of light transport operators.

Infinite-dimensional analogue of SVD. For bounded operators over infinite dimensional Hilbert spaces, the notion of singular value decomposition (SVD) is replaced by a continuous analogue [Crane and Gockenbach 2020], due to the fact that such operators may not have a countable set of singular values. For compact operators however, the singular values are countable (See Gohberg *et al* [Gohberg and Kreĭn 1978] Chap.2, Sec.2) and the SVD remains a discrete sum known as the *Schmidt expansion*. For a compact operator A it takes the form of

$$Af = \sum_{k=0}^{\infty} s_k \langle v_k, f \rangle u_k, \quad (6)$$

where $\{u_k\}_{k \geq 0}$ and $\{v_k\}_{k \geq 0}$ are complete orthogonal sequences in the operator value space standing for the left and right singular vectors respectively, and the s_k are the singular values (also named *s-numbers*) of A .

When an operator is self-adjoint in a specific Hilbert space (self-adjointness depends on the dot product), its *s-numbers* are equal to its eigenvalues up to the sign, which in turn do not depend on the dot-product.

3.3 Definitions: Light transport operators

Assumptions and spaces considered. We limit our formulation to monochromatic light transport with surface reflection only. The space of radiance functions is $\mathcal{H} = L^2(S \times \Omega)$, where S is the set of surfaces (with bounded area) in a scene and Ω is the upper hemisphere of directions. $\mathcal{H}$ is a separable Hilbert space with the inner product [Soler *et al.* 2022]

$$\langle l_1, l_2 \rangle_{\mathcal{H}} = \int_S \int_{\Omega} l_1(\mathbf{x}, \omega) \bar{l}_2(\mathbf{x}, \omega) \cos \theta \, d\omega \, d\mathbf{x}. \quad (7)$$

The cosine in Equation 7 imbues valuable properties to some of the operators defined below. The notation in this inner product is depicted in Figure 3.

Reflectance operator K . The reflectance equation [Kajiya 1986] expresses radiance leaving a point $\mathbf{x}$ in direction ω as:

$$L(\mathbf{x}, \omega) = \int_{\Omega} \rho(\mathbf{x}, \omega, \omega') L_i(\mathbf{x}, \omega') \cos \theta' \, d\omega', \quad (8)$$

where $L(\mathbf{x}, \omega)$ is the exitant radiance at $\mathbf{x}$ in direction ω , L_i is incident radiance and ρ is the bidirectional reflectance distribution function (BRDF) at $\mathbf{x}$. Figure 3 summarize the geometric notations. For consistency we consider “mirror” reflections to be using a very sharp true function as opposed to a distribution.

$\mathbb{C}$	complex plane
S	surfaces in the scene
Ω	hemisphere of outgoing directions (local frame)
$\mathcal{L}$	4D domain $S \times \Omega$
$\mathcal{H}$	space of radiance distributions $L_2(\mathcal{L})$
$\mathcal{B}$	space of spatial distributions $L_2(S)$
$\mathcal{O}$	space of directional distributions $L_2(\Omega)$
$\langle \cdot, \cdot \rangle_{\mathcal{H}}$	dot product weighted by cosine (Eq. 7)
$\ \cdot \ _{\mathcal{H}}$	norm induced by $\langle \cdot, \cdot \rangle_{\mathcal{H}}$
$K : \mathcal{H} \rightarrow \mathcal{H}$	global reflectance operator
$K_{\mathbf{x}} : \mathcal{O} \rightarrow \mathcal{O}$	local reflectance operator at $\mathbf{x}$
$G : \mathcal{H} \rightarrow \mathcal{H}$	re-parameterization operator
$T : \mathcal{H} \rightarrow \mathcal{H}$	light transport operator ($T = KG$)
$T_b : \mathcal{B} \rightarrow \mathcal{B}$	radiant exitance transport operator
$\rho(\mathbf{x}, \omega, \omega')$	BRDF at $\mathbf{x}$ in directions $(\omega, \omega') \in \Omega^2$
$\bar{\rho}(\mathbf{x})$	albedo at $\mathbf{x}$ (Eq. 11)
$v(\mathbf{x}, \mathbf{y})$	binary visibility function between points $\mathbf{x}$ and $\mathbf{y}$
$\kappa_b(\mathbf{x}, \mathbf{y})$	integration kernel of T_b
$\kappa(\mathbf{x}, \omega, \mathbf{y}, \omega')$	integration kernel of T
$p_{S/\Omega}$	functions mapping $(\mathbf{x}, \omega)$ to the point and direction seen from $\mathbf{x}$ in direction ω (Fig.3)

Fig. 2. Notation used in this paper.

Equation 8 defines an integral operator $K : \mathcal{H} \rightarrow \mathcal{H}$ which maps an incident radiance field into an exitant field: $L = KL_i$. Since the integral is only over Ω , the reflectance operator K is a *partial integral operator*. We may express K as an operator tensor product of $K_{\mathbf{x}} : L^2(\Omega) \rightarrow L^2(\Omega)$ at $\mathbf{x}$ and the identity $I : L^2(S) \rightarrow L^2(S)$ where for all $\mathbf{x} \in S$, $K_{\mathbf{x}}(L_i(\mathbf{x}, \cdot)) = (KL_i)(\mathbf{x}, \cdot)$. As will be seen later on, $K_{\mathbf{x}}$ is a well behaved integral operator over $\mathcal{O} = L^2(\Omega)$.




Fig. 3. Notation for Equations 7,8,9 and 12. We deviate from convention in the literature [Veach 1997] by always considering outgoing directions in the local coordinate system. Consequently $\omega'' \neq -\omega'$ in general.

Reparameterisation operator G. When radiance is constant along directions in free space (absence of volumetric scattering), the incident radiance field $L_i(\mathbf{x}, \omega')$ is equal to the exitant radiance at some location $\mathbf{y} = p_S(\mathbf{x}, \omega')$ along the direction expressed within that point's own local frame as $\omega'' = p_\Omega(\mathbf{x}, \omega')$. This reparameterisation of the radiance field is captured via an operator $G : \mathcal{H} \rightarrow \mathcal{H}$ where

$$L_i(\mathbf{x}, \omega') = (GL)(\mathbf{x}, \omega') = L(p_S(\mathbf{x}, \omega'), p_\Omega(\mathbf{x}, \omega')).$$

Transport operator T. The light transport operator T combines the reparameterisation and reflectance operators as follows:

$$(TL)(\mathbf{x}, \omega) = \int_{\Omega} \rho(\mathbf{x}, \omega, \omega') (GL)(\mathbf{x}, \omega') \cos \theta' d\omega'. \quad (9)$$

Thus $T : \mathcal{H} \rightarrow \mathcal{H}$ transports a radiance field through one-bounce. Since we consider outgoing directions as the upper hemisphere in the local coordinate system at each point, our definitions of K and G differ slightly from those in the literature [Arvo 1996; Veach 1997]. However the operator form of the rendering equation remains as

$$L = E + TL = E + KGL, \quad (10)$$

where $E \in \mathcal{H}$ is the emitted radiance field and $L \in \mathcal{H}$ is the unknown radiance.

Lambertian transport operator T_b . When all materials in a scene are Lambertian, the BRDF ρ is reduced to a directionless *albedo* $\bar{\rho}(\mathbf{x})$:

$$\bar{\rho}(\mathbf{x}) = \int_{\Omega} \rho(\mathbf{x}, \omega, \omega') \cos \theta d\omega = \pi \rho(\mathbf{x}, \cdot, \cdot). \quad (11)$$

Also, reflected radiance is independent of directions and radiant exitance $B(\mathbf{x}) = \int_{\Omega} L(\mathbf{x}, \omega) \cos \theta d\omega$, is considered instead. Rewriting Equation 9 in this space defines the radiant exitance transport operator $T_b : L^2(S) \rightarrow L^2(S)$ [Sillion and Puech 1994] as

$$(T_b B)(\mathbf{x}) = \int_S v(\mathbf{x}, \mathbf{y}) \frac{\bar{\rho}(\mathbf{x}) \cos \theta' \cos \theta''}{\pi \|\mathbf{x} - \mathbf{y}\|^2} B(\mathbf{y}) d\mathbf{y}, \quad (12)$$

where the integral is over surfaces rather than angle and $v(\mathbf{x}, \mathbf{y})$ is sometimes referred to as the “visibility function”, which restricts the integral to points $\mathbf{y}$ that are visible from $\mathbf{x}$.

Kernel expressions of T and T_b . We define a geometric term

$$g(\mathbf{x}, \mathbf{y}) = v(\mathbf{x}, \mathbf{y}) \frac{\cos \theta' \cos \theta''}{\|\mathbf{x} - \mathbf{y}\|^2}.$$

From Equation 12 we write the kernel expression of T_b as:

$$T_b B = \int_S \kappa_b(\mathbf{x}, \mathbf{y}) B(\mathbf{y}) d\mathbf{y} \quad \text{with} \quad \kappa_b(\mathbf{x}, \mathbf{y}) = \frac{\bar{\rho}(\mathbf{x})}{\pi} g(\mathbf{x}, \mathbf{y}).$$

To express T as a proper kernel-integral operator, the integral needs to be defined over the entire domain $S \times \Omega$. This requires a Dirac distribution to make the integral over S “select” pairs of elements of $\mathcal{L}$ that can actually exchange light:

$$(TL)(\mathbf{x}, \omega) = \int_{S \times \Omega} L(\mathbf{y}, \omega'') \kappa(\mathbf{x}, \omega, \mathbf{y}, \omega'') d\mathbf{y} d\omega'', \quad (13)$$

where the operator “kernel” is the distribution

$$\kappa(\mathbf{x}, \omega, \mathbf{y}, \omega') = g(\mathbf{x}, \mathbf{y}) \rho(\mathbf{x}, \omega, p_\Omega(\mathbf{y}, \omega')) \delta_{\mathbf{x}}(p_S(\mathbf{y}, \omega')).$$

3.4 Review: Properties of light transport operators

General properties. In the absence of refractive materials, operators K and G are both self-adjoint due to the cosine in the dot product of Eq. 7 [Veach 1997] (See Appendix A for a proof using our own notations). From conservation of energy, we know that $\|K\|_2 < 1$ (See Arvo [1996] Sec. 6.3.3) and $\|G\|_2 \leq 1$, where the equality occurs when the scene is “closed”. There, Arvo also proves

$$\forall p \geq 1 \quad \|T\|_p < 1 \text{ and } \|T_b\|_p < 1. \quad (14)$$

Non-compact operators. Neither K nor T are compact as they are both partial integral operators [Kalitvin and Zabrejko 1991; Soler et al. 2022]. The product of a bounded operator and a compact operator is always compact (see Kato [1995] p158), but $T = KG$ is known to be non-compact. Therefore G cannot be compact.

Conditionally-compact operators. The operator $K_{\mathbf{x}}$ is a Hilbert-Schmidt (hence compact) operator, due to its bounded kernel (the sliced BRDF in Eq. 8). Thus, the non-compactness of K is due to the “larger” space that it acts on ($\mathcal{H}$ instead of $\mathcal{O}$).

T_b is compact only in scenes where κ_b is everywhere bounded, which imposes conditions of not having abutting edges (or corners) or contact points between surfaces. At an abutting edge, it is indeed possible to construct a bounded sequence of light distributions with increasing spatial frequency content whose image under T_b does not contain any converging subsequence, proving that T_b is non compact. Fortunately, T_b coincides with a compact operator everywhere away from such edges [Soler et al. 2022].

4 Compact approximations of operators

In this section we explain that although neither T nor T_b is compact in general, they admit—not necessarily finite-dimensional—strongly converging compact approximations in the sense of Eq. 2. We also draw connections between these compact approximations and popular ‘fixes’ adopted in the practical computation of light transport in order to avoid arbitrary large values of the kernel. For example, to avoid speckles (white spots) in path-traced images and images rendered with virtual point lights [Keller 1997], multiple authors proposed to cap the geometric factor [Kollig and Keller 2006], or use a pre-integrated point-to-virtual-surface element geometry factor [Hašan et al. 2009]. Another example is vertex merging [Georgiev et al. 2012] as used in Galerkin approximations within non-Lambertian scenes, where couples of surfaces and directional elements are paired despite their position or directions not matching exactly.

Strong convergence in this context implies that the radiance function resulting from the action of a non-compact transport operator can be obtained as the limit of the value resulting from an approximating sequence of compact operators. Since compact operators form a closed set in the operator norm, these approximations will not be uniformly converging (operator-norm convergence in the sense of Eq. 1). In practice a sequence of operators T_n that is strongly converging but not uniformly converging to the light transport operator, allows to find for any $\epsilon > 0$ a sequence of light distributions $\{L_n\}$ for which $\|T_n L_n - TL_n\| > \epsilon$ for any $n \geq 0$. We will see that this difference materializes as compact approximations filtering out high

frequencies where the original operator would actually preserve them.

Compact approximations of T_b and T . As recalled in Section 3.4, the radiant exitance operator T_b is not compact since its kernel is not bounded at abutting edges. One way to overcome this is by parameterizing the kernel via a small threshold $\epsilon > 0$ so that

$$\kappa_{b,\epsilon}(\mathbf{x}, \mathbf{y}) = \frac{\bar{\rho}(\mathbf{x})}{\pi} f_\epsilon(\mathbf{x}, \mathbf{y}), \quad (15)$$

where $f_\epsilon(\mathbf{x}, \mathbf{y})$ is the capped point-to-point geometry factor³

$$f_\epsilon(\mathbf{x}, \mathbf{y}) = v(\mathbf{x}, \mathbf{y}) \min\left(\frac{1}{\epsilon}, \frac{\cos \theta' \cos \theta''}{\|\mathbf{x} - \mathbf{y}\|^2}\right). \quad (16)$$

Because $\kappa_{b,\epsilon}$ is bounded, its use as the integration kernel (instead of κ_b) leads to a Hilbert-Schmidt operator $T_{b,\epsilon}$, which is therefore compact.

In the non-Lambertian case, there are two causes for the operator T not being compact: abutting edges and the partial integration which causes the $\mathbf{1}_x$ function to appear in Equation 13. One possible way to overcome the latter while ensuring strong convergence is to replace $\mathbf{1}_x$ in κ by a Gaussian with parameterized sharpness such as

$$\kappa_\epsilon(\mathbf{x}, \omega, \mathbf{y}, \omega'') = g_\epsilon(\omega \cdot p_\Omega(\mathbf{y}, \omega'')) \rho(\mathbf{x}, \omega, p_\Omega(\mathbf{y}, \omega'')) f_\epsilon(\mathbf{x}, \mathbf{y}),$$

where g_ϵ is a normalized Gaussian of variance ϵ centered at 1. In Appendix B we prove the strong convergence of T_ϵ (resp. $T_{b,\epsilon}$) to T (resp. T_b) in the L_1 norm for bounded light distributions. While we cannot prove strong convergence for the L_2 norm, strong convergence in L_1 for bounded light distributions is sufficient in all arguments that require strong continuity in this paper. Note in particular that eigenfunctions of light transport operators are bounded because of energy conservation.




Fig. 4. Effect of Equation 15 on the first bounce of indirect illumination in the Cornell box. The flattening of curves along geometrically continuous regions is a practical consequence of the energy reduction in high frequencies by compact approximations.

Finite rank Galerkin approximations of T_b and T [Sillion et al. 1991], are alternative examples of compact approximations of these operators with strong convergence when their basis functions form a complete Schauder basis for $\mathcal{H}$ [Chatelin 2011]. This explains why finite element methods can yield a converging approximation to solve equation 10 for a specific emission function E and also explains why (due to non-uniform convergence) adaptive meshing is required for controlled error across different choices of E .

³This arbitrary choice of a capping method leads a simpler convergence proof.




Fig. 5. A toy scene with a compact (no abutting edges in the scene) radiant exitance operator that attenuates high frequencies. It is impossible to design an emission function E such that the transported radiance $L = T_b L$ contains a discontinuity. The estimate of the solution E “compensates” for the filtering by introducing arbitrarily high spatial frequencies with large intensities. Due to the strong convergence of the Galerkin approximation, this experiment provides a valid insight into the behavior of the non-discretized operator.

Compact approximations favor low frequencies. Although there is no general definition of spatial or angular “frequency” of a function, we adopt common use of this terminology as the non-zero coefficients of functions projected onto a basis with local support in the Fourier domain. For example, spatial or directional wavelets [Peers et al. 2009], Fourier bases [Durand et al. 2005] or spherical harmonics [Soler et al. 2015]. If $A : X \rightarrow X$ is a compact operator over a separable Hilbert space, and $\{\varphi_n\}_n$ is any orthonormal Schauder basis of $L^2(X)$, we prove in appendix D that

$$\lim_{n \rightarrow \infty} \|A\varphi_n\| = 0. \quad (17)$$

In the specific situation of a basis with local frequency support, Eq. 17 directly proves that A attenuates higher frequencies⁴. This result applies to compact approximations $T_{b,\epsilon}$ and T_ϵ . This effect is illustrated on Figure 4.

Equation 17 does not imply that compact operators are low-pass filters. Rather, that the action of the operator results in stronger attenuation of higher frequencies. Thus, Lambertian light transport tends to be “low frequency” away from edges. For example, consider assigning an emission distribution E to a planar emitter some distance from a parallel planar receiver. If we wish the receiver to contain a step edge in the spatial distribution of reflected radiance L (e.g. as seen in a hard shadow), solving for E such that $L = T_b E$ is not possible. An attempt to approximate E would contain arbitrarily high spatial frequencies at large energy to compensate for Equation 17. This is illustrated in Figure 5. This theory underpins the reason that numerical approximations of light transport inherently attenuate high spatial and angular frequencies.

Since compact approximations have eigenvalues that converge to 0, when the corresponding eigenfunctions form a complete system, they necessarily increase in spatial (or angular) frequency with the index of the corresponding eigenvalue. We prove this in Appendix C.

⁴Note that a sequence of functions with *decreasing* frequency would not be complete, so the same argument cannot be used to show that compact operators filter low frequencies.

Convergence of eigenelements. While the spectrum of a parameter dependent approximation is generally "continuous" w.r.t. to norm-convergence (see Conway [1991] p431), less can be expected when only strong convergence is satisfied. Generally speaking, strong convergence of a holomorphic family of parameter-dependent compact operators (such as T_ϵ and $T_{b,\epsilon}$) may display very unexpected behaviors. For example, the approximation could have an infinite number of eigenvalues even though the operator (when $\epsilon = 0$) might have no eigenvalues at all, or an uncountable point-spectrum, etc (See Kato [1995], VII.4, p371).

A useful result however is that when a parameter dependent approximation A_ϵ of an operator A strongly converges to A , then the isolated eigenvalues of A are the limit of eigenvalues of A_ϵ [Chatelin 1981, 2011], neglecting multiplicity. This theory simplifies even further when the operator and its approximations are both self-adjoint and densely defined, in which case the convergence properties cover both eigenvalues and eigenprojections (See Chatelin [1981] definition 2.1 and Section 2.3). Consequently, both eigenvalues and eigenprojections of $T_{b,\epsilon}$ converge toward those of T_b (We show in Sec. 5.1 that T_b is self-adjoint in a particular Hilbert space).

Contrary to T_b , having no formal proof that the point-spectrum of T is countable, the previous conditions do not apply. We are currently limited to observing the convergence of eigenvalues in a number of typical geometric configurations, which tends to suggest that the point-spectrum of T is countable (See Figure 6). Numerical experiments indeed show that Galerkin approximations of T are stable (In the sense of Chatelin [Chatelin 1981]). A formal proof of such stable convergence (left for future work) would allow to apply Proposition 2.2 of F.Chatelin's paper [Chatelin 1981] in order to show that eigenvalues of these approximations actually converge to isolated elements of the point-spectrum of T .

The strong convergence of approximations T_ϵ and $T_{b,\epsilon}$ also mean that partial eigenfunction expansions of these operators strongly converge, which validates the use of eigenfunctions and eigenvalues of these compact approximations to represent the behavior of the operators over finite dimensional sets of light distributions.

Usage for eigenfunction expansions. Expanding elements of $\mathcal{H}$ with the eigenfunctions of an operator requires the family of eigenfunctions of that operator to be dense in $\mathcal{H}$ (accounting for eigenvalue 0). This is the case for T_b (See Section 5), but this is yet to be proven for T . Regardless, the compactness of T_ϵ is not sufficient to ensure completeness of its eigenfunctions⁵. The eigenfunctions Λ_i of T_b (resp. $T_{b,\epsilon}$) are

$$\forall L \in \mathcal{B} \quad L = \sum_{i=0}^{\infty} \langle \Lambda_i, L \rangle_S \Lambda_i. \quad (18)$$

For T_ϵ , we can at least say that a similar equation holds for elements selected inside the span of its eigenfunctions.

In summary we have justified in this section that the eigendecomposition of the compact approximations of light transport can be used to estimate the countable part of the operators' point spectrum, and that the largest eigenvalues correspond to the energy exchanges of lowest frequency.

⁵A notable counterexample in finite dimensions is that of defective matrices.

5 Point-Spectrum of transport operators

In this section, we analyze the point-spectra of operators T , K , G and T_b . For each operator, we define the elements that are contained within its point-spectrum as a theorem accompanied by a proof. Figure 9 provides a table summarizing our results.

5.1 Spectrum of the radiant exitance transport operator

Although T_b is the infinite-dimensional analog of the radiosity matrix, its non-compactness in the general case precludes direct extensions of the properties of the latter. However, a common 'trick' to make the radiosity matrix symmetric positive definite [Baranoski et al. 1997] is useful towards proving that T_b is symmetrizable.

Assuming that the albedo $\bar{\rho}$ is never null, let $S : \mathcal{B} \rightarrow \mathcal{B}$ be the trivial self-adjoint linear operator defined as

$$(SB)(\mathbf{x}) = \frac{1}{\bar{\rho}(\mathbf{x})} B(\mathbf{x}). \quad (19)$$

This operator is also positive since

$$\forall B \in \mathcal{B} \quad \langle SB, B \rangle = \left\langle \frac{1}{\bar{\rho}} B, B \right\rangle = \int_S (\rho(\mathbf{x})^{-1/2} |B|)^2 d\mathbf{x}.$$

It is therefore possible to define a new dot-product

$$\forall f, g \in L^2(S) \quad \langle f, g \rangle_S = \langle f, Sg \rangle,$$

and we have

$$\langle T_b f, g \rangle_S = \langle T_b f, Sg \rangle = \langle f, ST_b g \rangle = \langle f, T_b g \rangle_S.$$

The operator ST_b is trivially self-adjoint over $\mathcal{B}$, and $ST_b = T_b S$. In other words, T_b is symmetrizable on $\mathcal{B}$, which also means that it is self-adjoint on the Hilbert space $\mathcal{B}_S = (L^2(S), \langle \cdot, \cdot \rangle_S)$. We obtain the following result:

THEOREM 5.1.1: *The point-spectrum of T_b has the following properties:*

- (1) $\sigma(T_b) = \sigma_p(T_b)$;
- (2) $\sigma_p(T_b) \subset (-1, 1)$;
- (3) $\sigma_p(T_b)$ is countable.
- (4) eigenfunctions of T_b for different eigenvalues are orthogonal w.r.t to the dot product $\langle \cdot, \cdot \rangle_S$;
- (5) eigenfunctions of T_b span $\mathcal{B}$.
- (6) the first eigenfunction of T_b is positive everywhere

PROOF: Property (2) is the consequence of the point-spectrum of self-adjoint operators being real (See Lax [Lax 2002], chap.28) combined with Equation 14. Other properties are direct consequences of the *spectral theorem*: A self-adjoint operator A on Hilbert space has no residual spectrum, proving (1), and its eigenfunctions for different eigenvalues are orthogonal (4) (See for instance Reed [1972], Th. VI.8). Because $\mathcal{B}$ is separable these eigenfunctions can therefore only be countable (3) – in a separable space every orthogonal sequence is countable. From the multiplication form of the spectral theorem ([Reed and Simon 1972], Th. VII.3), we deduce that the eigenfunctions of T_b form a Schauder basis of $\mathcal{B}$ (5). The last property is a consequence of T_b being a "positive operator" (in the sense of Lax [Lax 2002], p253). $\square$




Fig. 6. Convergence and stability of the largest eigenvalues of a Galerkin approximation of the non-Lambertian transport operator T in a toy scene made of two squares with the MERL gold-paint material with an increasing number of elements. Left: as opposed to the Lambertian case, the eigenvalues of T are not necessarily real, because of the inherent non-symmetry of T . Zooming on the red squared region of the complex plane (*center*) where we show four eigenvalues computed with an increasing number of surface elements, suggests that they display consistent convergence. Since pairs of surface elements are used to define directional finite elements, the effective matrix size is the square of the indicated numbers (e.g. max size here is $300^2 \times 300^2$). Note that eigenvalues come in this scene by pairs of opposite signs because of the symmetry of the scene.

Self-adjointness alone does not imply that eigenfunctions of T_b span $\mathcal{B}$. The operator's point-spectrum needs to be countable which happens in this case when the space is separable, or more generally when the operator is compact. It is possible that 0 belongs to $\sigma_p(T_b)$, for instance when the geometry includes some finite surface that is not visible from any other surface. If 0 is in the point spectrum, a complete system of eigenfunctions of T_b includes its corresponding eigenfunctions. Finally, a complete system of eigenfunctions with no null eigenvalue does not make the operator invertible, since the span of eigenfunctions is only dense in $\mathcal{B}$ and not equal to $\mathcal{B}$.

Distribution of the eigenvalues. When T_b is compact, in addition to being countable, $\sigma_p(T_b)$ can only have 0 as an accumulation point. When it is not compact, eigenvalues are still countable (since the operator is self-adjoint and the space is separable), and will therefore necessarily accumulate at extra non zero values within $(-1, 1)$. These eigenvalues correspond to eigenfunctions for which most the energy is concentrated next to an abutting edge. This explains why compact approximations of T_b are practical for image synthesis: while they remove non negligible eigenvalues, the corresponding energy exchanges are mostly local.

When T_b is compact, it is also Hilbert-Schmidt as are its compact approximations $T_{b,\epsilon}$. Since it is self-adjoint in $\mathcal{B}_S$, we can apply Eq. 4 to an orthogonal basis of eigenfunctions in $\mathcal{B}_S$ to get $\sum_i \lambda_i^2 < \infty$. This brings the following upper bound on the asymptotic behavior of its eigenvalues: $\lambda_n = o(1/\sqrt{n})$. Figure 7 demonstrates this using a simple scene with no abutting edges. Determining the actual asymptotic behavior of the spectrum is non-trivial and depends on the specific operator (see for instance Jeribi [2021] Sec 11.4).

Finally, while finite multiplicity of eigenvalues is not a limitation in this paper, it is fairly straightforward to construct scenes where the eigenvalues of T_b have multiplicity greater than 1: given a closed scene S_1 (no light leaks), we build a new scene by duplicating the geometry of S_1 multiple times, to obtain as many closed and separate environments. Every eigenpair (λ, Λ) of T_b in S_1 will give birth to as many eigenpairs in the new scene where each closed subpart will have Λ while other subparts are 0. Having an infinite number of scaled copies which total surface area stays finite would allow




Fig. 7. Toy example of a scene where T_b is compact showing the upper bound on the asymptotic decrease of eigenvalues.

to construct a scene where eigenvalues have infinite multiplicities, making the operator non-compact even if the original scene had no abutting edges. This however is outside of physical constraints.

Singular value decomposition of T_b . Since T_b is self-adjoint in $\mathcal{B}_S$, and given that its eigenvalues $\{\lambda_k\}_{k \geq 0}$ are real and countable, and its eigenfunctions $\{\Lambda_k\}_{k \geq 0}$ form a complete orthogonal basis of $\mathcal{B}_S$ (See Section 5.1), T_b has the following “Schmidt expansion” in $\mathcal{B}_S$ ⁶:

$$T_b L = \sum_{k=0}^{\infty} |\lambda_k| \langle L, \Lambda_k \rangle_S \underbrace{\text{sign}(\lambda_k) \Lambda_k}_{u_k}.$$

The sign of λ_k has been transferred to the left singular functions u_k in order to keep the singular values non negative. This expansion is however not uniformly converging since the eigenvalues of T_b may have non zero accumulation points when the operator is not compact. In $\mathcal{B}$ however, there is no such trivial expression for a

⁶Note that the term *schmidt expansion* is not perfectly appropriate here since the convergence is not uniform.

Schmidt expansion of T_b , and numerical computations bring singular values that are different from its eigenvalues, mostly because the eigenfunctions of T_b are not orthogonal in this space.

5.2 Spectrum of the reflectance operators

The local and global scattering operators K_x and K have different spectra. While they share eigenvalues, the eigenvalues of the latter have infinite dimensionality.

THEOREM 5.2.1: *The spectrum of K_x is equal to its point-spectrum, and has the following characteristics:*

- (1) $\sigma_p(K_x) \subset (-1, 1)$
- (2) the eigenfunctions of K_x span $\mathcal{O}$

PROOF: This operator is Hilbert-Schmidt (therefore compact) and self-adjoint. Its eigenvalues are consequently real and equal to its singular values, and its eigenfunctions—including eigenvalue 0, if present—span the half-space of directions $\mathcal{O}$. Furthermore the eigenvalues of K_x lie in the open interval $(-1, 1)$ since energy conservation imposes $\|K_x\| < 1$. Since K_x is compact, the only possible accumulation point for its eigenvalues is zero, which may or not be part of the point-spectrum even if the rank of K_x is infinite. $\square$

Figure 8 shows the largest eigenvalues of K_x for some elements of the MERL database. Since the integral kernel of K_x is the BRDF, we sometimes refer to its eigenvalues as the “eigenvalues of the BRDF”.

THEOREM 5.2.2: *The point-spectrum of K has the following characteristics:*

- case 1:** if there is no measurable area where the BRDFs all share the same eigenvalue then $\sigma_p(K) = \emptyset$;
- case 2:** else if the scene can be partitioned into a finite number of measurable areas with respective BRDFs ρ_i , then $\sigma_p(K) = \bigcup_{i \geq 0} \sigma_p(\rho_i)$. Its eigenvalues have infinite dimensionality and its eigenfunctions span $\mathcal{H}$. This case is typical of common light transport scenes;
- case 3:** else, the point spectrum of K will be limited to the material eigenvalues that are shared across measurable subsets of S , each having infinite dimensionality. In this case the eigenfunctions of K do not span $\mathcal{H}$.

PROOF: Because K is a tensor product between K_x and an identity operator over $L^2(S)$, the existence of an eigenpair (k, ϕ) of K also implies that $K_x \phi(x, \cdot) = k \phi(x, \cdot)$ wherever $\phi(x, \cdot) \neq 0$, which means that $(k, \phi(x, \cdot))$ is also an eigenpair of the BRDF at x . And, in order to be a non-zero function in $\mathcal{H} = L^2(S \times \Omega)$, the property $\phi(x, \cdot) \neq 0$ should remain valid over a measurable subset $A \subset S$. On the contrary, when the BRDFs across $A \subset S$ all share the same eigenvalue k with eigenfunction ρ_x , any function of the form $(x, \omega) \mapsto \rho_x(\omega)f(x)$ over A and 0 elsewhere is an eigenfunction of K , and from Th. 5.2.1, $L^2(A \times \Omega)$ is spanned by these eigenfunctions. $\square$

5.3 Spectrum of the reparametrization operator

Let $\mathcal{L}_1 \subset \mathcal{L}$ denote the set of points $(x, \omega) \in \mathcal{L}$ for which $p_x(x, \omega)$ is defined, $\mathcal{H}_1 = L^2(\mathcal{L}_1)$ and $P_1 : \mathcal{H} \rightarrow \mathcal{H}_1$ the orthogonal projection on $\mathcal{H}_1$. We also call $G_1 : \mathcal{H}_1 \rightarrow \mathcal{H}_1$ the restriction of G to $\mathcal{H}_1$. Note that because p_x is reversible wherever it is defined, using $\mathcal{H}_1$ for




Fig. 8. *Top:* Point-spectrum of some reflectance operators from the MERL database. *Bottom:* eigenfunctions 0, 6 and 20 of the gold-paint material, computed by projecting the BRDF onto a basis of spherical harmonics up to degree 30 (green means negative). As theory predicts (See Lax [2002], chap.28), only the first eigenfunction is entirely positive.

the output space of G_1 makes sense. With these settings, we have

$$G = G_1 P_1 \text{ and } G_1^2 = I. \quad (20)$$

Similarly to G operator G_1 is self-adjoint, and $\|G_1\| = 1$ which makes it an isometry on $\mathcal{H}_1$. Its eigenvalues are therefore ± 1 . Furthermore, to any eigenpair (λ, V_1) of G_1 corresponds an eigenpair (λ, V) of G with $V(x, \omega) = 0$ for any $(x, \omega) \notin \mathcal{L}_1$.

Finally, any function $V \in \mathcal{H}$ such that $P_1 V = 0$ satisfies $GV = 0$. Therefore 0 is also in the point-spectrum of G . No other eigenvalue can exist, as it would automatically be an eigenvalue of G_1 . In closed scenes, $P_1 = I$, $\mathcal{L}_1 = \mathcal{L}$, and $G_1 = G$, and the part of the point-spectrum corresponding to the nul eigenvalue vanishes.

THEOREM 5.3.1: *Spectrum of the reparametrization operator*

- (1) $\sigma_p(G) = \{-1, 1\}$ in closed scenes;
- (2) $\sigma_p(G) = \{-1, 0, 1\}$ otherwise;
- (3) the eigenfunctions of G form a complete family of $\mathcal{H}$.

PROOF: In order to prove (3), let $\{\varphi_n\}_{n \geq 0}$ (resp. $\{\psi_n\}_{n \geq 0}$) be an orthogonal Schauder basis of $\mathcal{H}_1$ (resp. $\mathcal{H} \setminus \mathcal{H}_1$)⁷. Assuming that both functions can be extended by 0 in $\mathcal{L}$, we define the sequence:

$$\begin{aligned} G_{3n} &= \frac{1}{2}(\varphi_n + G\varphi_n) \\ G_{3n+1} &= \frac{1}{2}(\varphi_n - G\varphi_n). \\ G_{3n+2} &= \psi_n. \end{aligned}$$

⁷These spaces are separable, and any such combination of spatial and directional wavelets would work.

Since $G^2 = I$ we obtain

$$GG_{3n} = G_{3n} \quad GG_{3n+1} = -G_{3n+1} \quad GG_{3n+2} = 0.$$

After noting that $\varphi_n = G_{3n} + G_{3n+1}$ we see that $\{G_n\}_{n \geq 0}$ spans $\mathcal{H}$, which makes it a complete—yet not orthogonal inside each eigenspace—family of eigenfunctions of G . An fully orthogonal family of eigenfunction can further be obtained by Gram-Schmidt orthogonalization in each eigenspace. $\square$

5.4 Spectrum of the radiance transport operator

Unlike the point-spectrum of T_b (Section 5.1), much less can be determined about the point-spectrum of T because of the combination of the following properties:

- (1) T is not compact;
- (2) T is not normal;
- (3) T cannot be symmetrized.

The non-compactness of T has been discussed in Section 3.4. To demonstrate non-normality, since K and G are self-adjoint, we have

$$\begin{aligned} TT^* &= KGGK = K^2, \\ T^*T &= GKKG = GK^2G. \end{aligned}$$

That is, reflecting twice is not equivalent to transporting, reflecting twice and transporting back. Note that neither is T quasi-normal nor hyponormal, which would yield beneficial properties [Conway 1991]. (3) is a side effect of the Dirac in Equation 13, which precludes multiplication of T by a trivial operator (such as S in Eq. 19) to make it self-adjoint.

THEOREM 5.4.1: *The spectrum of T has the following properties:*

- (1) $\sigma_p(T)$ is not empty; the largest eigenvalue is positive with a positive eigenfunction;
- (2) $\sigma_p(T)$ lies in the disk of radius $\|K\|$;
- (3) the eigenfunctions of T for non-zero eigenvalues do not necessarily span $\mathcal{H}$.

PROOF: Since T is a bounded operator it has a non-empty spectrum (See Kato [1995] Ch3, §6.2), but this does not guarantee the same property for its point-spectrum. But since its kernel is positive (assuming non-zero reflectance), the Krein-Rutman theorem [Phát and Dieu 1994] applied to the cone of light distributions of constant sign ensures that when the scene is closed we have $\sigma_p(T) \neq \emptyset$ and its maximal eigenvalue is a positive real value of unit multiplicity, associated to an all-positive eigenfunction. When the scene is not closed, restricting T to $\mathcal{H}_1$ preserves the conditions to apply this theorem, hence proving (1). Note that the same applies to T_b which coincides with T when materials are Lambertian. (2) is straightforward: let (V, λ) be an eigenpair of T (which means $\|V\| = 1$), then $\|TV\| = |\lambda|$ and $\|TV\| \leq \|K\|\|G\|$. Using Equation 14 we get $|\lambda| \leq \|K\|$.

Since $T = KG$, any eigenfunction V of T is in the range of K , so is the span of the eigenfunctions, which proves (3). The extreme case is in Lambertian scenes, where eigenfunctions for non-zero eigenvalues are directionally constant, and therefore cannot represent any non-directionally constant function in $\mathcal{H}$. $\square$

The following might be useful in practical contexts where one set of eigenfunctions might be computed easily, either for transporting light or importance:

THEOREM 5.4.2: *The spectrum of T and T^* are connected via the operators G and K :*

- (1) The eigenfunctions of T and T^* are mutually orthogonal for different eigenvalues;
- (2) G turns the former into the later, while K turns the later into the former;
- (3) $\sigma_p(T) = \sigma_p(T^*)$.

PROOF: Orthogonality is a classical result. In order to prove (2) and (3), let (λ, L) be an eigenpair of T^* . We have $GKL = \lambda L$. Left multiplying by G leads to $KL = \lambda GL$. Reporting the former equation in the later, we get $K(GKL) = \lambda G(GKL)$, which simplifies to $T(KL) = \lambda(KL)$. Conversely, if (λ, V) is an eigenpair of T , we have $KGL = \lambda L$. Left-multiplying by G directly leads to $T^*(GL) = \lambda(GL)$. $\square$

Note that the above properties do not imply that eigenvalues are real, and indeed, numerical experiments tend to show that even in the simplest cases they are not (See Figure 6). Finally, we lack a formal proof that $\sigma_p(T)$ is countable. It seems to be a reasonable conjecture given the stable convergence of eigenvalues of its Galerkin approximations [Chatelin 2011] (See Figure 6). We leave such a proof for future work.

5.5 Singular value decomposition of radiance transport operator

Interestingly, T admits a simple explicit singular value decomposition in scenes with a finite number of non-spatially-varying materials m_i ⁸ which we briefly recall here for completeness [Soler et al. 2022]. This result will be used to derive the spectral properties of T . We call ρ_{ij} the j^{th} eigenvalue of material m_i and its corresponding eigenfunction $r_{ij} \in \mathcal{O}$. We define $\{\phi_k^i\}_{k \geq 0}$ to be an orthogonal basis of the subset of S where the material is m_i (wavelets are such an example). Finally we name $g : j \mapsto (n_1, n_2)$ an arbitrary bijection between $\mathbb{N}$ and $\mathbb{N}^2$, and use $j = g(n_1, n_2)$. Thus, we have

$$TL = \sum_{i=1}^n \sum_{j>0} |\rho_{ij}| \langle L, \phi_j^i \rangle \psi_j^i, \quad (21)$$

where for all $(n_1, n_2) \in \mathbb{N}^2$,

$$\psi_j^i(\mathbf{x}, \omega) = \phi_{n_1}^i(\mathbf{x}) r_{n_2}^i(\omega) \quad (22)$$

$$\phi_j^i(\mathbf{x}, \omega) = \text{sign}(\rho_{ij}) G \psi_j^i(\mathbf{x}, \omega). \quad (23)$$

The above equation takes the form of Equation 6 after reordering the elements in the sum.

6 Light transport formulations of the spectrum

We now derive connections between the eigenvalues of compact approximations of the light transport operators and light paths and their integrals. In some cases, the connection is direct such as the structure of the operators' output space being governed by its eigenelements (Section 6.1). In Section 6.2 we discover a relationship between eigenvalues of the operator and linear combinations of multiple bounces of light. In Section 6.3 we explain how the measure of single light paths of fixed length, as a pure geometric quantity (with no specific light distribution involved), is also closely related

⁸This is the most general situation, assuming that surfaces with spatially-varying materials can further be separated in sub-regions with constant materials

operator	space	normality	class	point-spectrum	completeness
T_b	$\mathcal{B}$	symmetrizable	bounded	countable	yes
$T_{b,\epsilon}$	$\mathcal{B}$	symmetrizable	Hilbert-Schmidt	countable $\rightarrow 0$	yes
K_x	$\mathcal{O}$	self-adjoint	Hilbert-Schmidt	countable $\rightarrow 0$	yes
K	$\mathcal{H}$	self-adjoint	bounded	countable (∞ -dim)	yes
G	$\mathcal{H}$	self-adjoint	unitary	$\{0, 1, -1\}$	yes
T	$\mathcal{H}$	none	bounded	countable (conjectured)	conjectured
T_ϵ	$\mathcal{H}$	none	Hilbert-Schmidt	countable $\rightarrow 0$	conjectured

Fig. 9. Summary of spectral properties of the different operators involved in light transport (See definitions in Section 3.1). Completeness accounts for eigenfunctions corresponding to eigenvalue 0.

to the eigenvalues. Finally in 6.4 we show that eigenfunctions can be expressed as an integral over solutions of a parameter-dependent generalized light transport problem.

All results derived in this section remain valid of any compact approximation of T and T_b , including such as those described in Section 4. We use the generic notation $\tilde{T}$ to designate them. Note that while some of the following results would be straightforward for matrices, careful proofs are still required for these infinite dimensional operators.

6.1 Structure of the operator's output space

Projecting $\tilde{T}$ onto a finite N -dimensional subspace of its output space, results in a matrix whose eigenvalues estimate the largest elements of $\sigma_p(\tilde{T})$. This formulation is inspired from the Monte-Carlo approximations of large matrices proposed by Halko [2011].

Let $\mathcal{H}_N$ be the N -dimensional subspace of $\mathcal{H}$ spanned by the N first elements of a complete sequence of functions of finite support in the frequency domain (wavelets, spherical harmonics, etc). Consider p light distributions $L_1, \dots, L_p$ in this space sampled from a Gaussian distribution. For some n , let $q_1, \dots, q_n$ be an orthogonal basis of the n -dimensional span Q of $\{\tilde{T}L_1, \tilde{T}L_2, \dots, \tilde{T}L_p\}$ (most of the time $n = p$). We also define

$$Q : \mathbb{R}^n \rightarrow Q$$

$$(a_1, \dots, a_n) \mapsto a_1 q_1 + \dots + a_n q_n.$$

Using this notation, the adjoint $Q^* : Q \rightarrow \mathbb{R}^n$ of Q computes the a_i coordinates for any $L \in Q$. Both Q and Q^* can be plunged into $\mathcal{H}$ so that QQ^* is the orthogonal projection on Q , whereas Q^*Q is the identity over $\mathbb{R}^n$. Given this, the following theorem holds.

THEOREM 6.1.1: *For any $\epsilon > 0$ there exists $(n, N) \in \mathbb{N}^2$ such that the $n \times n$ matrix $M = Q^* \tilde{T} Q$ satisfies*

$$(Mv = \lambda v) \Rightarrow \|\tilde{T}Qv - \lambda Qv\| \leq \epsilon,$$

with probability $1 - 10^{-n}$. That is, the eigenvalues of M statistically approximate some eigenvalues of $\tilde{T}$ using sufficiently many input light distributions L_i .

PROOF: Let $\epsilon > 0$ and $T_N = P_N \tilde{T}$ where P_N is the orthogonal projection on $\mathcal{H}_N$. Since $\tilde{T}$ is compact, there exists N such that $\|T_N - \tilde{T}\| \leq \epsilon/3$. We therefore have

$$\|\tilde{T} - QQ^* \tilde{T}\| \leq \frac{\epsilon}{3} + \|T_N - QQ^* T_N\| + \|QQ^* (T_N - \tilde{T})\|.$$

The 3rd term above verifies $\|QQ^* (T_N - \tilde{T})\| \leq \|QQ^*\| \|T_N - \tilde{T}\| \leq \epsilon/3$ since, as an orthogonal projection $\|QQ^*\| = 1$. The second term is majorated using the framework of Halko [2011] (Sections 4.3 and 4.4) which we rephrase here in our context:

PROPOSITION 6.1.1 (HALKO, 2011). *Let A be a $N \times P$ matrix, and $\omega_0, \omega_1, \dots$ a sequence of random Gaussian vectors, and Q the matrix built as above. As soon as n consecutive vectors verify $\|(I - QQ^*)A\omega_k\| \leq \epsilon/(10\sqrt{\frac{2}{\pi}})$, then with probability $1 - 10^{-n}$ we have $\|(I - QQ^*)A\| \leq \epsilon$.*

We apply this result to the operator $T_N P_N$, that is a finite rank operator over the finite dimensional space $\mathcal{H}_N$ which therefore can be assimilated to a matrix, to determine the existence of n such that $\|T_N - QQ^* T_N\| \leq \epsilon/3$. In summary, there exists N and n such that

$$\|\tilde{T} - QQ^* \tilde{T}\| \leq \epsilon. \quad (24)$$

Let (λ, v) be an eigenpair of M . We have

$$\|\tilde{T}Qv - \lambda Qv\| \leq \|\tilde{T}Qv - QQ^* \tilde{T}Qv\| + \underbrace{\|QQ^* \tilde{T}Qv - \lambda Qv\|}_{QMv}.$$

The first term on the right hand side is less than ϵ because of Eq. 24 and since $Mv = \lambda v$, the second term is equal to 0, which yields $\|\tilde{T}Qv - \lambda Qv\| < \epsilon$. $\square$

This theorem should be understood as follows: $\mathcal{H}_N$ being the space where the L_i are sampled, choosing n large enough will reduce the approximation error $\|\tilde{T}Qv - \lambda Qv\|$ down to the threshold imposed by N . However, using a very large N will require n to be large as well in order to reduce $\|T_N - QQ^* T_N\|$. Halko [2011] conjectures that n should be slightly larger than N in order to satisfy Eq. 24. In other words, controlling the error requires sampling a small number of low frequency light distributions, or a larger number of higher-frequency light distributions.

As defined in the theorem, the matrix $M = Q^* \tilde{T} Q$ resembles a product of matrices, but it is not. Operator Q may be viewed as an "infinite matrix" with n columns each of which is one of the light distributions q_i . However, we may estimate M by sampling these light distributions at a finite (yet sufficiently large) number of points and directions.

6.2 Linear combinations of powers of $\tilde{T}$

Here we derive that the minimization of a linear combination of powers of T results in a polynomial whose roots estimate the largest

elements of the point-spectrum of $\mathbf{T}$. While such a connection seems obvious for a matrix (choosing $1 + P(x)$ to be proportional to the characteristic polynomial of that matrix) the following theorem shows that it holds for our compact approximations.

THEOREM 6.2.1: *For any light distribution $L \in \mathcal{H}$ with a finite support on the sequence of eigenfunctions of $\tilde{\mathbf{T}}$ there exists a finite set of coefficients $\alpha_1, \dots, \alpha_n$ such that*

$$L + \sum_{k=1}^n \alpha_k \tilde{\mathbf{T}}^k L = 0. \quad (25)$$

Furthermore, the roots of polynomial $1 + \sum_k \alpha_k x^k$ are the corresponding eigenvalues of $\tilde{\mathbf{T}}$.

PROOF: If $\{\Lambda_i\}$ is the sequence of eigenfunctions of $\tilde{\mathbf{T}}$ (which we do not assume to be complete nor orthogonal), and L has a finite support on this basis, then there exist $n + 1$ coefficients $\beta_i \neq 0$ (without loss of generality we assume no gap in the related Λ_i) such that

$$L = \sum_{p=0}^n \beta_p \Lambda_p. \quad (26)$$

Let $\alpha_0, \dots, \alpha_n \in \mathbb{C}$. Applying $\tilde{\mathbf{T}}^k$ to equation 26, and exchanging sums (both are finite) we immediately get

$$\sum_{k=0}^n \alpha_k \tilde{\mathbf{T}}^k L = \sum_{p=0}^n \beta_p P(\lambda_p) \Lambda_p. \quad (27)$$

where P is the polynomial $\sum_{k=0}^n \alpha_k X^k$, hence defining operator $P(\tilde{\mathbf{T}})$ (See Conway [1990] Sec.4.10)⁹.

Because of equation 27 the necessary and sufficient condition for $P(\tilde{\mathbf{T}})L = 0$ is that $P(\lambda_p) = 0$ for each p . Such a polynomial being defined up to a factor, we choose $\alpha_0 = 1$. Polynomial $\prod_{i=0}^n (1 - \frac{1}{\lambda_i} X)$ is in this case the unique solution, whose expansion provides the α_i coefficients for Equation 25. $\square$

In practice, this theorem proves useful to estimate eigenvalues because of the increasing oscillatory nature of eigenfunctions of $\tilde{\mathbf{T}}$, which means that any finite-frequency light distribution will have a limited support on the sequence of eigenfunctions, or at least very small values of β_n beyond some N . Consequently even when the coefficients do not exactly vanish, optimizing

$$(\alpha_1, \dots, \alpha_n) = \operatorname{argmin} \|L + \sum_{k=0}^n \alpha_k \tilde{\mathbf{T}}^k L\|^2 \quad (28)$$

produces a polynomial whose roots will approach the largest eigenvalues. We illustrate this in Figure 10.

Theorem 6.2.1 is related to theorem 6.1.1 in the following way: let's choose $\{L_k\}$ (in theorem 6.1.1) to belong to the span $\mathcal{M}$ of the first eigenfunctions of $\tilde{\mathbf{T}}$, and call C the characteristic polynomial of matrix $\mathbf{M}$. Because $\mathcal{M}$ is stable by $\tilde{\mathbf{T}}$ we have $C(\mathbf{M}) = \mathbf{Q}^* C(\tilde{\mathbf{T}}) \mathbf{Q}$. Since $C(\mathbf{M}) = 0$ (Cailey-Hamilton theorem), we also have $\mathbf{Q}^* C(\tilde{\mathbf{T}}) \mathbf{Q} = 0$. Consequently, for any $L \in \mathcal{M}$ in theorem 6.2.1 the

⁹Note that self-adjointness is not required here

corresponding polynomial P is C (or a factor of C depending on the support of L over the $\{\Lambda_i\}$).

6.3 The measure of circular light paths

In this subsection we reveal an unexpected connection between the measure of circular light paths in a scene and the eigenvalues of $\tilde{\mathbf{T}}$.

THEOREM 6.3.1: *The eigenvalues of $\tilde{\mathbf{T}}$ are the reciprocals of the poles of the entire function $d(\lambda) = \sum_{n=0}^{\infty} (-1)^n a_n \lambda^n$, where*

$$a_n = -\frac{1}{n} \sum_{k=2}^n (-1)^k p_k a_{n-k}, \quad (29)$$

and p_n is the measure of circular light paths of length n in the scene, with $p_0 = a_0 = 1$ and $p_1 = a_1 = 0$.

PROOF: In Section 4 we saw that the compact approximations $\mathbf{T}_\epsilon$ and $\mathbf{T}_{b,\epsilon}$ of light transport operators also happen to be Hilbert-Schmidt. Assuming $\tilde{\mathbf{T}}$ is one of these, we can express the Hilbert-Carleman determinant (See Gohberg *et al.* [2012] p176) of $\tilde{\mathbf{T}}$ as

$$\operatorname{Det}(\mathbf{I} + \lambda \tilde{\mathbf{T}}) = \sum_{n=0}^{\infty} a_n \lambda^n \quad \text{with} \quad (30)$$

$$a_n = \frac{1}{n!} \begin{vmatrix} 0 & n-1 & 0 & \dots & 0 & 0 \\ \operatorname{Tr}(\tilde{\mathbf{T}}^2) & 0 & n-2 & 0 & \dots & 0 \\ \operatorname{Tr}(\tilde{\mathbf{T}}^3) & \operatorname{Tr}(\tilde{\mathbf{T}}^2) & 0 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ \operatorname{Tr}(\tilde{\mathbf{T}}^n) & \operatorname{Tr}(\tilde{\mathbf{T}}^{n-1}) & \operatorname{Tr}(\tilde{\mathbf{T}}^{n-2}) & \dots & \operatorname{Tr}(\tilde{\mathbf{T}}^2) & 0 \end{vmatrix}$$

with $a_0 = 1$ and $a_1 = 0$. Since $\tilde{\mathbf{T}}$ is Hilbert-Schmidt, all operators $\tilde{\mathbf{T}}^k$ are also trace-class for $k \geq 2$, hereby justifying the use of the trace in Equation 30. Just like the more classical *Fredholm determinant* (Gohberg *et al.* [2012] Chp.VI), the Hilbert-Carleman determinant is an entire function over $\mathbb{C}$ which zeros are the $\lambda = -\frac{1}{\lambda_i}$, where λ_i are the eigenvalues of $\tilde{\mathbf{T}}$.

For $k \geq 2$, denoting by κ_k the kernel of the trace-class operator $\tilde{\mathbf{T}}^k$, thanks to Equation 5 we have [Brislaw 1988]

$$\begin{aligned} \operatorname{Tr}(\tilde{\mathbf{T}}^k) &= \int_S \kappa_k(\mathbf{x}, \mathbf{x}) d\mathbf{x} \\ &= \int_S \int_{S^{k-1}} \kappa(\mathbf{x}, \mathbf{x}_1) \kappa_{k-1}(\mathbf{x}_1, \mathbf{x}) d\mathbf{x}_1 d\mathbf{x} \\ &\dots \\ &= \int_{S^k} \kappa(\mathbf{x}_1, \mathbf{x}_2) \kappa(\mathbf{x}_2, \mathbf{x}_3) \dots \kappa(\mathbf{x}_k, \mathbf{x}_1) d\mathbf{x}_1 \dots d\mathbf{x}_k \end{aligned} \quad (31)$$

In the non lambertian case the $\mathbf{x}_i$ in these equations should be understood as couples of points and directions. Because of the arbitrarily sharp Gaussian g_ϵ in the kernel expression of $\tilde{\mathbf{T}}$, only the ones that actually constitute a correct light path are non zero. In other words for both Lambertian and non-Lambertian scenes the above equations prove that $\operatorname{Tr}(\tilde{\mathbf{T}}^k)$ corresponds to the integral of all circular light paths of length k in the scene. Henceforth, we use the




Fig. 10. Illustration of Th.6.2.1, where we show that the largest eigenvalues of $\tilde{T}$ can be approximated as roots of a polynomial which coefficients α_k zero the first bounces of light $\tilde{T}^k L$.

notation $p_k = \text{Tr}(\tilde{T}^k)$. Rewriting the determinant in Equation 30 using Laplace expansion along the first column leads to the recursive expression of Equation 29 (See Appendix E for further details). $\square$

The simplest method to compute p_k is Monte-Carlo integration, via a random starting point and direction followed by importance sampling of the path up to length k . The last vertex is then connected to the first¹⁰. Doing so, the probability density of sampling the circular path $\pi_k = (x_1, \dots, x_k)$ is

$$P(\pi_k) = \frac{1}{S} \kappa(x_1, x_2) \dots \kappa(x_{k-1}, x_k),$$

and we have

$$\begin{aligned} p_k &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \frac{1}{P(\pi_k)} \kappa(x_1, x_2) \kappa(x_2, x_3) \dots \kappa(x_k, x_1) \\ &= \lim_{N \rightarrow \infty} \frac{S}{N} \sum_{i=1}^N \kappa(x_k, x_1). \end{aligned}$$

Even though the calculation of p_n is stable, estimating eigenvalues via a finite number of these path integrals converges slowly. The roots of polynomials are generally highly sensitive to noise in their coefficients, especially when these roots are evenly spaced [Wilkinson 1959]. This is why in a practical calculation scenario truncating $\text{Det}(I + \lambda \tilde{T})$ to its first terms causes instability when determining its roots. Figure 11 shows practical values of p_n and λ_n in a simple scene where the values vanish rapidly with increasing n .

6.4 Parameter-dependent light transport solutions

The previous subsections derived connections between path integrals and eigenvalues. We now explain how it is also possible to express eigenfunctions as a sum of solutions to a parameter-dependent light transport problem in the complex domain.

¹⁰Note that implementation efficiency is not an issue at this point, and we acknowledge that better sampling strategies that reuse the same path for multiple p_i are available




Fig. 11. The measure of n -bounce circular light paths p_n shown for n up to 16 on an example scene (inset). The scene is composed asymmetrically of 3 Lambertian squares (albedo 0.5) so that κ is bounded, making T_b compact.

THEOREM 6.4.1: Any eigenfunction Λ_i of $\tilde{T}$ corresponding to an eigenvalue λ_i with unit geometric multiplicity can be expressed as a path integral

$$\Lambda_i = -\frac{1}{2i\pi} \oint_{\Gamma} L(\alpha) d\alpha, \quad (32)$$

where $L(\alpha)$ is the solution to the “generalized light transport problem”

$$\alpha L = E + \tilde{T}L, \quad (33)$$

and $\Gamma \subset \mathbb{C}$ is any simple closed curve enclosing the sole eigenvalue λ_i , and E any light distribution such that $\langle \Lambda_i, E \rangle \neq 0$.

PROOF: The explanation to this intriguing theorem relies on the fact that because $L(\alpha) = (\alpha - \tilde{T})^{-1} E = R(\tilde{T}; \alpha) E$, Eq. 32 corresponds to the Riesz integral of the resolvent, that is known to be a projection onto the union of eigenspaces of the operator which corresponding eigenvalues are enclosed by Γ (See for instance Gohberg [1978] ch.I). The integrated “light” distribution in Equation 32 will therefore be equal to an (unnormalized) eigenfunction of $\tilde{T}$, or 0 depending on whether $\langle E, \Lambda_i \rangle \neq 0$ or not. $\square$

Equation 33 can be viewed as a generalization of the traditional light transport equation 10 by introducing a parameter α . While

the existence of a solution to this equation is guaranteed by the Fredholm alternative theorem [Zemyan 2012] (since $\alpha \in \Gamma \subset \rho(\tilde{T})$), its solution need not be positive nor even real. Thus, it would not represent a physically meaningful distribution of light. Nevertheless, this is helpful to effectively compute an eigenfunction of $\tilde{T}$.

A practical, yet non computationally efficient way to compute Λ_i is to apply a Galerkin approximation to $\tilde{T}$ and solve for $(\alpha - \tilde{T})L(\alpha) = E$ using a biconjugate gradient. Computing this solution using the traditional Monte-Carlo light transport paradigm is however not feasible, since the series obtained by expanding $(\alpha - \tilde{T})^{-1}$ only converges for $|\alpha| > \|\tilde{T}\|$. This precludes any curve Γ that would only encircle a strict subset of the point-spectrum of $\tilde{T}$.

7 Practical computation of the spectrum

In this section, we suggest a couple of practical methods to compute the eigenvalues of the light transport operators T_b and T .

7.1 Lambertian operator T_b : Galerkin is best

Following the work of Chatelin [1981], since the point-spectrum of T_b is countable, a converging approximation of the eigenvalues of T_b can be obtained by computing the eigenvalues of an n -element Galerkin approximation with matrix A_n of that operator. The scene is partitioned into n surfaces of area S_i and the light transport operator is projected onto a finite element basis. For a basis of piecewise constant functions, the matrix entry $a_n^{i,j}$ of the projected transport operator is [Sillion and Puech 1994]

$$a_n^{i,j} = \frac{1}{S_i} \int_{S_i} \int_{S_j} \kappa_b(\mathbf{x}_i, \mathbf{x}_j) d\mathbf{x}_i d\mathbf{x}_j \quad (34)$$

where S_i the area of surface element i and κ_b the kernel of T_b (See Section 3.3). The eigenvalues of the matrix can then be calculated using a standard eigensolver.

7.2 Non-Lambertian operator T : Monte-Carlo approach

In non-Lambertian scenes, a Galerkin approach would require discretization of surfaces and directions so that Equation 34 integrates over pairs of surface element-direction pairs $\ell_i = S_k \times \Omega_i$

$$a_n^{i,j} = \frac{1}{\mu(\ell_i)} \int_{\ell_i} \int_{\ell_j} \kappa(\mathbf{x}, \omega, \mathbf{y}, \omega'') d\mathbf{x} d\omega d\mathbf{y} d\omega'' \quad (35)$$

where μ is the measure associated with ℓ . Unfortunately, this approach is computationally impractical. Compared to the Lambertian case where the eigensolver would be applied on a matrix that is about 1000×1000 , the non-Lambertian case for comparable accuracy would be $10^6 \times 10^6$ or larger. The matrix in Eq.35 is very sparse, and even more when reflectance functions are highly glossy. Computing a_{ij} is however expensive because the correlation between the geometric factor in κ and the BRDF cannot be neglected while maintaining physical consistency.

Since Equation 35 is an integral, $a_n^{i,j}$ is the expected value of an appropriately scaled random variable:

$$a_n^{i,j} = \mu(\ell_i) \mathbb{E}(\kappa(\mathbf{x}, \omega, \mathbf{y}, \omega'')).$$

Consequently we form the following random matrix:

$$A_n = \frac{2\pi S}{n} \left\| \kappa(\mathbf{x}_i, \omega_i, \mathbf{y}_j, \omega_j'') \right\|_{ij},$$

where points and directions are randomly sampled. In practice, we suggest to select n random points $\mathbf{x}_i$, and outgoing directions for each point are implicitly defined by the connection to all other points at the cost of a bias. This also ties sparsity to the glossiness of reflectance functions only, since in this formulation only the directional integration is implicitly considered. For compact approximations T_ϵ of T (see Section 4) the eigenvalues of A_n form an estimator of the eigenvalues of T_ϵ [Chatelin 1981]. The estimator is biased since the eigenvalues of A_n do not linearly depend on its coefficients, but the bias converges to 0 when n tends to infinity. Note that because of the lack of a proof that $\sigma_p(T)$ is actually countable and T_ϵ strongly converges to T , we can only ensure that this method computes converging approximations of the isolated elements of $\sigma_p(T)$ (See Section 4).

We use an eigensolver that only requires computing the product of vectors by A_n or its transpose [Lehoucq et al. 1997], which is performed efficiently as a triple loop over the $\mathbf{x}_i$. Figure 12 shows an example of using this technique on a non-Lambertian scene. The method presented in this section shares similarity with random sub-matrix sampling as proposed by Frieze [2004] to compute singular values of large matrices.

8 Discussion

8.1 Discussion of results

Challenges specific to light transport operators. Despite the large array of mathematical tools to analyse the spectra of linear operators, light transport operators T and T_b present unique challenges. Although their non-compactness can be mitigated by approximations, these approximations inherit some of the bad properties of the operators: non normality (or any sub-normality class), discontinuous integration kernel, non-compact resolvent. These severely limit the set of applicable methods. Although they are not positive operators or symmetric, the fact that they are integral operators with a strictly positive kernel allows some of the analyses presented.

Relationship between T and its adjoint. A curious result from Theorem 5.4.1 is that the adjoint T^* has the same point spectrum (eigenvalues) as T . The adjoint operator has a practical significance in light transport, as it represents the transport of importance or tracking light backwards from the eye to the scene. Our result confirms that the two operators share similar spectral properties and therefore similar computational challenges with respect to convergence. Although T is not self-adjoint, $T = KG$ where K and G are self-adjoint operators. As a result, T and T^* have mutually orthogonal eigenfunctions and operators K and G can be used to transform one into the other and vice versa. This could potentially have practical implications in precomputed light transport where one of them is diagonalized.

Relationship between T and T_b for diffuse scenes. In Lambertian scenes, any eigenfunction Λ_b of T_b can be turned into an eigenfunction of T as $\Lambda(\mathbf{x}, \omega) = \Lambda_b(\mathbf{x})$. Conversely, any eigenfunction of T with a non-zero eigenvalue will have a constant directional component (by means of $T\Lambda = \lambda\Lambda$) and is similarly tied to an eigenfunction of T_b . Additionally, 0 is an eigenvalue of T with infinite




Fig. 12. *Left*: validation of the Monte-Carlo estimation of eigenvalues of Section 7.2 on the two-squares scene of Figure 6 (Lambertian version with albedo 0.8). When increasing the number of points n , eigenvalue histograms show sharper peaks that exactly align with the reference eigenvalues. *Middle*: Monte-Carlo estimation of eigenvalues in the CornellBox scene. *Right*: 3D histogram in the complex plane of Monte-carlo estimation of eigenvalues in the non-Lambertian scene of Figure 6.

dimensionality, corresponding to all directional functions of mean value zero.

Generally T is not compact as a partial integral operator, it is therefore counterintuitive that it may have the point-spectrum of a compact operator plus eigenvalue 0 (T_b is compact when the scene has no abutting edges). In Lambertian scenes however T can be restricted to the space of directionally constant functions for which a single value represent the whole function, which in effect makes T a non partial integral operator. Note that in this case, the eigenspaces of T (including eigenvalue 0) are complete and its eigenvalues are countable.

Specific case of $\rho_{ij} \geq 0$. An interesting observation is that because of Equation 21, when all the ρ_{ij} are positive, T can be left-symmetrized by the operator

$$\begin{aligned} R : T(\mathcal{H}) &\rightarrow T(\mathcal{H}) \\ RTL &= \sum_{i=1}^n \sum_{j>0} \langle L, \varphi_j^i \rangle \psi_j^i, \end{aligned} \quad (36)$$

which in this case makes RT self-adjoint (the proof involves showing that $\langle \varphi_j^i, \psi_l^k \rangle = \langle \varphi_l^k, \psi_j^i \rangle$ which only happens when φ_j^i does not include the negative sign of some of the ρ_{ij}). Since R has eigenvalues $|\rho_{ij}|^{-1}$, it is bounded below by 1 (which means it has a bounded inverse) and we can in this case apply Th. 9.1 of Silberstein [1962] to deduce that T is symmetrizable over $T(\mathcal{H})$. Consequently its eigenvalues are real, and its eigenfunctions are orthogonal for the dot product $\langle f, g \rangle_R = \langle \sqrt{R}f, \sqrt{R}g \rangle$ [Silberstein 1962]. Furthermore orthogonality implies countability because $T(\mathcal{H}) \subset \mathcal{H}$ is separable. Since G is invertible on $\mathcal{H}$, we know that $K(\mathcal{H}) = T(\mathcal{H})$ and $\ker(T) = G^{-1}(\ker(K))$. Since K is self-adjoint, its eigenfunctions (including those in $\ker(K)$) span $\mathcal{H}$, and therefore eigenfunctions of T span $\mathcal{H}$, proving a strong result for the spectrum of T .

This happens at least in Lambertian scenes, where the only eigenvalues of K correspond to the albedo of all materials in the scene. In classical scenes, it seems that BRDFs have both positive and negative eigenvalues (See Figure 8). Although the effect of flipping the sign of negative eigenvalues of BRDFs is unclear, this would preserve the frequency content of material reflectances [Durand et al. 2005]. Perhaps this case would be useful in the context of analysis of light fields via a symmetrizable operator.

Estimating eigenfunctions from α_k . The result from Theorem 6.2.1 means that if L belongs to the span of a finite set of eigenfunctions of a compact approximation of T or T_b , the roots of polynomial P given by the theorem gives the corresponding eigenvalues. Given one such λ_i , one can also expand $Q(x) = P(x)/(x - \lambda_i)$ and obtain $Q(T)L = \beta_i \Lambda_i$. While a random L will generally not be in the span of a finite set of eigenfunctions of T , the above expression only approximate an eigenfunction, assuming that eigenvalues decrease fast enough for the contribution of $\{\Lambda_i\}_{i>n}$ to be small enough.

Total response of the generalized transport system. The measures of circular light paths p_k defined in Section 6.3 can be used as a measure of the total response of the generalized transport equation $\alpha L = E + TL$ to border conditions (E in this expression). This quantity is generally defined as

$$d = \int_{\mathcal{L}} r_{\alpha}(x, x) dx,$$

where r_{α} is the integral kernel [Brislaw 1988] of the resolvent operator $R(T; \alpha)$. Recalling the calculation leading to Eq. 31, the resolvent of $\tilde{T}$ is clearly not trace class due to the terms $\alpha I + \tilde{T}$. Dropping the first two terms still allows to define a regularized total response $\tilde{d}$, which turns into

$$\tilde{d} = \alpha^{-1} p_2 + \alpha^{-2} p_3 + \alpha^{-3} p_4 + \dots \quad (37)$$

This gives a formal proof of the intuitive following fact: for compact approximations of T , the larger the measure of circular light paths, the more energy will be in the solution of the generalized equation.

8.2 Open questions

Eigenfunctions of T . The spectrum of T diagonalizes the operator and so estimation of the eigenvalues and eigenfunctions of T would simplify the calculation of higher powers (multiple-bounce transport), inversion (relighting, inverse rendering), etc. In this paper, we have analyzed the spectra (eigenvalues and eigenfunctions) mathematically and estimated eigenvalues numerically. However, stable methods for the estimation of the eigenfunctions of T remain elusive.

Quantifying scene complexity using p_n . The complexity of a scene with respect to computing light transport solutions is difficult to quantify. There are a variety of factors that make computation

challenging, including number of light sources, reflectances in the scene, visibility, geometric elements, etc. We conjecture that the measure of circular light paths in a scene may be useful for quantitative assessment of the interplay between these factors. As explained in Section 6.3, p_n is a single number per length of circular paths n obtained as the trace of $\mathbf{T}^n$. Its distribution over n encodes information about the relative contributions of different path lengths to the light transport solution. Thus, scenes where p_n has a heavy or long tail may be more challenging to solve accurately than scenes where p_n has a sharp peak for lower values of n .

Is the point-spectrum of $\mathbf{T}$ countable? Despite our efforts, we were unable to derive a formal proof for this challenging question. Since $\mathcal{H}$ is a separable Hilbert space, completeness of the eigenfunctions of $\mathbf{T}$ would prove countability of its point-spectrum. However $\mathbf{T}$ does not meet the conditions for Goh'berg theorems on completeness of eigenfunctions ([Gohberg and Kreĭn 1978] Chp.V), nor is it dissipative, quasinormal or hyponormal [Gohberg and Kreĭn 1978]. Exploiting the characteristics of the convergence of $\mathbf{T}_\epsilon$ toward $\mathbf{T}$ through asymptotic perturbation theory may be an interesting direction to consider (See Chatelin [2011], and Kato [1995] Chp.8).

New class of light transport algorithms? Solving the generalized light transport equation, $\lambda L = E + \mathbf{T}L$ for $|\lambda| < \|\mathbf{T}\|$, with path tracing could be an important step toward practical calculation of the eigenfunctions of $\mathbf{T}$, using the Riesz integral in Equation 32. As explained in Section 6.4, a generic formulation of its solution as an integral over light paths remains elusive. Such a formulation may reveal fresh perspective of the rendering equation and potentially lead to a new class of light transport algorithms.

9 Conclusion and future work

In this paper we have analyzed the spectral characteristics of light transport operators $\mathbf{K}$, $\mathbf{K}_x$, $\mathbf{G}$, $\mathbf{T}_b$ and $\mathbf{T}$. We demonstrated that, although neither $\mathbf{T}_b$ nor $\mathbf{T}$ are compact, these two operators can still be strongly approximated by compact operators. Further, we highlight that these approximations actually correspond to the ones used in practice while implementing path tracing. We have also revealed various connections between the point spectrum of these operators and path integrals in a scene, combining circular light paths, powers of the operator and projection on the operators' output space. We identify important questions that remain unanswered and leave them as future work in this area.

In conclusion, we believe that there is much to be discovered at the confluence of stochastic approximations of large matrices [Dimov et al. 2015; Frieze et al. 2004; Halko et al. 2011; Kobayashi et al. 2001] and light transport estimation.

Acknowledgments

Kartic Subr was supported by a Royal Society University Research Fellowship.

A Self-adjointness of $\mathbf{G}$ and $\mathbf{K}$

Although $\mathbf{K}$ and $\mathbf{G}$ are long known to be self-adjoint [Veach 1997], we derive this using our notations. For two points $\mathbf{x}$ and $\mathbf{y}$ that are

visible to each other (via local directions ω' and ω'' as in Figure 2),

$$\langle \mathbf{G}L_1, L_2 \rangle = \int_S \int_\Omega (\mathbf{G}L_1)(\mathbf{x}, \omega') \overline{L_2}(\mathbf{x}, \omega') \cos \theta' d\mathbf{x} d\omega'.$$

Now we change variables from $(\mathbf{x}, \omega')$ into $(\mathbf{y}, \omega'')$. We have $d\omega' = dy \cos \theta'' / r^2$ and $d\omega'' = dx \cos \theta' / r^2$, from which we get $dy \cos \theta'' d\omega'' = dx \cos \theta' d\omega'$. Using the fact that $(\mathbf{G}L_1)(\mathbf{x}, \omega') = L_1(\mathbf{y}, \omega'')$ and $\overline{L_2}(\mathbf{x}, \omega') = (\overline{\mathbf{G}L_2})(\mathbf{y}, \omega'') = \overline{\mathbf{G}L_2}(\mathbf{y}, \omega'')$, the above equation turns into

$$\begin{aligned} \langle \mathbf{G}L_1, L_2 \rangle &= \int_S \int_\Omega L_1(\mathbf{y}, \omega'') \overline{\mathbf{G}L_2}(\mathbf{y}, \omega'') \cos \theta'' d\mathbf{y} d\omega'' \\ &= \langle L_1, \mathbf{G}L_2 \rangle. \end{aligned}$$

Similarly,

$$\begin{aligned} \langle \mathbf{K}L_1, L_2 \rangle &= \int_S \int_\Omega (\mathbf{K}L_1)(\mathbf{x}, \omega) \overline{L_2}(\mathbf{x}, \omega) \cos \theta d\mathbf{x} d\omega \\ &= \int_S \int_\Omega \left(\int_\Omega \rho(\mathbf{x}, \omega, \omega') L_1(\mathbf{x}, \omega') \cos \theta' d\omega' \right) \overline{L_2}(\mathbf{x}, \omega) \cos \theta d\mathbf{x} d\omega \end{aligned}$$

Since $\rho(\mathbf{x}, \omega, \omega') = \rho(\mathbf{x}, \omega', \omega)$ due to Helmholtz reciprocity and $\mathbf{K}L_2 = \overline{\mathbf{K}L_2}$, the two definite integrals over Ω can be exchanged to yield

$$\langle \mathbf{K}L_1, L_2 \rangle = \langle L_1, \mathbf{K}L_2 \rangle.$$

Both operators are self-adjoint thanks to the cosine in the dot product. $\square$

B Strong convergence of $\mathbf{T}_\epsilon$ and $\mathbf{T}_{b,\epsilon}$ in L_1 norm

Let L be a bounded light distribution. For any $\epsilon > 0$ we denote by $S_\epsilon(\mathbf{x})$ the set of points $\mathbf{y}$ for which $\kappa_b(\mathbf{x}, \mathbf{y}) > \frac{1}{\epsilon}$, and M_ϵ the set

$$M_\epsilon = \left\{ \mathbf{x} \in S \mid \exists \mathbf{y} \in S \mid \kappa(\mathbf{x}, \mathbf{y}) \geq \frac{1}{\epsilon} \right\}.$$

Clearly for all $\mathbf{x} \in S$ we have $S_\epsilon(\mathbf{x}) \subset M_\epsilon$. Using successively $S_\epsilon \subset S$, Hölder's inequality and the non-negativity of κ_b we get

$$\begin{aligned} |((T_b - T_{b,\epsilon})L)(\mathbf{x})| &\leq \int_{S_\epsilon(\mathbf{x})} \left| \left(\kappa_b(\mathbf{x}, \mathbf{y}) - \frac{1}{\epsilon} \right) L(\mathbf{y}) \right| d\mathbf{y} \\ &\leq \|L\|_\infty \int_{S_\epsilon(\mathbf{x})} \left| \kappa_b(\mathbf{x}, \mathbf{y}) - \frac{1}{\epsilon} \right| d\mathbf{y} \\ &\leq \|L\|_\infty \int_{S_\epsilon(\mathbf{x})} \kappa_b(\mathbf{x}, \mathbf{y}) d\mathbf{y}. \\ &\leq \|L\|_\infty \int_{M_\epsilon} \kappa_b(\mathbf{x}, \mathbf{y}) d\mathbf{y}. \end{aligned} \quad (38)$$

Therefore

$$\begin{aligned} \|(T_b - T_{b,\epsilon})L\|_1 &\leq \|L\|_\infty \int_S \int_{M_\epsilon} \kappa_b(\mathbf{x}, \mathbf{y}) d\mathbf{y} d\mathbf{x} \\ &\leq \|L\|_\infty \underbrace{\int_{M_\epsilon} \int_S \kappa_b(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y}}_{\leq 1} \\ &\leq \|L\|_\infty \mu(M_\epsilon) \end{aligned}$$

In classical physically meaningful scenes, where edges and corners constitute a non measurable set of points, the measure $\mu(M_\epsilon)$ verifies

$\lim_{\epsilon \rightarrow 0} \mu(M_\epsilon) = 0$, therefore

$$\lim_{\epsilon \rightarrow 0} \|(T_b - T_{b,\epsilon})L\|_1 = 0.$$

The strong convergence of T_ϵ to T is immediate after noticing that for any directional function $\omega \mapsto h(\omega)$ we have whenever h is continuous at ω ,

$$\lim_{\epsilon \rightarrow 0} \int_{\Omega} h(\omega') g_\epsilon(\omega \cdot \omega') d\omega' = h(\omega). \quad (39)$$

Once again, the "proper" geometry of scenes ensuring that light distributions are only non continuous at a set of points which measure is 0, and applying the preceding proof for the remaining term $f_\epsilon(\mathbf{x}, \mathbf{y})$, we deduce the strong convergence of T_ϵ to T almost everywhere. $\square$

C Eigenfunctions increase in frequency

In this proof T denotes a compact approximation of one of the transport operators. Let $\{\varphi_n\}_{n \geq 0}$ be an orthogonal basis of the space with non decreasing frequency content (e.g. directional/spatial wavelets). Let $\{\Lambda_k\}_{k \geq 0}$ be the eigenfunctions of T assumed to form a complete system.

Assume, first, that eigenfunctions are orthogonal. Then, for all n

$$\varphi_n = \sum_{k=0}^{\infty} \alpha_{nk} \Lambda_k \quad \text{with} \quad \alpha_{nk} = \langle \varphi_n, \Lambda_k \rangle.$$

The convergence of that sum implies that

$$\forall n \in \mathbb{N} \quad \lim_{k \rightarrow \infty} \alpha_{nk}^2 = 0.$$

Selecting the first $N + 1$ values of n allows us to apply the limit to the supremum of the α_{nk} :

$$\forall N \in \mathbb{N} \quad \forall n \in [0, N] \quad \lim_{k \rightarrow \infty} \sup_{n \leq N} \alpha_{nk}^2 = 0,$$

which exactly states that when k increases, the coefficients of Λ_k over $\{\varphi_n\}$ uniformly converge to 0, over every interval $[0, N]$. Since $\sum_n \alpha_{nk}^2 = 1$, that means most of the energy in Λ_k lies beyond N .

Now if eigenfunctions are not orthogonal, we can still apply Gram-Schmidt orthogonalization to $\{\Lambda_k\}$ (which remains valid for countably infinite sequences) to obtain an orthogonal sequence $\{\Lambda'_k\}_{k \geq 0}$ such that

$$\varphi_n = \sum_{k=0}^{\infty} \alpha'_{nk} \Lambda'_k \quad \text{with} \quad \alpha'_{nk} = \langle \varphi_n, \Lambda'_k \rangle.$$

Therefore, Λ'_k increases in frequency with k . But since Λ'_k is built using Gram-Schmidt orthogonalization, we know that for all k , Λ_k is a linear combination of $\Lambda'_0, \dots, \Lambda'_k$, and thus the increase in frequency applies to Λ_k . $\square$

D Proof of Equation 17

For any $N \in \mathbb{N}$ we call $P_N : \mathcal{H} \rightarrow \mathcal{H}$ the projection on the finite dimensional space spanned by $\{\varphi_n\}_{n < N}$, and we define the finite rank operator A_n to be

$$A_n = AP_n.$$

Since A is compact and the sequence $\{\varphi_n\}_{n > 0}$ is complete, we know that A_n converges to A in the operator norm [Chatelin 1981]:

$$\forall \epsilon > 0 \quad \exists p \in \mathbb{N} \quad \forall n \geq p \quad \|A_n - A\| < \epsilon.$$

That means

$$\forall \epsilon > 0 \quad \exists p \in \mathbb{N} \quad \forall n \geq p \quad \forall L \in \mathcal{A} \quad \|(A_n - A)L\| < \epsilon \|L\|.$$

We now restrict the above equation to $L = \varphi_{n+1}$ for every possible $n > p$. We have in this case $(A - A_n)\varphi_{n+1} = A\varphi_{n+1}$ and $\|\varphi_{n+1}\| = 1$, which gives after shifting $n + 1$ to n :

$$\forall \epsilon > 0 \quad \exists p \in \mathbb{N} \quad \forall n > p \quad \|A\varphi_n\| < \epsilon. \quad \square$$

E Proof of Equation 29

Let A_n be the matrix in Eq. 30 and let $c_n = \det(A_n)$. In order to prove Eq. 29 one needs to notice that when expanding this determinant along the first column, the k^{th} minor m_k is

$$m_k = \begin{vmatrix} \mathbf{D}_{n,k} & \mathbf{0} \\ \mathbf{B}_{n,k} & \mathbf{A}_{n-k} \end{vmatrix},$$

where $\mathbf{D}_{n,k}$ is a $(k - 1) \times (k - 1)$ lower-triangular matrix with $n - 1, n - 2, \dots, n - k + 1$ on its diagonal, making the bottom-left submatrix $\mathbf{B}_{n,k}$ irrelevant in this calculation. Following this, we get

$$c_n = \sum_{k=2}^n (-1)^{k+1} p_k m_k = \sum_{k=2}^n (-1)^{k+1} \frac{(n-1)!}{(n-k)!} p_k c_{n-k}$$

Substituting $n! a_n = c_n$ directly leads to Eq. 29.

References

- James Arvo. 1996. *Analytic Methods for Simulated Light Transport*. Ph.D. Dissertation. Yale University.
- Ian Ashdown. 2001. *Eigenvector radiosity*. Ph.D. Dissertation. University of British Columbia. <https://doi.org/10.14288/1.0051491>
- Jiamin Bai, Manmohan Chandraker, Tian-Tsong Ng, and Ravi Ramamoorthi. 2010. A Dual Theory of Inverse and Forward Light Transport. In *Computer Vision – ECCV 2010 (Lecture Notes in Computer Science)*, Kostas Daniilidis, Petros Maragos, and Nikos Paragios (Eds.). Springer, Berlin, Heidelberg, 294–307. https://doi.org/10.1007/978-3-642-15552-9_22
- Gladimir V. G. Baranoski, Randall Bramley, and Jon G. Rokne. 1997. Eigen-Analysis for Radiosity Systems. In *Proceedings of the Sixth International Conference on Computational Graphics and Visualization Techniques*. Vilamoura, Algarve, Portugal.
- Simon Barry. 2000. *Trace Ideals and Their Applications: Second Edition* (second edition ed.). Mathematical Surveys and Monographs, Vol. 120. American Mathematical Society. <https://bookstore.ams.org/surv-120-s>
- Laurent Belcour, Thomas Deliot, Wilhem Barbier, and Cyril Soler. 2022. A Data-Driven Paradigm for Precomputed Radiance Transfer. *Proc. ACM Comput. Graph. Interact. Tech.* 5, 3 (July 2022), 26:1–26:15. <https://doi.org/10.1145/3543864>
- Daniele Boffi. 2010. Finite element approximation of eigenvalue problems. *Acta Numerica* 19 (May 2010), 1–120. <https://doi.org/10.1017/S0962492910000012> Publisher: Cambridge University Press.
- Chris Brislawn. 1988. Kernels of Trace Class Operators. *Proc. Amer. Math. Soc.* 104, 4 (1988), 1181–1190. <https://doi.org/10.2307/2047610> Publisher: American Mathematical Society.
- F. Chatelin. 1981. The Spectral Approximation of Linear Operators with Applications to the Computation of Eigenlements of Differential and Integral Operators. *SIAM Rev.* 23, 4 (Oct. 1981), 495–522. <https://doi.org/10.1137/1023099>
- Françoise Chatelin. 2011. *Spectral Approximation of Linear Operators*. Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9781611970678>
- John Conway. 1991. *The Theory of Subnormal Operators*. Mathematical Surveys and Monographs, Vol. 36. American Mathematical Society, Providence, Rhode Island. <https://doi.org/10.1090/surv/036>
- J. B. Conway. 1990. *A Course in Functional Analysis* (2 ed.). Springer-Verlag, New York. <https://www.springer.com/gp/book/9780387972459>
- Daniel K. Crane and Mark S. Gockenbach. 2020. The Singular Value Expansion for Arbitrary Bounded Linear Operators. *Mathematics* 8, 8 (Aug. 2020), 1346. <https://doi.org/10.3390/math8081346> Number: 8 Publisher: Multidisciplinary Digital Publishing Institute.
- Ivan Dimov, Sylvain Maire, and Jean Michel Sellier. 2015. A new Walk on Equations Monte Carlo method for solving systems of linear algebraic equations. *Applied Mathematical Modelling* 39, 15 (Aug. 2015), 4494–4510. <https://doi.org/10.1016/j.apm.2014.12.018>

- Frédo Durand, Nicolas Holzschuch, Cyril Soler, Eric Chan, and François X. Sillion. 2005. A Frequency Analysis of Light Transport. In *ACM SIGGRAPH 2005 Papers (SIGGRAPH '05)*. ACM, New York, NY, USA, 1115–1126. <https://doi.org/10.1145/1186822.1073320>
- Emma D'Aniello and Martina Maiuriello. 2022. On the spectrum of weighted shifts. *Revista de la Real Academia de Ciencias Exactas, Físicas y Naturales. Serie A. Matemáticas* 117, 1 (Oct. 2022), 4. <https://doi.org/10.1007/s13398-022-01328-z>
- Alan Frieze, Ravi Kannan, and Santosh Vempala. 2004. Fast monte-carlo algorithms for finding low-rank approximations. *J. ACM* 51, 6 (Nov. 2004), 1025–1041. <https://doi.org/10.1145/1039488.1039494>
- Gaurav Garg, Eino-Ville Talvala, Marc Levoy, and Hendrik P. Lensch. 2006. Symmetric Photography: Exploiting Data-sparseness in Reflectance Fields. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques (EGSR '06)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 251–262. <https://doi.org/10.2312/EGWR/EGSR06/251-262>
- Iliyan Georgiev, Jaroslav Krivánek, Tomáš Davidovič, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 192:1–192:10. <https://doi.org/10.1145/2366145.2366211>
- Israel Gohberg, Seymour Goldberg, and Nahum Krupnik. 2012. *Traces and Determinants of Linear Operators*. Birkhäuser. Google-Books-ID: 6DMDCAAAQBAJ.
- Israel Gohberg and Mark Grigorovich Krein. 1978. *Introduction to the Theory of Linear Nonselfadjoint Operators*. American Mathematical Soc. Google-Books-ID: GOJ-FyTscyrkC.
- N. Halko, P. G. Martinsson, and J. A. Tropp. 2011. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Rev.* 53, 2 (2011), 217–288. <https://doi.org/10.1137/090771806>
- Miloš Hašan, Jaroslav Krivánek, Bruce Walter, and Kavita Bala. 2009. Virtual spherical lights for many-light rendering of glossy scenes. *ACM Transactions on Graphics* 28, 5 (Dec. 2009), 1–6. <https://doi.org/10.1145/1618452.1618489>
- Miloš Hašan, Fabio Pellacini, and Kavita Bala. 2007. Matrix Row-column Sampling for the Many-light Problem. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07, Vol. 26)*. ACM, New York, NY, USA, 26:1–26:10. <https://doi.org/10.1145/1275808.1276410>
- Fu-Chung Huang and Ravi Ramamoorthi. 2010. Sparsely Precomputing The Light Transport Matrix for Real-Time Rendering. *Comput. Graph. Forum* 29 (2010), 1335–1345. <https://doi.org/10.1111/j.1467-8659.2010.01729.x>
- Aref Jeribi. 2021. *Perturbation Theory for Linear Operators: Denseness and Bases with Applications*. Springer, Singapore. <https://doi.org/10.1007/978-981-16-2528-2>
- James T. Kajiya. 1986. The rendering equation. *ACM SIGGRAPH Computer Graphics* 20, 4 (Aug. 1986), 143–150. <https://doi.org/10.1145/15886.15902>
- Gerhard K. Kalisch. 1972. On Operators on Separable Banach Spaces with Arbitrary Prescribed Point Spectrum. *Proc. Amer. Math. Soc.* 34, 1 (1972), 207–208. <https://doi.org/10.2307/2037930> Publisher: American Mathematical Society.
- A. Kalitvin and P.P. Zabrejko. 1991. On the Theory of Partial Integral Operators. *Journal of Integral Equations and Applications* 3 (Sept. 1991), 351–382. <https://doi.org/10.1216/jiea/1181075630>
- Tosio Kato. 1995. *Perturbation Theory for Linear Operators*. Springer Science & Business Media. Google-Books-ID: 8ji2kN_D3BwC.
- Alexander Keller. 1997. Instant radiosity. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 49–56. <https://doi.org/10.1145/258734.258769>
- Mei Kobayashi, Georges Dupret, Oliver King, and Hikaru Samukawa. 2001. Estimation of Singular Values of Very Large Matrices using Random Sampling. *Computer & Mathematics with Applications* 42, 10–11 (2001), 1331–1352.
- Thomas Kollig and Alexander Keller. 2006. Illumination in the Presence of Weak Singularities. In *Monte Carlo and Quasi-Monte Carlo Methods 2004*, Harald Niederreiter and Denis Talay (Eds.). Springer, Berlin, Heidelberg, 245–257. https://doi.org/10.1007/3-540-31186-6_15
- Peter D. Lax. 2002. *Functional Analysis*. Wiley-Interscience. Open Library ID: OL7618836M.
- R. B. Lehoucq, D. C. Sorensen, and C. Yang. 1997. *ARPACK: Solution of Large Scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods*. Available from netlib@ornl.gov.
- Christian Lessig and Eugene Fiume. 2010. On the Effective Dimension of Light Transport. In *Proceedings of the 21st Eurographics Conference on Rendering (EGSR'10)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 1399–1403. <https://doi.org/10.1111/j.1467-8659.2010.01736.x>
- Bradford J. Loos, Lakulish Antani, Kenny Mitchell, Derek Nowrouzezahrai, Wojciech Jarosz, and Peter-Pike Sloan. 2011. Modular Radiance Transfer. In *Proceedings of the 2011 SIGGRAPH Asia Conference (SA '11)*. ACM, New York, NY, USA, 178:1–178:10. <https://doi.org/10.1145/2024156.2024212>
- Manabu Machida. 2014. Singular eigenfunctions for the three-dimensional radiative transport equation. *JOSA A* 31, 1 (Jan. 2014), 67–74. <https://doi.org/10.1364/JOSAA.31.000067>
- Dhruv Mahajan, Ira Kemelmacher Shlizerman, Ravi Ramamoorthi, and Peter Belhumeur. 2007. A Theory of Locally Low Dimensional Light Transport. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07, Vol. 26-3)*. ACM, New York, NY, USA, 62:1–62:10. <https://doi.org/10.1145/1275808.1276454>
- Mohammed Hichem Mortad. 2022. *Counterexamples in Operator Theory*. Birkhäuser. <https://link.springer.com/book/10.1007/978-3-030-97814-3>
- Tian-Tsong Ng, Ramanpreet Singh Pahwa, Jiamin Bai, Kar-Han Tan, and Ravi Ramamoorthi. 2012. From the Rendering Equation to Stratified Light Transport Inversion. *Int. J. Comput. Vision* 96, 2 (Jan. 2012), 235–251. <https://doi.org/10.1007/s11263-011-0467-6>
- Derek Nowrouzezahrai, Patricio Simari, Evangelos Kalogerakis, and Eugene Fiume. 2007. Eigentransport for Efficient and Accurate All-frequency Relighting. In *Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia (GRAPHITE '07)*. ACM, New York, NY, USA, 163–169. <https://doi.org/10.1145/1321261.1321290>
- Jiawei Ou and Fabio Pellacini. 2011. LightSlice: Matrix Slice Sampling for the Many-lights Problem. In *Proceedings of the 2011 SIGGRAPH Asia Conference (SA '11)*. ACM, New York, NY, USA, 179:1–179:8. <https://doi.org/10.1145/2024156.2024213>
- Pieter Peers, Dhruv K. Mahajan, Bruce Lamond, Abhijeet Ghosh, Wojciech Matusik, Ravi Ramamoorthi, and Paul Debevec. 2009. Compressive light transport sensing. *ACM Trans. Graph.* 28, 1 (Feb. 2009), 3:1–3:18. <https://doi.org/10.1145/1477926.1477929>
- Vũ Ngọc Phât and Trinh Cong Dieu. 1994. On the Krein-Rutman theorem and its applications to controllability. In *Proceedings of the American Mathematical Society*, Vol. 120. 495–500. <https://doi.org/10.1090/S0002-9939-1994-1182706-X> ISSN: 0002-9939, 1088-6826 Issue: 2 Journal Abbreviation: Proc. Amer. Math. Soc..
- MICHAEL Reed and BARRY Simon. 1972. *Methods of Modern Mathematical Physics. Vol. I Functional Analysis*. Academic Press. <https://doi.org/10.1016/B978-0-12-585001-8.50001-5>
- J. P. O. Silberstein. 1962. Symmetrisable operators. *Journal of the Australian Mathematical Society* 2, 4 (Dec. 1962), 381–402. <https://doi.org/10.1017/S1446788700027427>
- François Sillion and Claude Puech. 1994. *Radiosity and Global Illumination*. Morgan Kaufmann.
- Françis X. Sillion, James R. Arvo, Stephen H. Westin, and Donald P. Greenberg. 1991. A global illumination solution for general reflectance distributions. In *Proceedings of the 18th annual conference on Computer graphics and interactive techniques (SIGGRAPH '91)*. Association for Computing Machinery, New York, NY, USA, 187–196. <https://doi.org/10.1145/122718.122739>
- Peter-Pike Sloan, Jan Kautz, and John Snyder. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics* 21, 3 (July 2002), 527–536. <https://doi.org/10.1145/566654.566612>
- Peter-Pike Sloan, Ben Luna, and John Snyder. 2005. Local, deformable precomputed radiance transfer. In *ACM SIGGRAPH 2005 Papers (SIGGRAPH '05)*. Association for Computing Machinery, New York, NY, USA, 1216–1224. <https://doi.org/10.1145/1186822.1073335>
- Cyril Soler, Mahdi M. Bagher, and Derek Nowrouzezahrai. 2015. Efficient and Accurate Spherical Kernel Integrals Using Isotropic Decomposition. *ACM Trans. Graph.* 34, 5 (Nov. 2015), 161:1–161:14. <https://doi.org/10.1145/2797136>
- Cyril Soler, Ronak Molazem, and Kartic Subr. 2022. A Theoretical Analysis of Compactness of the Light Transport Operator. In *ACM SIGGRAPH 2022 Conference Proceedings (SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/3528233.3530725>
- Jiguang Sun and Aihui Zhou. 2016. *Finite Element Methods for Eigenvalue Problems*. CRC Press. <https://doi.org/10.1201/9781315372419> Journal Abbreviation: Finite Element Methods for Eigenvalue Problems Publication Title: Finite Element Methods for Eigenvalue Problems.
- Eric Veach. 1997. *Robust Monte-Carlo Methods for Light Transport Simulation*. Ph.D. Dissertation. Stanford University.
- Jiaping Wang, Yue Dong, Xin Tong, Zhouchen Lin, and Baining Guo. 2009. Kernel Nyström method for light transport. In *ACM SIGGRAPH 2009 papers (SIGGRAPH '09)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/1576246.1531335>
- Rui Wang, Jiajun Zhu, and Greg Humphreys. 2007. Precomputed Radiance Transfer for Real-time Indirect Lighting Using a Spectral Mesh Basis. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 13–21. <https://doi.org/10.2312/EGWR/EGSR07/013-021>
- J. H. Wilkinson. 1959. The evaluation of the zeros of ill-conditioned polynomials. Part I. *Numer. Math.* 1, 1 (Dec. 1959), 150–166. <https://doi.org/10.1007/BF01386381>
- Stephen M. Zeman. 2012. *The Classical Theory of Integral Equations: A Concise Treatment* (1 ed.). Birkhäuser Boston.