A Theoretical Analysis of Compactness of The light Transport Operator Supplementary material

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

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, Ronak Molazem, and Kartic Subr. 2022. A Theoretical Analysis of Compactness of The light Transport Operator Supplementary material. In *ACM Transations on Graphics*. ACM, New York, NY, USA, Article 31, 5 pages. https://doi.org/10.1145/3528233.3530725

Here we derive detailed proofs for the results in the main paper. In Section 1 we analyze connections between the spectra of the global ($\mathbf{K}$) and local ($\mathbf{K}_x$) reflectance operators that are defined in the main paper. In Section 2 we prove that despite $\mathbf{T}$ not being compact, it has a converging Schmidt expansion. In Section 3 we prove that $\mathbf{T}_b$ is not compact in scenes with edges. We examine the continuity of the integral operator's kernel κ in Section 4. Finally we prove that $\mathbf{T}_b$ is not invertible and locally smoothes high frequencies.

1 LOCAL VS. GLOBAL REFLECTANCE: $\mathbf{K}$ AND $\mathbf{K}_x$

We recall that the global operator $\mathbf{K}$ is a linear map from $\mathcal{H}$ to itself, the local operator $\mathbf{K}_x$ is a linear map from O to itself. We first prove two preliminary results.

Proposition 1.1: *In a scene with a unique material, the eigenvalues of* $\mathbf{K}$ *have infinite multiplicity.*

For physically-based models[1], the BRDF is a square-integrable function on $\Omega \times \Omega$. Therefore the integral operator $\mathbf{K}_\mathbf{x}$ is Hilbert-Schmidt and thus compact. Due to Helmoltz's reciprocity principle [Veach 1997], $\mathbf{K}_x$ is also Hermitian for the following dot product [Arvo 1996]:

$$\langle f_1, f_2 \rangle_\Omega = \int_\Omega f_1(\omega) f_2(\omega) \cos\theta d\omega.$$

In conclusion, $\mathbf{K}_x$ has countable and real eigenvalues and its eigenfunctions form a complete orthonormal sequence of $(O, \langle ., . \rangle_\Omega)$, which we name $\{r_i\}_{i>0}$, with eigenvalues ρ_i. In the general case, materials produce $\mathbf{K}_\mathbf{x}$ operators with an infinite number of non zero eigenvalues, all of which verify $|\rho_i| < 1$ because of energy conservation. We thus have

$$\forall \mathbf{x} \in S \;\; \forall n > 0 \quad \mathbf{K}_\mathbf{x} r_n = \rho_n r_n$$

Let $\{b_j\}_{j>0}$ be a function basis of $\mathcal{B}$ (we assume that S is bounded). For the Hilbert space $\mathcal{H}$ of square-integrable functions over $S \times \Omega$ with dot product defined using Equation 5 in the main paper, a suitable basis is therefore $\{\psi_n\}_{n>0}$ with

$$\psi_n(\mathbf{x}, \omega) = b_i(\mathbf{x}) r_j(\omega) \text{ with } n = (i, j). \tag{1}$$

Any bijective numbering system that turns the countable set of (i, j) pairs into a single integer n will do, for instance $n = (i+j)(i+j+1)/2 + j$. With this setting, we have

$$\begin{aligned}(\mathbf{K}\psi_n)(\mathbf{x}, \omega) &= \int_\Omega b_i(\mathbf{x}) r_j(\omega') \rho(\mathbf{x}, \omega, \omega') \cos\theta' d\omega' \\ &= b_i(\mathbf{x})(\mathbf{K}_\mathbf{x} r_j)(\omega) \\ &= b_i(\mathbf{x}) \rho_j r_j(\omega) \\ &= \rho_j \psi_n(\mathbf{x}, \omega).\end{aligned}$$

Since there is always an infinite set of pairs (i, j) for any value of j, the above derivation proves that ρ_j is an eigenvalue of $\mathbf{K}$ of infinite multiplicity. Consequently $\mathbf{K}$ is not compact. □

This "undesirable" behavior is explained by the fact that $\mathbf{K}$ is a partial integral operator. In particular $\mathbf{K}$ leaves the spatial component of elements of $\mathcal{H}$ untouched, just like the identity operator does. In other words, $\mathbf{K}$ is the product of the identity along surfaces and the local reflectance operator $\mathbf{K}_x$ at every point $\mathbf{x}$.

Proposition 1.2: *In a scene composed of finite areas with a single material each, the eigenvalues of the global reflectance operator are the union of eigenvalues of all materials, each of which has infinite multiplicity.*

We suppose in this section that the scene is composed of several materials $m_k : \Omega \times \Omega \to \mathbb{R}$, each of which occupies a region S_k of non zero area of the scene. Supposing that (λ, ψ) is an eigenpair of $\mathbf{K}$, we have

$$\forall \mathbf{x} \in S \quad \forall \omega \in \Omega \quad (\mathbf{K}\psi)(\mathbf{x}, \omega) = \lambda \psi(\mathbf{x}, \omega).$$

This equality applies to the entire scene and remains true for each particular surface S_k because $\mathbf{K}$ acts independently on all points of the scene. Consequently the restriction of ψ to S_k is either null, or equal to an eigenfunction of the restriction of $\mathbf{K}$ to S_k. In the latter case, λ is necessarily one of the eigenvalues ρ_{j_0} of m_k, and we have for all points $\mathbf{x} \in S_k$

$$\forall \omega \in \Omega \;\; (\mathbf{K}\psi)(\mathbf{x}, \omega) = \rho_{j_0} \psi(\mathbf{x}, \omega).$$

Let $\{r_j\}_{j>0}$ be the eigenbasis of $\mathbf{K}_\mathbf{x}$ for $\mathbf{x} \in S_k$, with eigenvalues $\{\rho_j\}_{j \geq 0}$. Without loss of generality we assume that these eigenvalues have unit multiplicity. Writing the above equation on the orthogonal basis formed by $\psi_n(\mathbf{x}, \omega) = b_i(\mathbf{x}) r_j(\omega)$ (defined in Equation 1 above), and calling γ_{ij} the coefficients of ψ on this basis leads

[1] Purely specular reflection is usually modeled via a Dirac distribution as a shortcut for a otherwise very sharp reflectance lobe.

Siggraph, 0730-0301/2022/1-ART

to

$$\sum_{i,j} \gamma_{ij}\rho_j b_i r_j = \rho_{j_0} \sum_{i,j} \gamma_{ij} b_i r_j.$$

Equating coefficients on both sides, we get

$$\forall i \;\; \forall j \;\; \rho_j \gamma_{ij} = \rho_{j_0}\gamma_{ij}.$$

When $j \neq j_0$, we have $\gamma_{ij} = 0$. Otherwise the above equation would lead to $\rho_j{=}\rho_{j_0}$ which is not possible. Therefore we have

$$\psi(\mathbf{x},\omega) = r_{j_0}(\omega) \sum_i \gamma_{ij_0} b_i(\mathbf{x}).$$

In other words, ψ is the product of an eigenfunction of m_k and one arbitrary basis element of S_k. The proof generalises to the case where the multiplicity of eigenvalues of m_k is larger than one. □

Note that this situation corresponds to the most common case in lighting simulation: materials span over objects or surfaces of finite non zero areas (even scenes with spatially varying materials can be considered to belong to this class, after splitting the scene to follow spatial variations). More generally, the only possible eigenvalues of **K** are those that are common to materials over entire regions of measurable areas. Consequently in a scene where materials are different at every point, it is most likely that the point-spectrum of **K** is empty.

2 SCHMIDT EXPANSION OF T

We now develop the proof for theorem 3.5 in the main paper.

Proposition 2.1: *In a closed scene with a single non-spatially varying material, the light transport operator can be written as*

$$\mathbf{T}L = \sum_{i,j>0} s_j \rho_j \langle \varphi_n, L\rangle_{\mathcal{H}} \psi_n$$

$$\text{with } n = g(i,j) = j + \frac{(i+j)(i+j+1)}{2},$$

where ψ_n is defined in proposition 1.1 and $\varphi_n = s_j \mathbf{G}\psi_n$ with $s_j = \pm 1$ such that $s_j\rho_j > 0$. Furthermore, both sets $\{\varphi_n\}_{n>0}$ and $\{\psi_n\}_{n>0}$ form complete orthonormal sequences in $\mathcal{H}$.

Let $g : \mathbb{N}\times\mathbb{N} \longrightarrow \mathbb{N}$ be a one-to-one mapping. The choice $g(i,j) = (i+j)(i+j+1)/2 + j$ is such a mapping. For a scene with the same material everywhere, choosing two indices $n_1 = g(i_1, j_1)$ and $n_2 = g(i_2, j_2)$, we have

$$\begin{aligned}\langle \psi_{n_1}, \psi_{n_2}\rangle_{\mathcal{H}} &= \int_S b_{i_1}(\mathbf{x}) b_{i_2}(\mathbf{x}) d\mathbf{x} \int_\Omega r_{j_1}(\omega) r_{j_2}(\omega)\cos\theta d\omega \\ &= \delta_{i_1,i_2}\delta_{j_1,j_2} \\ &= \delta_{n_1,n_2}\end{aligned}$$

Similarly we have

$$\begin{aligned}\langle \varphi_{n_1}, \varphi_{n_2}\rangle_{\mathcal{H}} &= \int_S \int_\Omega (s_{n_1}\mathbf{G}\psi_{n_1})(\mathbf{x},\omega)(s_{n_2}\mathbf{G}\psi_{n_2})(\mathbf{x},\omega)\cos\theta d\omega d\mathbf{x} \\ &= s_{n_1}s_{n_2}\int_S\int_\Omega \psi_{n_1}(\mathbf{y},\omega')\psi_{n_2}(\mathbf{y},\omega')\cos\theta d\omega d\mathbf{x},\end{aligned}$$

with $(\mathbf{y},\omega') = p(\mathbf{x},\omega)$, where p turns a point and direction into the closest visible point and corresponding direction (both directions in the local frame). This change of variables assumes that for any $L \in \mathcal{H}$, $\mathbf{G}L$ is defined everywhere, hence the condition that the scene needs to be closed. Here we have $\|\mathbf{x}-\mathbf{y}\|^2 = d\mathbf{x}\cos\theta/d\omega' = d\mathbf{y}\cos\theta'/d\omega$, which leads to $d\mathbf{x}\cos\theta d\omega = d\mathbf{y}\cos\theta' d\omega'$ and allows the change of variables to integrate along $(\mathbf{y},\omega')$:

$$\begin{aligned}\langle \varphi_{n_1}, \varphi_{n_2}\rangle_{\mathcal{H}} &= s_{n_1}s_{n_2}\int_S\int_\Omega \psi_{n_1}(\mathbf{y},\omega')\psi_{n_2}(\mathbf{y},\omega') d\mathbf{y}\cos\theta' d\omega' \\ &= s_{n_1}s_{n_2}\langle \psi_{n_1},\psi_{n_2}\rangle_{\mathcal{H}} \\ &= \delta_{n_1,n_2}\end{aligned}$$

Since $\{r_i\}_{i>0}$ and $\{b_i\}_{i>0}$ are complete sequences in their respective spaces, their tensor product will be a complete sequence for the set of square-integrable functions defined over the tensor product space $\mathcal{H} = \mathcal{B}\times\mathcal{O}$. In closed scenes, $\mathbf{G}^2 = \mathbf{I}$ which makes $\mathbf{G}$ an isometry [Arvo 1996], so $\{\varphi_n\}_{n>0}$ that is the image of $\{\psi_n\}_{n>0}$ by $\mathbf{G}$ is also complete. Any light distribution L can therefore be expressed as

$$L = \sum_{i,j>0} \langle \varphi_n, L\rangle_{\mathcal{H}}\varphi_n \text{ with } n = g(i,j).$$

Applying $\mathbf{T}$ on both sides preserves convergence of the infinite sum because $\mathbf{T}$ is bounded. Using $\mathbf{G}^2 = \mathbf{I}$ (here again assuming that $\mathbf{G}$ is defined everywhere because the scene is closed),we also have $s_j\mathbf{G}\varphi_n = \psi_n$ and

$$\begin{aligned}\mathbf{T}L &= \sum_{i,j>0} s_j\langle\varphi_n, L\rangle_{\mathcal{H}}\mathbf{KG}\varphi_n \\ &= \sum_{i,j>0} s_j\langle\varphi_n, L\rangle_{\mathcal{H}}\mathbf{K}\psi_n \\ &= \sum_{i,j>0} s_j\rho_j\langle\varphi_n, L\rangle_{\mathcal{H}}\psi_n,\end{aligned}$$

which completes the proof. □

The above proposition generalizes to scenes with multiple materials m_i, each of which occupies a finite region, applying the same re-numbering than in proposition 1.2. It is trivial to verify that in these conditions φ_n (resp. ψ_n) are the eigenfunctions of $\mathbf{T}^*\mathbf{T} = \mathbf{G}^*\mathbf{K}^*\mathbf{KG} = \mathbf{GK}^2\mathbf{G}$ (resp. $\mathbf{TT}^* = \mathbf{KG}^2\mathbf{K} = \mathbf{K}^2$).

We therefore have proved that in closed scenes at least, the light transport operators has a countable set of singular values and that its singular vectors can be built from the eigenvectors of the local reflectance operators, with infinite multiplicity. Since these eigenvalues are not necessarily positive whereas a proper singular value decomposition should have positive singular values, we had to introduce $s_j = \pm 1$ depending on the sign of ρ_j, and multiplying φ_i accordingly.

3 NON-COMPACTNESS OF $\mathbf{T}_b$ AND ABUTTING EDGES

We prove this by explicitly constructing a situation where a bounded sequence of radiant exitance distributions B_n in a way that $\mathbf{T}B_n$ does not admit any converging subsequence. Although we demonstrate this on a specific geometric configuration, this proof generalizes to any scene with unbounded κ, since it exploits an inherent property of the integral operator which preserves high frequencies next to the poles of κ. The configuration we choose however enables a simple exposition.

We consider a scene with two orthogonally abutting diffuse half-planes with a point x_0 on the common edge and albedo ρ. We define S_0 to be a region on the first square with unit measure. We construct

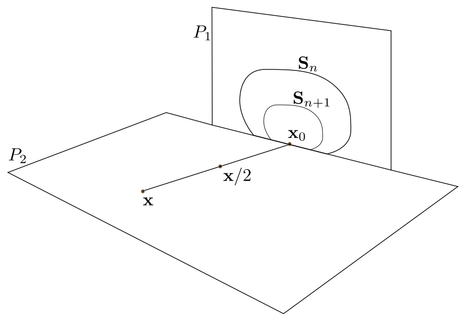

Fig. 1. Notations for the analysis of $\mathbf{T}_b$.

S_n to be scaled image of S_0 by a factor of 2^{-n} using $\mathbf{x}_0$ as the common origin on each plane (See Figure 1). Let $s_n = 1/4^n$ denote the area of S_n. We define a sequence of radiant exitance functions B_n as

$$B_n(\mathbf{y}) = 2^n \mathbf{1}_{S_n}(\mathbf{y}). \quad (2)$$

Each distribution B_n is of unit norm since

$$\|B_n\|_2^2 = \int_S B_n(\mathbf{y})^2 d\mathbf{y} = 4^n s_n = 1$$

For any point $\mathbf{x} \in P_2$ let $F_n(\mathbf{x})$ be the form factor from $\mathbf{x}$ to S_n. Using Equation 2, the transported radiant exitance $\mathbf{T}B_n(\mathbf{x})$ is

$$\mathbf{T}B_n(\mathbf{x}) = \rho 2^n F_n(\mathbf{x}).$$

Since scaling preserves form factors, we have $F_n(\mathbf{x}) = F_{n-1}(2\mathbf{x})$. Therefore

$$\begin{aligned}\mathbf{T}B_n(\mathbf{x}) &= \rho 2^n F_n(\mathbf{x}) \\ &= \rho 2^n F_{n-1}(2\mathbf{x}) \\ &= 2\mathbf{T}B_{n-1}(2\mathbf{x}) \\ &= 2^n \mathbf{T}B_0(2^n \mathbf{x}).\end{aligned}$$

Now we compute the distance between pairs of elements in the sequence $\mathbf{T}B_n$. Let n and p be two non negative integers:

$$\begin{aligned}\|\mathbf{T}B_n - \mathbf{T}B_{n+p}\|_2^2 &= \int_S \left|\mathbf{T}B_n(\mathbf{x}) - \mathbf{T}B_{n+p}(\mathbf{x})\right|^2 d\mathbf{x} \\ &= \int_S |2^n \mathbf{T}B_0(2^n\mathbf{x}) - 2^n \mathbf{T}B_p(2^n\mathbf{x})|^2 d\mathbf{x}\end{aligned}$$

With the change of variables $\mathbf{y} = 2^n\mathbf{x}$, we have $d\mathbf{y} = (2^n)^2 d\mathbf{x} = 4^n d\mathbf{x}$ (since those are surface elements), this yields

$$\|\mathbf{T}B_n - \mathbf{T}B_{n+p}\|_2^2 = \int_S |\mathbf{T}B_0(\mathbf{y}) - \mathbf{T}B_p(\mathbf{y})| d\mathbf{y},$$

which does not depend on n. Let's call r_p that distance:

$$r_p = \|\mathbf{T}B_n - \mathbf{T}B_{n+p}\|_2,$$

and estimate the limit of r_p when p tends to infinity. We know that

$$\begin{aligned}\mathbf{T}B_p(\mathbf{y}) &= 2^p \ \mathbf{T}B_0(2^p \mathbf{y}) \\ &= \rho \ 2^p \ F_0(2^p \mathbf{y}).\end{aligned}$$

As p tends to infinity, the point $2^p\mathbf{y}$ moves arbitrarily far away from the origin, and the form factor $F_0(2^p\mathbf{y})$ is equivalent to s_0 times the point-to-point form factor between $2^p\mathbf{y}$ and the center of S_0. Because the center of scaling is on the edge, the cosine $\cos\theta$ with the normal at $\mathbf{y}$ involved in the form factor tends to 0, $\cos\theta'$ with the normal at the center of S_0 tends to 1:

$$\begin{aligned}\mathbf{T}B_p(\mathbf{y}) &\sim \rho 2^p \frac{\cos\theta\cos\theta'}{\|2^p\mathbf{y}\|^2} s_0 \\ &\sim \rho \frac{\cos\theta}{2^p\|\mathbf{y}\|^2}\end{aligned}$$

which converges to 0 as $p \to \infty$. We deduce that

$$\lim_{p\to\infty} r_p = \|\mathbf{T}B_0\|_2,$$

which is a non negative number. If P_2 was infinite, it would be possible to compute this value analytically, but we do not explicitly need to do so.

Now we show that the sequence $\{\mathbf{T}B_n\}_{n>0}$ has no converging subsequence. Let q be a strictly increasing mapping of $\mathbb{N}$ to itself, such that $u_n = \mathbf{T}B_{q(n)}$ is a supposedly converging subsequence. Then such a subsequence is also a Cauchy sequence and therefore:

$$\forall\epsilon > 0 \ \ \exists n_\epsilon \in \mathbb{N} \ \ \forall n > n_\epsilon \ \ \forall p > 0 \quad \|u_n - u_{n+p}\|_2 < \epsilon. \quad (3)$$

However, we have

$$\|u_n - u_{n+p}\|_2 = \|\mathbf{T}B_{q(n)} - \mathbf{T}B_{q(n+p)}\|_2 = r_{q(n+p)-q(n)}.$$

Since $\lim_{p\to\infty} r_p > 0$, and $r_p = 0$ only for $p = 0$, r_p has a non negative minimum value v over $\mathbb{N} - \{0\}$. Because q is strictly increasing we know that $q(n+p) - q(n) > 0$ and therefore $r_{q(n+p)-q(n)} > v$ for all $p > 0$. This contradicts equation 3. We deduce that the sequence $\mathbf{T}B_n$ has no converging subsequence. □

When the two surfaces P_1 and P_2 are not orthogonal, the proof still stands, since $\cos\theta'$ is bounded whatsoever and $\cos\theta$ still converges to 0. All other derivations and arguments remain valid.

This proof leverage the fact that in the presence of an edge, $\mathbf{T}_b$ preserves high frequencies. It therefore locally acts very similarly to an infinite dimensional isometry and is therefore not compact. The proof actually works for any L_p norm with $p \geq 1$. We presented it using the L_2 norm to keep consistent with the choices in the main paper that make $\mathcal{B}$ be a Hilbert space.

4 CONTINUITY OF $\kappa(\mathbf{x}, \mathbf{y})$ AT $\mathbf{x} = \mathbf{y}$

In this section we prove that κ is bounded on any curved/flat surface (more specifically the surface needs to be of class G_2). It is however continous in only one very specific situation: flat surfaces and portions of spheres. In the general case no value can be attributed to $\kappa(\mathbf{x}, \mathbf{x})$.

In order to find the value in the limit, we consider the case of a point $\mathbf{x}$ on a surface that of class G_2 at this point, and compute a series expansion of $\kappa(\mathbf{x}, \mathbf{y})$ at a nearby point $\mathbf{y}$. Without loss of generality, we use 2D coordinates to represent both points using an order 2 differentiable parametrization of the surface in the neighborhood of $\mathbf{x}$, along the tangent plane at $\mathbf{x}$ to the surface of the scene. We call h the distance between $\mathbf{x}$ and the projection of $\mathbf{y}$ on this tangent plane. The figure below shows a cross-section of the surface containing both points (their normals are in the same plane, since we consider the limit as $\mathbf{x}$ and $\mathbf{y}$ are arbitrarily close to each other):

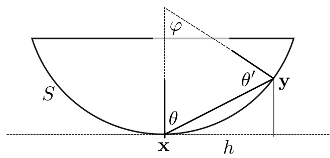

With this setting the coordinates of $\mathbf{x}$ are $(0,0)$, and the coordinates of $\mathbf{y}$ are $(h, \alpha_1 h + \frac{\alpha_2}{2}h^2 + O(h^3))$. Because the normal is vertical at $\mathbf{x}$, we have $\alpha_1 = 0$. The value $\alpha_2 = 0$ corresponds to a locally flat surface (or at least locally flat surfaces fall in this case). Coefficient α_2 is actually equal to the surface linear curvature at $\mathbf{x}$. When $\alpha_2 \neq 0$, the considered section of S is locally equivalent to a circle. We have therefore $\theta = \theta'$ and

$$\cos\theta = \cos\theta' = \frac{\frac{1}{2}\alpha_2 h^2 + O(h^3)}{\sqrt{h^2 + (\frac{1}{2}\alpha_2 h^2 + O(h^3))^2}}.$$

So

$$\begin{aligned}\kappa(\mathbf{x},\mathbf{y}) &= \rho(x)\frac{\cos\theta\cos\theta'}{\|\mathbf{x}-\mathbf{y}\|^2}\\ &= \rho(\mathbf{x})\frac{(\cos\theta)^2}{\|\mathbf{x}-\mathbf{y}\|^2}\\ &= \rho(\mathbf{x})\frac{(\frac{1}{2}\alpha_2 h^2 + O(h^3))^2}{\left(h^2 + (\frac{1}{2}\alpha_2 h^2 + O(h^3))^2\right)^2}.\end{aligned}$$

Assuming that $\alpha_2 \neq 0$, the limit of the kernel can therefore be studied when h tends to 0. We have

$$\lim_{h\to 0}\kappa(\mathbf{x},\mathbf{y}) = \rho(\mathbf{x})\frac{1}{4}\alpha_2^2$$

If the surface around $\mathbf{x}$ is differentiable, the curvature α_2 of the cross section will depend on the direction $\mathbf{v}$ of that cross-section, involving the shape operator $\nabla\mathbf{n}(\mathbf{x})$ of the surface at $\mathbf{x}$. The projection of $\mathbf{y}$ onto the tangent plane is now a 2D vector $\mathbf{h}$ that we write in this plane using polar coordinates as $\mathbf{h} = [h\cos\beta, h\sin\beta]$. Involving the two principal curvatures k_1 and k_2, we obtain

$$\begin{aligned}\lim_{h\to 0}\kappa(\mathbf{x},\mathbf{y}+\mathbf{h}) &= \frac{\rho(\mathbf{x})}{4}\frac{\mathbf{h}^\top\nabla\mathbf{n}(\mathbf{x})\mathbf{h}}{\|\mathbf{h}\|^2}\\ &= \frac{\rho(\mathbf{x})}{4}(k_1\cos^2\beta + k_2\sin^2\beta)^2.\end{aligned}$$

This expression obviously depends on the direction β at which $\mathbf{y}$ approaches $\mathbf{x}$, which proves that even if κ is bounded, it has no limit for $\mathbf{x} = \mathbf{y}$ when the two principal curvatures are not equal.

In summary, this completes the proof that: (1) On differentiable surfaces the kernel of the light transport operator is bounded, but generally not continuous; and (2) it is continous whenever the two principal curvatures of the surface are equal.

Specific case of the interior of a sphere. When $k_1(\mathbf{x}) = k_2(\mathbf{x}) = k(\mathbf{x})$, κ becomes a continuous function at $\mathbf{y} = \mathbf{x}$ which makes $\mathbf{T}_b$ a trace-class operator in this case. In particular the trace of $\mathbf{T}$ can be computed by applying the theorem for trace-class operators with continous kernels [Conway 1990]:

$$\begin{aligned}\mathrm{Tr}(\mathbf{T}_b) &= \int_S \kappa(\mathbf{x},\mathbf{x})d\mathbf{x}\\ &= \int_S \frac{1}{4}\rho(\mathbf{x})k(\mathbf{x})^2 d\mathbf{x}\end{aligned}$$

Inside a sphere of radius r, the two principal curvatures are constant and equal to $1/r$. Applying the above formula we get

$$\mathrm{Tr}(\mathbf{T}_b) = \int_S \frac{1}{4r^2}\rho(\mathbf{x})d\mathbf{x} = \overline{\rho},$$

where $\overline{\rho}$ is the average albedo inside the sphere.

In this very specific geometric configuration however, solving for the eigenvalues of the light transport operator is straightforward: First, for any $\mathbf{x}$ and $\mathbf{y}$,

$$\begin{aligned}\kappa(\mathbf{x},\mathbf{y}) &= \rho(\mathbf{x})\frac{\cos\theta\cos\theta'}{\pi\|x-y\|^2}\\ &= \frac{1}{\pi}\rho(\mathbf{x})\frac{(\cos\theta)^2}{2r^2(1-\cos\varphi)}\\ &= \frac{1}{\pi}\rho(\mathbf{x})\frac{(\sin\frac{\varphi}{2})^2}{2r^2 2\sin(\frac{\varphi}{2})^2}\\ &= \frac{\rho(\mathbf{x})}{4\pi r^2}.\end{aligned}$$

So when applying $\mathbf{T}_b$ to a function B in the sphere, we get

$$\begin{aligned}(\mathbf{T}_b B)(\mathbf{x}) &= \int_S \kappa(\mathbf{x},\mathbf{y})B(\mathbf{y})d\mathbf{y}\\ &= \rho(\mathbf{x})\frac{1}{4\pi r^2}\int_S B(\mathbf{x})d\mathbf{x},\end{aligned}$$

meaning that an eigenfunction can only be a multiple of ρ. This means that the only eigenspace is of dimension 1, and assuming eigenvectors are normalized, the only eigenfunction of $\mathbf{T}_b$ is in this case

$$\Lambda(\mathbf{x}) = \frac{1}{\|\rho\|}\rho(\mathbf{x}).$$

Then, solving for the corresponding eigenvalue λ, we get

$$\begin{aligned}\lambda\rho(\mathbf{x})\frac{1}{\|\rho\|} &= \rho(\mathbf{x})\frac{1}{4\pi r^2}\int_S \frac{\rho(\mathbf{x})}{\|\rho\|}d\mathbf{x}\\ \lambda &= \frac{1}{4\pi r^2}\int_S \rho(\mathbf{x})d\mathbf{x}\\ &= \overline{\rho}\end{aligned}$$

It is interesting to see the effect of discretization on such a scene. We subdivide a tetraheron p times, to produce a tesselated approximation of a sphere with 4^{p+1} surface elements. We set the albedo to be 0.8 on each surface element. The trace of the discretized operator is always zero since self-form-factors for flat surface elements are zero. However, it should intuitively converge to 0.8!

We tabulate the largest eigenvalues of the matrix, computed using arpack++, as a function of the number of subdivision steps:

	2 (64)	3 (256)	4 (1024)	5 (4096)	6 (16384)
λ_0	0.7985581	0.7996539	0.7999077	0.7999821	0.7999841
λ_1	-0.1176031	-0.0502393	-0.0220877	-0.0101415	-0.0048050
λ_2	-0.1175783	-0.0502362	-0.0220863	-0.0101414	-0.0048050
λ_3	-0.1175607	0.0253634	0.0140502	-0.0101412	0.0041005
λ_4	0.5862844	0.0253376	0.0140497	0.0077256	0.0041004
λ_5	0.0586174	-0.0238030	0.0140491	0.0077254	0.0041004
λ_6	-0.0477238	-0.0238003	-0.0112667	-0.0052971	-0.0025860
λ_7	-0.0476994	-0.0231174	-0.0112622	-0.0052969	-0.0025860
λ_8	0.0473600	-0.0231033	-0.0107773	-0.0052968	-0.0025860
λ_9	-0.0447665	-0.0215689	-0.0107756	-0.0052241	-0.0024441

As expected for a light transport matrix, the diagonal is 0 and therefore the trace of the matrix is always 0. This is counterintuitive because eigenvalues of the matrix are supposed to converge to the actual eigenvalues of $\mathbf{T}$ [Chatelin 1981] although the trace apparently does not. This is possible because the number of eigenvalues increases with the size of the matrix and although their sum compensates the first eigenvalue $\lambda_0 = 0.8$ for a null trace, each of them still converges to 0 when the number of discretization elements grows to infinity.

5 PROOF THAT $\mathbf{T}_b$ IS NOT INVERTIBLE

Proposition 5.1: *Over open subsets of S where $\mathbf{y} \mapsto \kappa(\mathbf{x}, \mathbf{y})$ is bounded, the values of $\mathbf{T}_b$ coincide with the values of a compact operator.*

Suppose that such a subset $A \subset S$ exists, then

$$\exists M > 0 \;\; \forall \mathbf{x} \in A \;\; \forall \mathbf{y} \in S \;\; \kappa(\mathbf{x}, \mathbf{y}) < M. \tag{4}$$

We define the following integral operator between radiant exitance distributions over S and A (denoting the later set $\mathcal{B}_A$):

$$\mathbf{T}'_b : \mathcal{B} \rightarrow \mathcal{B}_A$$
$$\mathbf{T}'_b(L)(\mathbf{x}) = \int_S \kappa(\mathbf{x}, \mathbf{y}) L(\mathbf{y}) d\mathbf{y}.$$

The image of this operator trivially coincides with the image of $\mathbf{T}_b$ over A since

$$\forall \mathbf{x} \in A \;\; \forall L \in \mathcal{B} \;\; \mathbf{T}'_b(L)(\mathbf{x}) = \mathbf{T}_b(L)(\mathbf{x}).$$

Following Equation 4, $\mathbf{T}'_b$ is an integral operator with bounded kernel over $A \times S$. $\mathbf{T}'_b$ is therefore a Hilbert-Schmidt operator, and consequently compact. □

In every possible scene κ is locally bounded for at least one subset $A \subset S$ as above. Indeed, any flat or convex finite region of the surfaces in the scene will stand for a suitable A. In other words, a scene cannot have concave abutting edges at every possible points. This result has two direct consequences (propositions 5.2 and 6.2).

Proposition 5.2: *$\mathbf{T}_b$ is not invertible.*

Suppose that we found an $A \subset S$ suitable for proposition 5.1. Considering that $\mathbf{T}'_b$ is compact, it is not invertible. Consequently there exists $L'_0 \in \mathcal{B}_A$ such that the equation $\mathbf{T}'_b X = L'_0$ has no solution in S. Now we arbitrarily construct $L_0 \in \mathcal{B}$ such that $L_0(\mathbf{x}) = L'_0(\mathbf{x})$ for any point $\mathbf{x} \in A$. If there is an $X_0 \in \mathcal{B}$ such that $\mathbf{T}_b X_0 = L_0$, then

$$\forall \mathbf{x} \in A \;\; (\mathbf{T}_b X_0)(\mathbf{x}) = (\mathbf{T}'_b X_0)(\mathbf{x}) = L_0(\mathbf{x}) = L'_0(\mathbf{x}).$$

This contradicts the construction of L'_0 since it would yield a solution $X = X_0$ to the inverse of L'_0 by $\mathbf{T}'_b$. □

6 PROOF THAT $\mathbf{T}_b$ FILTERS OUT HIGH FREQUENCIES

An interesting consequence of the image $\mathbf{T}_b$ coinciding with the image of a compact operator almost everywhere is that $\mathbf{T}_b$ cancels out high frequencies almost everywhere. We prove this in two steps.

Proposition 6.1: *Compact operators attenuate high frequencies.*

Let $\mathbf{B}$ be a compact operator from a Hilbert space $\mathcal{H}$ to an Hilbert space $\mathcal{I}$. Let $\{b_i\}_{i>0}$ be a complete orthonormal sequence of $\mathcal{H}$. We assume without loss of generality that the elements b_i are sorted by increasing frequency content. For any $n \in \mathbb{N}$ we call $\mathbf{P}_n : \mathcal{H} \rightarrow \mathcal{H}$ the projection on the finite dimensional space spanned by $\{b_i\}_{i<n}$, and we define the finite rank operator $\mathbf{B}_n$ to be

$$\mathbf{B}_n = \mathbf{B}\mathbf{P}_n.$$

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

$$\forall \epsilon > 0 \;\; \exists p \in \mathbb{N} \;\; \forall n \geq p \;\; \|\mathbf{B}_n - \mathbf{B}\| < \epsilon,$$

which means

$$\forall \epsilon > 0 \;\; \exists p \in \mathbb{N} \;\; \forall n \geq p \;\; \forall L \in \mathcal{B} \;\; \|(\mathbf{B}_n - \mathbf{B})L\| < \epsilon \|L\|.$$

We now restrict the above equation to $L = b_{n+1}$ for every possible $n > p$. We have $(\mathbf{B}_n - \mathbf{B})b_{n+1} = \mathbf{B}b_{n+1}$ and $\|b_{n+1}\| = 1$, which gives after shifting $n+1$ to n:

$$\forall \epsilon > 0 \;\; \exists p \in \mathbb{N} \;\; \forall n \geq p \;\; \|\mathbf{B}b_n\| < \epsilon.$$

Thus, $\mathbf{B}$ washes out high frequencies. □

Note that the reverse is not true: an operator that attenuate high frequencies may still be non compact, if its point-spectrum is decreasing but only converging to some small value that is strictly larger than 0.

Proposition 6.2: *In regions where κ is bounded, $\mathbf{T}_b$ attenuates high frequencies.*

Given $A \subset S$ that is suitable for proposition 5.1, we know that $\mathbf{T}'_b$ is compact. This operator therefore fits proposition 6.1. Since its output coincides with $\mathbf{T}_b$ over A, $\mathbf{T}_b$ washes out high frequencies over A. This proposition mathematically explains a phenomenon that is well known within the rendering community for many years.

REFERENCES

James Arvo. 1996. *Analytic Methods for Simulated Light Transport.* Ph. D. Dissertation. Yale University.

F. Chatelin. 1981. The Spectral Approximation of Linear Operators with Applications to the Computation of Eigenelements of Differential and Integral Operators. *SIAM Rev.* 23, 4 (Oct. 1981), 495–522. https://doi.org/10.1137/1023099

John B. Conway. 1990. *A Course in Functional Analysis* (2 ed.). Springer-Verlag, New York. https://www.springer.com/gp/book/9780387972459

Eric Veach. 1997. *Robust Monte-Carlo Methods for Light Transport Simulation.* Ph. D. Dissertation. Stanford University.