Geometric Information and Approximation in Recessed-Aperture Transport

Mostafa Nasr
Independent researcher, Giza, Egypt
mostafa.a.nasr@gmail.com

Working revision 14 for author review, 21 November 2026
Published parent (revision 13) Public version history

Abstract

How much of a recessed aperture must be measured to predict the power transmitted by a finite light source? We study direct rays in channels with absorbing walls. Two smooth circular tapers with identical end openings, depth, enclosed volume and on-axis transmission can transmit 1/36 and zero of the entering power over the same oblique source. For a channel with coaxial circular sections, the greatest convex function below its radius profile admits exactly the same rays. With a known curvature bound, uncertain radius measurements at a finite set of depths yield inner and outer channels that enclose transmission even when only part of the exit is visible. For three broad profiles under a specified annular source, five stations certify a one-percentage-point midpoint error with the computed interval, while a source-independent spacing rule gives nine. A narrower horn requires seventeen among the tested grids for either guarantee. For prismatic apertures we retain finite-source error floors for omitted aspect ratio and skew, certificates for rounding an affine shape record, and numerical checks showing when quadrature or angular resolution limits the prediction. The results concern mathematical and computational verification. Physical validation remains open.

1 Scientific question and contribution

Suppose a bore gauge gives the radii of the two openings and a few intermediate sections of a channel. A source illuminates the entrance obliquely. Can those readings establish how much of the entering power crosses the exit? A localized narrowing can screen the whole beam. A larger chamber elsewhere can restore the same total volume while leaving that obstruction in place. The question is which geometric measurements suffice for a specified prediction and accuracy.

We distinguish three mechanisms. *Information loss* gives the same recorded inputs to shapes that require different answers. *Representation error* arises when an instrument or a finite record approximates a sufficient geometry. *Evaluation error* arises when integrating the resulting response or distributing power over angular cells. A useful error budget keeps these contributions distinguishable. Physical testing additionally has to establish the optical properties of the wall, the incident source and the receiver.

We give both a failure and a sufficient measurement prescription. For smooth circular tapers, Section 7.4 constructs two volume-matched channels with different off-axis transmission and proves which outward wall features can be discarded exactly. Section 7.5 bounds partially transmitted power from measured radii and a bound on the curvature of the radius profile. Section 7.6 evaluates the bounds on four specified channels under finite illumination. For affine and polygonal channels, the earlier sections give error floors for omitted information and certificates for shape rounding. The quadrature, restricted-view and inverse studies then show how computational errors or an underdetermined observation can remain even when the geometry is specified.

Directional descriptions already underpin daylight simulation, fenestration and precomputed radiance transfer [12, 5, 7]. A system may retain shape in vertices, dimensions, a kernel or a precomputed operator. Those quantities count as information in our test; the deliberately reduced

records do not describe the complete inputs of every deployed system. The overlap model and diffuse form factors also have established precedents [3, 4]. Our results are bounded comparisons and certificates within that model, not a new universal transport law or solver.

2 Model and observable

2.1 Geometry and assumptions

The initial model is a channel with a constant parallelogram section. Section 7 extends the geometry to polygonal mouths, internal stops and axially varying circular sections. Each extension uses the same entrance-relative flux normalization. Let the rank-two matrix $B = [a\ b]$ contain the full edge vectors of a parallelogram mouth, with

$$A = \|a \times b\|_2, \qquad n = \frac{a \times b}{A}, \qquad U = [-1/2, 1/2]^2. \tag{1}$$

The inner mouth is BU; the outer mouth is its translation by v. Planar walls join corresponding edges. The geometry is $g = (B, v)$, with normal separation $d = n \cdot v > 0$. An oblique channel may have tangential offset. The flat case is defined separately by $v = 0$.

The ordered orthonormal frame and its shape coordinates are

$$x = a/\|a\|_2, \quad y = \frac{b - x(x \cdot b)}{\|b - x(x \cdot b)\|_2}, \quad Q = [x\ y\ n], \tag{2}$$

$$B = [x\ y] \begin{pmatrix} a_0 & c_0 \\ 0 & b_0 \end{pmatrix}, \quad a_0, b_0 > 0,\ c_0 \in \mathbb{R}, \quad A = a_0 b_0. \tag{3}$$

At fixed area and frame, aspect ratio and skew remain free. The walls are opaque and absorbing; there is no glazing, reflection, refraction, diffraction or participating medium. Incident radiance is uniform across the mouth for each direction. All lengths share a unit system. Uniform scaling preserves normalized transmission and scales power quadratically.

2.2 Directional overlap and transmitted power

The unit vector s points toward the source, with $\mu = n \cdot s > 0$. The displacement relative to the outer mouth and its edge coordinates are

$$\delta = \frac{d}{\mu} s - v, \qquad q = B^+ \delta, \qquad B^+ = (B^\mathsf{T} B)^{-1} B^\mathsf{T}. \tag{4}$$

A ray from Bu clears the convex channel exactly when $u, u + q \in U$. The visible mouth fraction is therefore

$$V_g(s) = [1 - |q_1|]_+ [1 - |q_2|]_+, \qquad [z]_+ = \max(z, 0). \tag{5}$$

Each factor is an interval-overlap fraction. Flat openings have $V = 1$ on the front hemisphere; rear and tangent directions supply no one-sided power.

For nonnegative angular radiance L, define

$$H_g(s) = A[n \cdot s]_+ V_g(s), \qquad \mathcal{R}(g, L) = \int_{S^2} L(s) H_g(s)\, \mathrm{d}\Omega, \tag{6}$$

$$P_0(g, L) = A \int_{S^2} L(s)[n \cdot s]_+\, \mathrm{d}\Omega, \qquad T(g, L) = \mathcal{R}(g, L)/P_0(g, L), \quad P_0 > 0. \tag{7}$$

Thus $T \in [0, 1]$ is the incident-flux-weighted average of V, reducing to $V(s)$ for a collimated direction. The target $|\widehat{T} - T| \le 0.01$ means one percentage point of unblocked power, not one percent relative error.

With radiometric units $\mathcal{R}$ is radiant power; consistent photometric weighting gives luminous flux. A small detector reading is not automatically whole-exit flux: collection and spatial integration must match the observable. Outgoing radiance, glare and a spatial receiver require additional information.

2.3 A geometry-dependent integration chart

Set $r(q) = v + Bq$ and $s(q) = r/\|r\|_2$. Since $n \cdot r = d$, the displacement support $[-1, 1]^2$ maps to front-facing directions with

$$\mathrm{d}\Omega = \frac{|r \cdot (a \times b)|}{\|r\|_2^3}\,\mathrm{d}q_1\,\mathrm{d}q_2 = \frac{Ad}{\|r\|_2^3}\,\mathrm{d}q_1\,\mathrm{d}q_2. \tag{8}$$

Consequently the same power is

$$\mathcal{R} = A^2 d^2 \int_{[-1,1]^2} \frac{L(s(q))(1-|q_1|)(1-|q_2|)}{\|v + Bq\|_2^4}\,\mathrm{d}q_1\,\mathrm{d}q_2. \tag{9}$$

This change of coordinates is exact for $d > 0$; the flat case is handled directly. Quadrant splitting resolves the tent creases, but not necessarily a source boundary or a narrow denominator peak. Sections 5.5 and 8 show why an easier integration boundary can leave a difficult integrand.

3 When is a geometric summary sufficient?

3.1 The same-input test

A descriptor $D(g)$ is the information about geometry supplied to a calculation. It might contain dimensions, an orientation frame, a lookup table or the coefficients of a response function. Calling it sufficient makes a claim about a *specified task*, not about geometry in isolation.

Definition 3.1 (Sufficiency for the stated response)**.** *For declared geometry and illumination families, D is sufficient if one calculation F satisfies*

$$\mathcal{R}(g, L) = F(D(g), L; \textit{fixed model data}) \tag{10}$$

for every allowed g and L. A shape-specific kernel, fitted parameter or family label used by F is part of its input, not fixed data shared by all shapes.

Because both summaries retain A, n, they determine P_0; their sufficiency for power and transmission is equivalent when $P_0 > 0$.

Proposition 3.2 (Identical inputs must require identical answers)**.** *A sufficient calculation exists exactly when*

$$D(g_1) = D(g_2) \quad \Longrightarrow \quad \mathcal{R}(g_1, L) = \mathcal{R}(g_2, L) \quad \textit{for every allowed } L. \tag{11}$$

A calculation given the same inputs must return the same answer. Conversely, if every geometry sharing a descriptor requires the same answer, that answer defines F. When all bounded non-negative angular fields are allowed, the criterion is equivalent to equal response kernels almost everywhere. Appendix A gives the formal argument and the corresponding restriction to a smaller lighting family.

This is deterministic sufficiency for a task, not statistical sufficiency or computational efficiency.

For template reuse, rotating a response is valid for a rigid copy. If deformation changes aspect ratio or skew while only Q, A, d are retained, the old response cannot in general describe the new

opening. Loss occurs at that storage boundary, not when computing Q from otherwise retained geometry.

The dual question fixes the geometry and reduces the lighting input. If $D_L(L_1) = D_L(L_2)$, the same argument gives a worst-error floor of $\frac{1}{2}|T(g, L_1) - T(g, L_2)|$ for every calculation given only g and D_L. In Section 6, D_L consists of all 145 incident patch fluxes.

3.2 A square and a rectangle that the summary cannot distinguish

Start with a square of side one and a rectangle of width two and height one half, both with depth one and normal extrusion. Let $\mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z$ denote Cartesian unit vectors and I the identity matrix. In common length units,

$$B_{\text{sq}} = [\mathbf{e}_x, \mathbf{e}_y], \qquad B_{\text{rect}} = [2\mathbf{e}_x, \tfrac{1}{2}\mathbf{e}_y], \qquad v = \mathbf{e}_z. \tag{12}$$

Both have area one, normal $\mathbf{e}_z$, depth one and frame $Q = I$. Thus both proposed summaries, (n, A, d) and (Q, A, d), are identical.

Illuminate them from $s_* = (1, 0, 2)/\sqrt{5}$. Across the depth, this direction shifts a ray sideways by one half of a length unit, with no vertical shift. That removes half the square's usable width, but only a quarter of the rectangle's. Equation (5) gives

$$T_{\text{sq}} = \frac{1}{2}, \qquad T_{\text{rect}} = \frac{3}{4}. \tag{13}$$

Figure 1 identifies the omitted width.

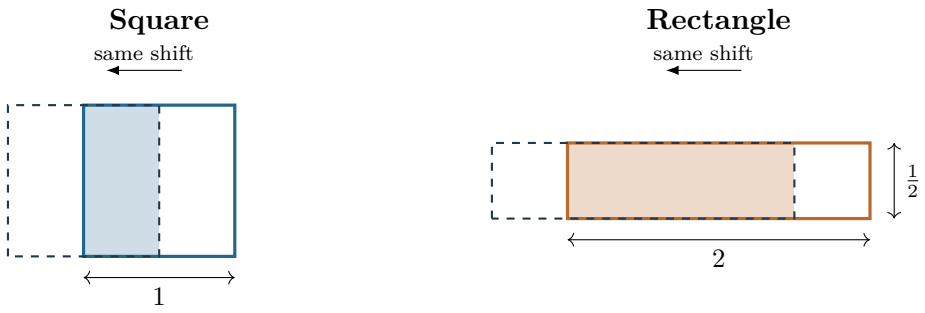


Figure 1: Why equal area does not imply equal transmission. Solid outlines are the inner mouths; dashed outlines are the projections of the outer mouths along the same ray direction. Shading marks the starting positions from which rays pass through. Both mouths have unit area and depth, and both projections shift by one half of a length unit. This is a geometric schematic of the directional calculation, not a measurement. Figure code was prepared with AI assistance.

To establish separation for a finite source, use a uniform circular cone of half-angle $\alpha = 2°$ about s_*, normalized to unit angular integral:

$$\mathcal{C}_\alpha = \{s : s \cdot s_* \geq \cos\alpha\}, \qquad L_\alpha(s) = \frac{\mathbf{1}_{\mathcal{C}_\alpha}(s)}{2\pi(1 - \cos\alpha)}. \tag{14}$$

The source is bounded and uniform across the mouth. Appendix B proves

$$T_{\text{rect}} - T_{\text{sq}} > \frac{5}{32} = 0.15625. \tag{15}$$

The exact expression for the gap is

$$T_{\text{rect}} - T_{\text{sq}} = \frac{1}{4} - \frac{2\alpha - \sin(2\alpha)}{(2/\sqrt{5})\pi \sin^2\alpha} \simeq 0.23343384, \tag{16}$$

with α in radians. Numerical integration estimates the individual transmissions as about 0.49172 and 0.72515. The analytic inequality, rather than confidence in those numerical digits, establishes the finite-source separation.

3.3 How much error does the missing information force?

Suppose two true transmissions are $r_1 < r_2$, but a common record forces the prediction to be one number y. Even the best choice of y cannot be closer than half the gap to both:

$$\max_i |y - r_i| \geq \frac{r_2 - r_1}{2}. \tag{17}$$

The midpoint attains this bound. For positive responses, the sharp relative-error bound is instead

$$\max_i \frac{|y - r_i|}{r_i} \geq \frac{r_2 - r_1}{r_1 + r_2}. \tag{18}$$

Its optimum is the harmonic mean $2r_1r_2/(r_1 + r_2)$; the two error criteria should not be interchanged.

For the square–rectangle cone, Eq. (15) forces an absolute error greater than 5/64, or 7.8125 percentage points, on at least one opening. A one-percentage-point target is therefore impossible for a predictor restricted to either common reduced record. The conclusion is independent of its fitting method or numerical sophistication.

3.4 Shear gives a second example

Aspect ratio is not the only missing information: an area-preserving shear also changes transmission. Let ℓ be a reference length and set

$$B_k = \ell \begin{bmatrix} 1 & k \\ 0 & 1 \\ 0 & 0 \end{bmatrix}, \qquad v = \ell \mathbf{e}_z, \qquad k \in \{0, 3/10\}. \tag{19}$$

The second mouth is a sheared square. Both have $A = \ell^2$, $d = \ell$, $n = \mathbf{e}_z$ and the same ordered frame $Q = I$. For a direction proportional to $(0.8, 0.4, 1)$, the edge-coordinate displacements are $(0.8, 0.4)$ and $(0.68, 0.4)$. Hence

$$V_0 = 0.20 \times 0.60 = 0.120, \qquad V_{0.3} = 0.32 \times 0.60 = 0.192. \tag{20}$$

The directional absolute-error floor is 0.036; the relative floor is 3/13, about 23.08%.

A uniform 2° cone about this direction also preserves the separation. The proven transmission gap exceeds 0.05, forcing a common-prediction error greater than 0.025. Numerical estimates are 0.1201756880 and 0.1919847478, a gap of about 7.18091 percentage points. Appendix B separates the finite-cone proof from these estimates; Appendix G records the independent ray checks.

3.5 What resolves the ambiguity—and what does not?

Shape is not orientation. A frame is useful when the shape is fixed. For a 1×0.5 rectangle of depth 0.5 under a direction proportional to $(0.6, 0.1, 1)$, rolls of 0°, 90° and 180° give visible fractions 0.63, 0.38 and 0.63. Here the frame changes, so (Q, A, d) distinguishes the orientations. In the preceding pairs the frame stays fixed while shape changes. Rotating a stored response for an old shape cannot, by itself, update that response to a new shape.

A failed scalar formula is a different failure. One might replace the entire angular response by a single attenuation factor multiplying P_0. That is exact for every illumination only if visibility is constant almost everywhere on the front hemisphere. A channel generally does not have that property. But disproving this scalar formula by changing the light does not prove that its *geometric inputs* are insufficient: another calculation receiving the full illumination could use those inputs differently. Our counterexamples instead keep the illumination fixed and change the geometry without changing the descriptor. Appendix A gives a concrete example of this distinction.

The answer depends on the lighting family. The complete B, v determine the response under the stated law, but need not be the smallest useful record. Suppose the allowed illuminations lie in a finite-dimensional space with basis $\varphi_1, \dots, \varphi_q$. Retaining the response moments

$$m_j(g) = \int_{S^2} \varphi_j(s) H_g(s)\, \mathrm{d}\Omega \tag{21}$$

is enough: for $L = \sum_j \alpha_j \varphi_j$, the answer is $\mathcal{R} = \sum_j \alpha_j m_j$. Spherical harmonics are angular basis functions ordered by degree. Nine exact response moments suffice when the illumination contains only degrees zero through two, even if the opening's response has sharp edges. Smoothness alone does not guarantee this. A finite cone is not in that lighting space, so success for the restricted family does not establish adequacy for a narrow source.

At the other extreme, all ideal angular responses contain much more information. For convex planar mouths with a common normal and known normal extrusion $v = dn$, $d > 0$, that complete family of responses determines shape up to translation and reflection in a point, by an established covariogram theorem [13]. Appendix C explains this connection. It is not a reconstruction claim for a few noisy measurements, nor a theorem that a particular number of stored coordinates is minimal.

4 Retaining shape at controlled precision

Once shape information is retained, rounding creates a separate question: how much can a decoded record change transmission at fixed pose and illumination? The following certificate lets an encoder test that change before discarding the source values.

4.1 Affine slope record

For $d > 0$, choose a shared orthonormal tangent basis E, put $G = E^{\mathsf{T}} B$, $t = E^{\mathsf{T}} v$ and $p = E^{\mathsf{T}} s/\mu$. Then

$$W = dG^{-1}, \qquad \beta = -G^{-1} t, \qquad q = Wp + \beta. \tag{22}$$

Centered channels require four coefficients; offsets add two. The flat case has $T = 1$ and a separate exact mode. With $F(q) = [1 - |q_1|]_+ [1 - |q_2|]_+$,

$$T = \int F(Wp + \beta)\, \mathrm{d}\nu_L, \qquad \mathrm{d}\nu_L = \frac{L(s)\mu\, \mathrm{d}\Omega}{\int_{\mu>0} L(s)\mu\, \mathrm{d}\Omega}. \tag{23}$$

The flux weights sum to one, so a uniform visibility-error bound controls every such weighted average.

4.2 Eight-corner certificate

The closed slope support is

$$\mathcal{P}(W, \beta) = W^{-1}([-1, 1]^2 - \beta). \tag{24}$$

Proposition 4.1 (Decoded-record bound)**.** *Let $W, \widetilde{W}$ be nonsingular and let $\mathcal{V}$ contain the vertices of both original and decoded supports. Write $\Delta W = W - \widetilde{W}$ and $\Delta\beta = \beta - \widetilde{\beta}$. At fixed pose and illumination,*

$$|\widetilde{T} - T| \le C = \min\left\{1, \max_{p\in\mathcal{V}} \|\Delta W p + \Delta\beta\|_1\right\}. \tag{25}$$

The clipped-tent product is 1-Lipschitz in the ℓ_1 norm. Its pointwise error is bounded by the affine-error norm, which is convex and bounded on either support by its largest vertex value. Outside their union both responses vanish. Averaging proves the claim; Appendix D gives details and grazing qualifications. Both supports are needed because rounding may admit formerly blocked directions. The certificate uses at most eight corners; it does not assert that visibility error itself peaks there.

The encoder decodes each candidate as a reader would, rejects overflow or rank loss, and accepts only a qualifying bound, falling back to exact input if necessary. Failure to certify is inconclusive. The guarantee concerns represented inputs in real arithmetic, with exact pose and normalization. A complete budget must separately account for evaluation, integration and input errors:

$$\epsilon_{\text{repr}} + \epsilon_{\text{eval}} + \epsilon_{\text{int}} + \epsilon_{\text{input}} \le \epsilon. \tag{26}$$

A 0.01 representation budget alone therefore does not guarantee 0.01 end-to-end accuracy. Absolute power also needs mouth area and incident irradiance.

4.3 Sharper but different guarantees

For an exact all-source maximum, split both supports at their edges and zero creases. On each resulting polygon the signed visibility difference is quadratic. Check vertices, edge stationary points and interior stationary points; vertices alone miss a half-shear example with zero vertex error and an edge maximum of 1/8. Appendix J.3 gives exact-rational calculations on four post-hoc fixtures. For g66 the directional maximum is about 0.00433468 versus corner bound 0.01149185. The tighter test certifies that record below 0.01, but cannot repair its inaccurate integration or change the original stream accounting.

A restricted source family permits a flux-weighted guarantee. If $D_L = \int_{\mu>0} L\mu \,\mathrm{d}\Omega > 0$ and $L/D_L \le \kappa$ uniformly over allowed sources, then

$$|T - \widetilde{T}| \le \kappa \int_{\mu>0} |V - \widetilde{V}|\mu \,\mathrm{d}\Omega.$$

A radiance cap alone is insufficient: concentration can shrink D_L while preserving normalized error. Appendix J.4 supplies the integral enclosure and tail treatment.

4.4 The 2,000-geometry study

Designed fixtures span moderate, shallow, slender, offset, strongly sheared and flat channels; their frequencies do not model manufacturing prevalence. The encoder chooses the smallest qualifying fixed-10, fixed-12, binary16, binary32 or exact-input record. All three adaptive streams were written, decoded and checked (Appendix F).

At budget 0.01 the stream uses 28,951 bytes versus 103,921 exact-input bytes; the largest accepted bound is below 0.009981. Sizes exclude pose, identifiers and random-access indexing. Six binary16 maps overflow. In the near-parallel stratum, 71 binary32 records have witnessed error above 0.01, including one unit-error probe. Fixed precision is not uniformly safe on this constructed family.

A post-hoc control rounds the primitive mouth matrix and depth instead. Neither chart dominates. Adding primitive-coordinate modes would reduce exact fallbacks from 89 to 86 and project 28,547 bytes, but no combined stream was emitted. The result supports error-controlled selection, not a universally optimal chart.

Table 1: Records satisfying the representation bound. Each column uses the stated absolute transmission budget. Counts are out of 2,000 geometries; stream sizes include headers, mode codes and payloads.

Record choice	0.001	0.005	0.01
Fixed 10-bit	17	298	643
Fixed 12-bit	211	861	899
Binary16 slope	1,122	1,428	1,485
Binary32 slope	1,845	1,893	1,911
Adaptive, including fallbacks	2,000	2,000	2,000
Adaptive stream, bytes	37,733	30,661	28,951
Exact-input fallbacks	155	107	89

5 Numerical adequacy after information is retained

5.1 Checks and decision rules

An independent implementation traces mouth-origin rays against wall triangles using Möller–Trumbore intersections [20], without the overlap product or pseudoinverse. The geometry gate's 558 seed-prefix comparisons and a separate 32-case integrated-ray panel meet their simultaneous sampling bounds. The latter has largest pooled discrepancy 0.00076752 versus radius 0.00319591. These verify implementations of the ideal model at stated precision, not physical detectors (Appendix G).

Earlier deterministic comparisons use a refined reference and empirical allowance u: discrepancy e passes if $e + u \le 0.01$, fails if $e - u > 0.01$, and otherwise remains unresolved. Later enclosures bound an ideal integral under stated arithmetic assumptions. A small refinement change is not such an enclosure; g66 below is a counterexample.

5.2 Shape, azimuth and angular resolution

The original matrix has 112 calibration cases on eight geometries and 240 held-out cases on twelve, including 40 flat controls. Calibration alone selects settings. It separates shape updates, azimuth retention and angular resolution. Here "zonal" means averaging around one aperture-normal axis, not the multi-axis zonal-lobe construction of deformable transfer [8]. No tested whole-kernel SH degree passes all calibration cases; degree 24 is the predeclared diagnostic ceiling.

Table 2: The 240 held-out cases at the declared settings. P/F/U means pass/fail/unresolved under the empirical allowance. Maximum errors are percentage points of unblocked power. "Updated" means recomputed for the current shape; "frozen" means retaining the family's reference-shape response.

Method	P/F/U	Maximum error (pp)	Negative predictions
Scalar attenuation	0/240/0	80.0175	0
Frozen zonal	0/239/1	94.9324	0
Updated zonal	89/151/0	23.1025	0
Frozen full SH	3/236/1	93.9473	0
Updated full SH	221/19/0	4.4984	3
Direct geometry	240/0/0	0.0138	0
Visibility lookup	240/0/0	0.1198	0

Updated full SH passes all 72 held-out 5° cones but only 64/72 at 2° and 61/72 at 1°. Higher degrees are not ruled out. Conversely, nine degree-two moments predict all 24 isotropic/polynomial

held-out cases with maximum discrepancy 4.745×10^{-6}, as the finite-lighting-space argument predicts. Physical range is not an accuracy test: all nineteen full-SH failures remain in $[0, 1]$, while its three negative predictions satisfy the absolute tolerance. Raw outputs are retained.

5.3 Retaining the boundary with SH

A local NumPy adapter implements the clipped-domain equations of Belcour et al. [11]: it integrates low-degree illumination over the projected opening at each mouth position, then integrates over that position. Inscribed polygons approximate circular sources, with omitted cap area bounded separately. The authors' C++/shader implementation was not run.

On twenty reused evaluation cases, direct and lookup each pass twenty; the adapter passes eighteen with two unresolved because of mouth-position integration uncertainty. The separate known worst whole-kernel SH case passes. This post-hoc comparison distinguishes expanding a whole kernel from retaining a boundary; it does not establish failure of the angular SH formula or fresh held-out validation. Appendix H gives equations, checks and local cost limits.

5.4 A sufficient record with an inaccurate integrator

Two separately seeded panels each contain 72 prisms and seventeen broad, cone and mixture queries, totaling 2,448. Inclination, offset and narrow/grazing responses test storage and integration separately.

Table 3: Finite-source panels. Each panel has 1,224 queries. Discrepancies are absolute transmission fractions, not percentage points, and maxima are rounded upward. All non-passing cases here are classified as failures. Binary16 refers to storage; evaluation uses wider arithmetic.

	Passes		Maximum discrepancy	
Method	Panel 1	Replication	Panel 1	Replication
Equal-area square surrogate	430	415	0.88646	0.90814
Packed degree-two SH, all fields	271	252	2.93194	3.14985
Full precision, coarse direct	1,194	1,193	0.62537	0.16059
Binary16, coarse direct	1,194	1,194	0.62498	0.16035
Binary16, intermediate direct	1,201	1,194	0.15842	0.07982
Initial source-routed rule	1,220	1,220	0.01507	0.02913
Revised source-routed rule	1,224	1,222	0.00505	0.02182
Binary16, fine paired integration	1,224	1,224	0.00161	0.00186

Both coarse full-precision and binary16 rules pass all 816 ordinary cases per panel; large errors occur among stress queries. Fine paired binary16 integration passes all queries under empirical allowances. These finite-panel results do not certify arbitrary query integration.

The source-routed rule uses moments for a broad component, direct integration for significant cones, and visibility 1/2 for a component carrying 1% of unblocked power. The last choice contributes at most 0.005 error, independently of the broad integral. After four failures per panel, a rule revised on panel 1 refines broad quadrature until twice the largest specified change is at most 0.001 or a ceiling is reached. It passes panel 1 but fails replication g66's isotropic and mixture queries. An indicator of 0.000359 accompanies an isotropic prediction near 0.168818 versus an independent angular estimate near 0.147004. Post-hoc enclosures establish the failures:

The mixture has 65% isotropic incident flux; its cone transmits nothing, so $T_{\mathrm{mix}} = 0.65 T_{\mathrm{iso}}$. One broad-integral failure causes both misses. Appendix I states the enclosure assumptions.

Table 4: Verified replication failures for g66. Intervals enclose the true ideal-model transmissions under the verifier's stated arithmetic assumptions. Error lower bounds are not empirical refinement indicators.

Source	Prediction, rounded	True transmission enclosure	Error exceeds
Isotropic	0.168818	[0.145911, 0.148097]	0.02072
65% isotropic mixture	0.109731	[0.094842, 0.096263]	0.01346

5.5 The geometric mechanism in g66

With $\mathsf{A} = G/d = W^{-1}$, the smooth displacement factor is

$$J(q) = \frac{1}{|\det W|[1 + \|\mathsf{A}(q-\beta)\|_2^2]^2}.$$

For g66, singular values approximately 54.7021 and 0.9142 create a narrow scale 0.01828 around $\beta \simeq (0.7734, -0.4171)$. Quadrant splitting removes tent creases but leaves this displaced, tilted strip unresolved. On the decoded record, orders 16/32 give 0.16863832/0.16881765, nearly agreeing overestimates; orders 24/64 give 0.08994343/0.14681339. The original rule did check each coordinate. Appendix J.1 relates the hazard to nearby complex poles without asserting a universal convergence law.

Analytically integrating the narrow coordinate for isotropic light gives estimates 0.147004102183 (original) and 0.146967585774 (decoded). The estimated representation change is 3.65×10^{-5}, while the saved output exceeds the decoded estimate by 0.02185. This post-hoc reduction resolves the observed mechanism; the enclosures, not agreement of decimal estimates, prove failure. Section 8 audits its wider performance and limit.

5.6 Workload-dependent cost

A post-main efficiency study fixes the 33×128 lookup, selects cheaper query rules on calibration and confirms direct/lookup accuracy on 48 fresh cases. A later cost-only amendment adds a reusable broad-field direct cache without changing predictions. Construction, source preparation and shape updates are charged. For 1,000 mixed queries, median seconds for fixed/changing shape are 0.093966/0.132537 (direct), 0.087464/0.140566 (cached direct), and 0.083372/0.484670 (lookup), from three batches per entry. These historical local measurements include setup and updates.

At high reuse the lookup and cached-direct medians differ by about 4.7% on one machine; rebuilding lookup per query costs substantially more. Degree-two SH is useful on the restricted broad-field workload, but degree-24 SH supplies no matched-tolerance mixed-workload speed claim because it fails that target. Later warmed panel timings exclude setup/transfer costs and have overlapping ranges. Neither set establishes a general speed or energy advantage. Appendix G gives memory and timing boundaries; Appendix K preserves chronology.

6 What a daylight angular patch retains

Mardaljevic evaluates the contribution of a luminous exterior category at a building aperture through the view lumen [29]. Our observable is the fraction of incident power transmitted through a recess. Both quantities weight the visible exterior by direction. Ward et al. describe the use of tabulated bidirectional scattering distribution functions (BSDFs) in Radiance [5]. Geisler-Moroder et al. treat the small apparent size of the direct sun by separating its peak from scattered light [14]. Their subsequent position paper discusses data generation and application-dependent

angular resolution [18]. We ask whether the incident power in each angular cell suffices to predict recess transmission for sources in the same cell.

Use straight rectangular recesses with $(a, b, d) = (1, 1, 1)$ and $(2, 1/2, 1)$. Put $d\nu = \cos\theta\, d\Omega$, $P_j = \int_{B_j} L\, d\nu$ and

$$T[L] = \frac{\int LV\, d\nu}{\int L\, d\nu}, \qquad \bar{V}_j = \frac{1}{\nu(B_j)} \int_{B_j} V\, d\nu, \qquad T_{\text{patch}} = \frac{\sum_j P_j \bar{V}_j}{\sum_j P_j}.$$

The projected-solid-angle average preserves exact transmitted flux whenever L is constant on each patch. Choose the full Klems cell $25° \le \theta \le 35°$, $|\phi| \le 9°$. Its ring has 20 equal azimuth sectors [30]. In turn place uniform circular caps of half-angle $\alpha = 0.266°$ at $\phi = 0$ and $\theta = 26°, 30°, 34°$. Normalize the radiance of each cap separately to give $\int L\, d\nu = 1$. Each cap lies entirely inside this cell. The nearest polar edge is at least 1° from the axis, and the distance to the nearest bounding meridian is $\arcsin(\sin\theta \sin 9°) > \alpha$. Thus the three 145-entry records P coincide exactly. Both geometries retain their own scalar patch response throughout.

Table 5: Averaging the scalar recess response over one full Klems incident cell. Errors are signed percentage points, $100(T_{\text{patch}} - T[L])$. The uniform-cell control has zero error by construction.

	Square		Rectangle	
Illumination	$T[L]$	Error	$T[L]$	Error
26° cap axis	0.511144	−11.08	0.752818	−10.78
30° cap axis	0.421688	−2.14	0.708088	−6.31
34° cap axis	0.324718	+7.56	0.659595	−1.46
Uniform cell	0.400298	0	0.645005	0

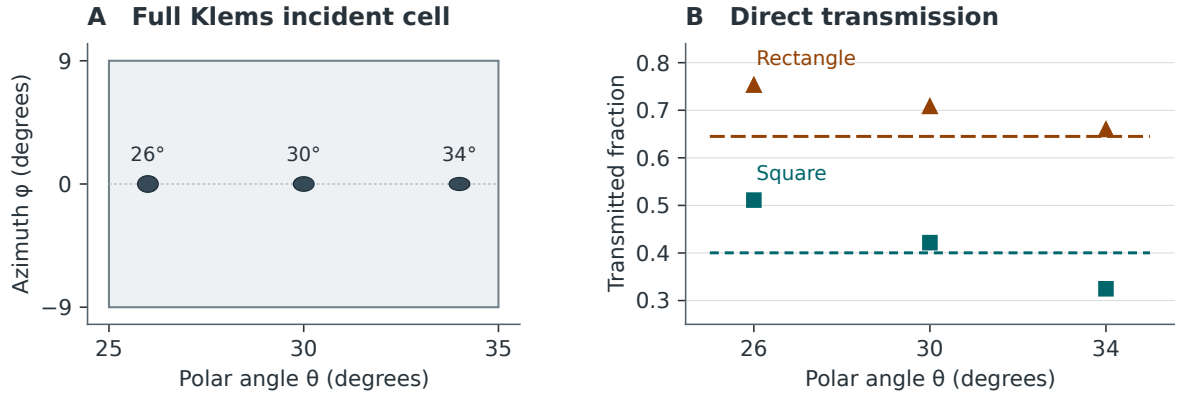


Figure 2: Finite sources in one incident angular cell. (A) Boundaries of the three 0.266° half-angle caps in polar–azimuth coordinates, plotted at their actual angular extents. Each source contributes the same total projected incident power to the shaded Klems cell. (B) Symbols give direct transmission for the individual caps. Dashed lines give the projected-solid-angle mean response of each aperture over the whole cell.

A second calculation bounds the error of every predictor that receives only the common incident record and the fixed geometry. For these caps at meridional angle t, s_x/s_z lies between $\tan(t-\alpha)$ and $\tan(t+\alpha)$, and $|s_y|/s_z \le \sin\alpha/\cos(t+\alpha)$. Projection onto the xz plane reduces the angle to the cap axis. With $[z]_+ = \max(0, z)$ the inequalities imply

$$l_{a,b}(t) = \left[1 - \frac{\tan(t+\alpha)}{a}\right]_+ \left[1 - \frac{\sin\alpha}{b\cos(t+\alpha)}\right]_+ \le T[L_t] \le u_a(t) = \left[1 - \frac{\tan(t-\alpha)}{a}\right]_+ .$$

Half of $l_{a,b}(26°) - u_a(34°)$ gives a lower bound on the largest absolute error at the two endpoints. Its value is approximately 8.583 percentage points for the square and 3.967 for the rectangle. Within each comparison the geometry is fixed. The information lost is the position of the source within the cell. Methods retaining the solar direction or finer incident data receive different inputs.

For the patch mean we integrate the rectangular overlap polynomial analytically over the cell. Each cap numerator uses composite Simpson quadrature on half the cap, reflected across $s_y = 0$, at 32, 64 and 128 subintervals in both coordinates. The denominator before radiance normalization is $\pi \sin^2 \alpha \cos t$. All six estimates change by less than 1.1×10^{-10} from 64 to 128. Cancellation of the equal-area product term provides an analytic check on the shape difference at each axis,

$$T_{2,1/2} - T_{1,1} = \tfrac{1}{2} \tan t - \frac{2\alpha - \sin(2\alpha)}{\pi \sin^2 \alpha \cos t}.$$

The numerical differences agree with this expression within 5×10^{-12}. The individual quadrature uncertainties remain empirical. This experiment concerns incoming-direction discretization for total direct power. Section 6.1 tests variable-resolution data and explicit solar directions on the same total-flux observable. Glare and outgoing radiance require different measurements.

As a separate software check, LBNL Radiance 6.0.2 traced 32,768 independent rays per shape and cap, using four black wall polygons and a unit-radiance rear target [32]. Entrance origins were uniform, and cap sampling used rejection to draw from projected solid angle. For all 196,608 rays the first surface and returned radiance agreed with an independent rear-plane intersection test. The six sampled transmissions differed from cap quadrature by at most 0.00642, or 2.56 binomial standard errors. Evaluating the distributed `klems_full.cal` bin function with `rcalc` put every direction in bin 25 (zero-based). This checks the direct geometric response and incident bin assignment in the chosen convention.

6.1 Crossing a patch boundary

We compare whole-exit flux for explicit solar directions and tabulated angular responses. The full Klems baseline integrates the actual disk flux in every overlapped cell and uses each cell's projected-area mean visibility; it does not assign the whole disk to its center cell.

The two shapes and four half-angles $0.1°, 0.266°, 1°, 2°$ cross three boundaries at 25 positions $u = -3°, -2.75°, \ldots, 3°$. Tracks are $(\theta, \phi) = (35° + u, 0)$, $(45° + u, 0)$, and $\theta = 30°$, $\phi = 9° + \arcsin(\sin u / \sin 30°)$. Here u is signed angular distance. The resulting 600 cases include 156 disks with positive area on both sides of a boundary.

Neutral rank-four Radiance XML transmission trees [39, 32] use the equal-projected-area disk–square map [38]. Maximum depth m gives 2^m incident-square cells per axis. Finest means use 8×8 Gauss nodes; a parent merges if the range of descendant means is at most 0.005. This controls means, not variation inside a finest cell. A retained leaf at depth ℓ has projected area $\pi 4^{-\ell}$; its binary32 density sends power uniformly into the matching outgoing leaf, with zero off-diagonal density. Whole-exit power is preserved up to rounding, but a solar image is spread.

Primary depths are 4, 6 and 8; depths 5 and 7 were added after those decisions. All ten serialized XMLs are retained. The deepest square/rectangle records have 36,760/28,612 incident leaves. Native Radiance readback at all 174,592 finest-cell centers agrees with decoded marginals below 10^{-28}. The center predictor evaluates one explicit ray; the finite-disk predictor integrates continuous visibility using analytic polar integration split at sign and clipping changes. Both assume known source support and geometry.

An independent verifier projects each cap to an ellipse with uniform projected-flux measure. Sector-centroid integration uses outward arithmetic, derivative bounds on smooth parts and Lipschitz bounds across creases. Every reference width is below 0.010 percentage points. For pre-

diction v and enclosure $[l, h]$, error lies in $[\max(0, l - v, v - h), \max(|v - l|, |v - h|)]$; pass/fail decisions compare its endpoints to 0.01. Separate XML overlap brackets and Klems/disk order checks retain every classification. The supplement specifies arithmetic assumptions and full protocols.

Table 6: Whole-exit transmission within one absolute percentage point, out of 150 cases per disk half-angle. The last column is the largest upper bound, in percentage points, on the error of a returned prediction, rounded upward. All decisions resolved. Depths 5 and 7 were added in a follow-up.

	0.1°	0.266°	1°	2°	All / 600	Worst bound
Full Klems	35	36	35	43	149	15.358
Tensor tree, depth 4	27	27	32	40	126	10.459
Tensor tree, depth 5	61	63	77	98	299	7.192
Tensor tree, depth 6	75	85	134	150	444	3.446
Tensor tree, depth 7	118	140	150	150	558	1.557
Tensor tree, depth 8	150	150	150	150	600	0.747
Solar center	150	150	103	97	500	2.405
Finite solar disk	150	150	150	150	600	0.008

Depth eight and finite-disk integration pass all 600 cases; the center ray passes all 150 solar-sized cases but only 97/150 at 2°. Among the 156 edge-spanning disks, Klems, center and depth seven pass 34, 109 and 154. Source extent and local visibility, not distance from a patch boundary alone, determine adequacy.

On the 35° meridian, both overlap factors remain active for $32° \le t \le 38°$ and $\alpha \le 2°$. Projected weighting gives $\mathbb{E}p_x = \tan t$ and

$$T_{\text{center}} - T_{\text{disk}} = \frac{1}{b}\mathbb{E}[(1 - p_x/a)|p_y|], \qquad \mathbb{E}|p_y| = \frac{2\alpha - \sin(2\alpha)}{\pi \sin^2 \alpha \cos t}.$$

Bounding p_x by $\tan(t \pm \alpha)$ gives continuous-sweep error below 0.795 points for the square at every tested radius. For the rectangle it is below 0.130/0.346 points at 0.1°/0.266° and above 1.039/2.028 at 1°/2°. At the square's 45° cutoff the center ray is zero while the 2° disk's enclosed transmission is 1.4535–1.4594%. Fifty-one post-hoc finer-offset decisions locate neighboring pass/fail positions; they are separate sampled results, not continuous guarantees or additional primary cases.

6.2 Flux into a restricted outgoing view cone

An accurate prediction of total transmission leaves open where the transmitted light goes. We used the ten Radiance XML records from Section 6.1 without rebuilding them. In the ideal straight absorbing recess, every transmitted ray preserves its incident propagation direction. In an XML leaf, transmitted power occupies a finite outgoing angular cell. A receiver that accepts only a cone of directions can distinguish these laws even when their total fluxes agree.

Two angular integrals. Let S be the incident solar disk, C the outgoing acceptance cone, and $\nu(D) = \int_D \cos\theta \, \mathrm{d}\Omega$ the projected solid angle in a common forward-propagation chart. Both angular coordinates are referred to the local aperture frame. The native Radiance incident and outgoing vector conventions are converted to this chart in the supplement. The solar disk has uniform radiance across directions and mouth positions. The two centered rectangular recesses have the same visibility under reversal of the transverse axes. With V given by Eq. (5), the ideal response is

$$B_C = \frac{1}{\nu(S)} \int_{S \cap C} V(s) \, \mathrm{d}\nu(s).$$

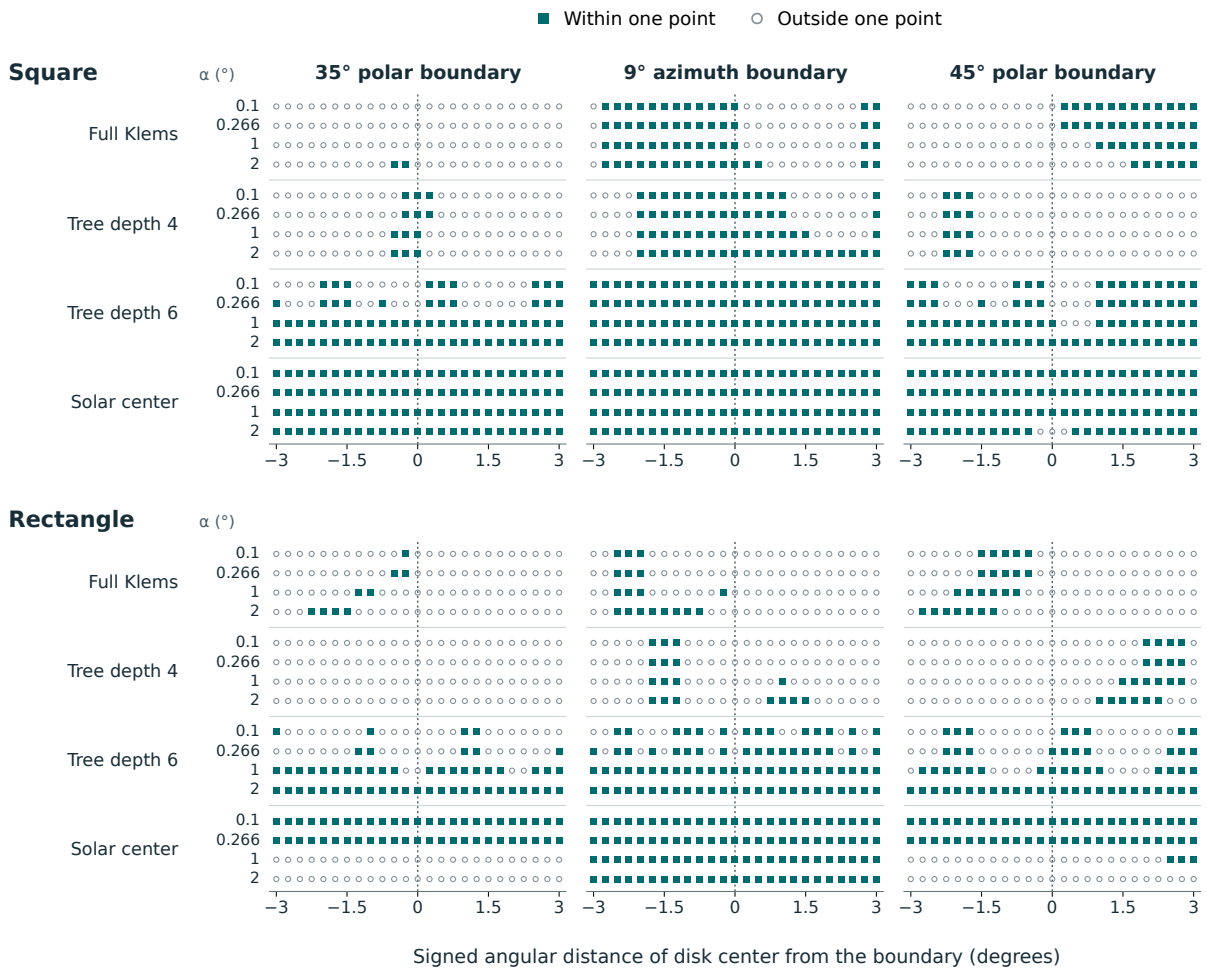


Figure 3: Where each returned transmission fell within one percentage point of the source-resolved reference. Each symbol represents one disk center, sampled at 0.25° intervals. The dotted centerline marks the Klems boundary. For each method the four rows give disk half-angle α. The polar crossings have azimuth zero; the azimuth crossing has polar angle 30°. The depth-eight tree and the explicit finite-disk calculation met the target for all 600 cases. Unmarked angles between sampled centers are not scored here.

Each stored leaf L_j has the same incident and straight-through outgoing cell in this chart. Its nonnegative BTDF density Y_j is constant on the corresponding cell pair. Its contribution is an integral over the pair of cells, giving

$$P_C = \frac{1}{\nu(S)} \sum_j Y_j \, \nu(S \cap L_j) \, \nu(C \cap L_j).$$

Both answers are fractions of the *whole unblocked incident solar power*. A one-point error means 0.01 on this common scale, including when $B_C = 0$. The model scores total power arriving within the angular acceptance over the whole exit. A particular room or spatial receiver would require additional geometry.

Solar centers (θ, ϕ) were $(35°, 0)$, $(30°, 9°)$ and $(45°, 0)$, with the same four half-angles $\alpha = 0.1°, 0.266°, 1°, 2°$. For each solar disk we used receiver half-angles $\beta = \alpha/2, \alpha, 2\alpha, 4°$, counting coincident radii once. We tested a centered receiver and displacements of $0.25, 0.5, 0.75, 1, 1.25$ times $\alpha + \beta$, in both directions along the local meridional and azimuthal tangents. Each displacement is a great-circle distance. This gives 1,890 source–view–shape configurations and 9,450

depth-specific comparisons. A displacement equal to $\alpha+\beta$ gives external contact and zero ballistic power. A greater displacement leaves the caps disjoint.

Reference and numerical checks. Orthographic projection turns the two caps into ellipses with uniform projected measure. The verifier encloses their intersection and visibility integral using spherical boundary arcs, derivative bounds on smooth parts and Lipschitz bounds at creases. All 1,890 independent meridional estimates lie within these reference intervals, each narrower than 0.025 percentage points. Actual XML leaves are decoded and hashed; cap masses are integrated separately. Two threshold-overlapping numerical sensitivity screens required refinement. A hemispherical receiver reproduces 120 preceding whole-exit outputs within 5.9×10^{-9}.

Native Radiance 6.0.2 sampling on selected cases is consistent under nominal simultaneous 99% intervals. Because its evaluator perturbs density slightly, a separate check queries all 88,456 diagonal leaves: a relative bound settles 9,412 decisions and native integration settles the remaining 38 with unchanged classifications. These checks remain distinct from physical validation and the ideal-reference enclosures.

Where accuracy survives angular restriction. At depths 4 to 8 the numbers within one point were 580, 647, 761, 1,105 and 1,307 out of 1,890. At depth 8, 24 of 24 centered 4° receivers passed. With a centered receiver *matching the solar radius*, only four of 24 passed. All four were the square at its 45° cutoff, where little incident power reaches the exit. Among the 987 configurations whose ballistic reference exceeded 1% of incident power, 599 passed at depth 8. These counts give equal weight to the designed configurations.

Figure 4 shows the role of direction and acceptance width. For the rectangle at $(35°, 0)$ with $\alpha = 0.266°$, the ballistic result for a centered matching cone was 64.68%. The depth-eight XML returned 34.99% there. Increasing only the receiver half-angle to 4° brought the XML result to 64.47% while the ballistic answer stayed fixed. The same spreading can put light into a view that contains no ballistic ray. At depth 8, 141 of 360 tangent views and 54 of 360 strictly exterior views received more than one percentage point of incident power. With the rectangle at $(35°, 0)$ and $\alpha = \beta = 0.266°$, a polar offset of $+1.25(\alpha + \beta)$ gives a strictly exterior view. Here depth 6 returned 2.246%, and depths 7 and 8 returned zero.

Widening a centered receiver beyond the source radius collects all ballistic light and monotonically increases nonnegative XML power. A post-hoc search brackets the first passing half-angle for each of the twenty failing depth-eight matched cases to below 0.012°, after checking the whole-exit upper tolerance. The archived tables retain every endpoint and screen; this is not a general angular-resolution rule.

A necessary resolution bound. In the Radiance basis convention, a tree of maximum depth m assigns each outgoing leaf a projected solid angle of at least $\pi 4^{-m}$. For a circular receiver wholly within the forward hemisphere ($\theta_C + \beta < \pi/2$), its projected area is $\pi \sin^2 \beta \cos \theta_C$. It can collect at most the following fraction of each leaf's uniformly distributed power, hence of the modeled whole-exit return P_H:

$$P_C \le f_m P_H, \qquad f_m = \min(1,\, 4^m \sin^2 \beta \cos \theta_C).$$

For coincident source and receiver disks the ideal restricted value equals the ideal whole-exit value B. Meeting tolerance ϵ on both observables then requires $f_m \ge (B - \epsilon)/(B + \epsilon)$ whenever $B > \epsilon$. For $\alpha = \beta = 0.1°$ at 35°, the depth-nine area factor is at most 0.654122. Independently enclosed full-disk lower bounds give required fractions above 0.9353 for the square and 0.9696 for the rectangle at $\epsilon = 0.01$. Thus maximum depth *at least ten* is necessary in this specific uniform-in-leaf diagonal encoding if it is to meet both targets. In the original depth-eight record the centered

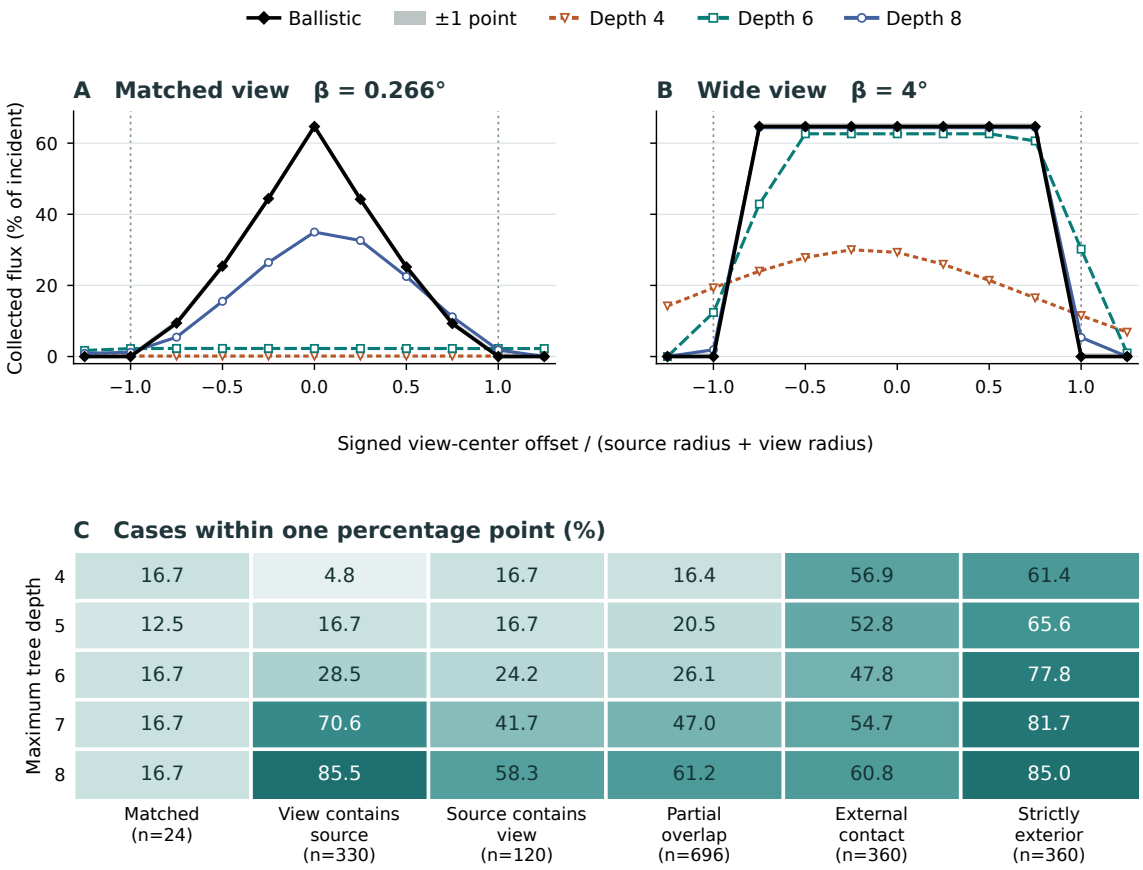


Figure 4: Angular collection by the rectangle for a 0.266° solar radius centered at $(35°, 0)$, and pass rates across the two-shape panel. (A, B) The view moves along the polar meridian. Symbols are the sampled positions, lines connect them, and shading marks one percentage point around the ballistic reference. The dotted lines at signed offsets ±1 mark external contact. (C) Each percentage counts positions whose screened absolute error is at most one percentage point of total incident disk flux. The sample sizes apply separately to each depth. Matched denotes coincident equal caps, and the containment columns exclude that case.

rectangle actually delivered 5.27% into the 0.1° view against a ballistic 64.87%. Section 6.3 tests the necessary condition on new records at depths nine and ten.

6.3 Do depths nine and ten resolve a matching solar disk?

The area bound in Section 6.2 leaves an experimental question. A depth-ten lattice admits cells small enough to clear that particular necessary bound, but the stored response can merge those cells. We therefore built a complete depth-nine and depth-ten XML for each recess using the same geometry, eight-node-per-axis cell integration, 0.005 merging criterion and binary32 diagonal density as before. As a continuity control the generator reproduced both depth-eight XMLs and their decoded whole-exit response tables byte for byte. We selected the same 24 centered, equal-radius source and view configurations as in Figure 4. A depth refers to the maximum allowed subdivision.

For a disk S that is also the view, put $p_j = \nu(S \cap L_j)/\nu(S)$, $A_j = \pi 4^{-\ell_j}$ and $M_j = Y_j A_j$. Here p_j is the incident disk fraction falling on leaf j, and M_j is its stored transmission integrated over the outgoing hemisphere. The two returns from one record are

$$H = \sum_j p_j M_j, \qquad P_S = \sum_j p_j M_j \, \frac{\nu(S) p_j}{A_j}.$$

The additional factor in P_S is the fraction of the outgoing leaf accepted by the matching view. An incident direction can lie inside S while a portion of its tabulated output falls outside. The ballistic return B still equals the ideal whole-exit integral. Thus $B - P_S = (B - H) + (H - P_S)$ separates the whole-exit discrepancy from power spread beyond the receiver.

All four XMLs are scored against independently enclosed ballistic targets. Refinement to numerical-screen width below 0.030 percentage points leaves all 48 decisions unchanged; native conditional-scattering/ray checks agree under the stated simultaneous sampling intervals.

Table 7: Absolute matched-view error in percentage points of unblocked incident disk power. Each shape and axis has four source/view half-angles at each maximum depth. Bold entries meet the 1.000-point target. The displayed errors use an independent ballistic quadrature estimate; classifications use the enclosed reference and the unrounded numerical screens.

		Maximum depth 9				Maximum depth 10			
(θ, ϕ)	Recess	0.1°	0.266°	1°	2°	0.1°	0.266°	1°	2°
$(35°, 0°)$	Square	23.527	8.350	2.456	1.096	23.527	8.350	2.456	1.096
	Rectangle	51.035	18.046	5.269	2.291	51.035	18.046	5.269	2.291
$(30°, 9°)$	Square	27.895	7.193	2.821	1.489	27.895	7.193	2.821	1.489
	Rectangle	41.789	10.779	4.225	2.237	41.789	10.779	4.225	2.237
$(45°, 0°)$	Square	**0.040**	**0.020**	**0.068**	**0.093**	**0.009**	**0.040**	**0.045**	**0.054**
	Rectangle	42.830	17.324	3.667	1.952	42.830	17.324	3.629	1.792

Both maximum depths passed exactly four of the 24 configurations. These were the square on the 45° cutoff at all four source radii. Its ballistic signals ranged from 0.074% to 1.455% of incident flux; the three narrower disks transmitted less than 1% of incident power in total. The entire rectangle panel missed, as did the square on the other two axes. The nearest miss was the square at $(35°, 0)$ with a 2° disk: its error was about 1.096 percentage points at each depth. The lower end of the numerical error screen exceeded 1.088 points at both depths. All 48 *whole-exit* predictions lay within 0.19 points of their respective ballistic targets.

Eighteen configurations encounter only level-nine leaves even in the depth-ten record; six encounter level-ten leaves without changing classification. Only 3,360 of the square's 147,820 leaves and 116,992 of the rectangle's 200,944 reach level ten. For the 0.1° rectangle at $(35°, 0)$, whole-exit $H = 64.79\%$ nearly matches ballistic $B = 64.87\%$, but matched $P_S = 13.84\%$: 50.95 points of incident power are assigned outside the receiver. A maximum-depth condition alone cannot guarantee angular delivery.

Separating cell merging from the finite lattice. A post-hoc diagnostic kept each stored incident marginal M_j fixed. It reassigned straight-through power to individual finest-grid cells $g \subseteq L_j$, each of projected area $A_g = \pi 4^{-m}$, instead of distributing power over the whole retained leaf. Write $p_g = \nu(S \cap g)/\nu(S)$ and $n_j = A_j/A_g$. In a matching view the increase in collection is

$$P_{\text{fine}} - P_S = \nu(S) \sum_j \frac{M_j}{A_g} \sum_{g \subseteq L_j} (p_g - p_j/n_j)^2 \geq 0.$$

This follows by expanding the squares and using $\sum_{g \subseteq L_j} p_g = p_j$. The whole-exit value H is identical by construction. For the 0.1° rectangle example, subdivision at the nominal depth-ten

spacing raised collection to 47.50%, leaving a 17.37-point deficit against the ballistic reference (Figure 5). Both retained-leaf merging and the finite angular cells themselves contributed to the loss.

The diagnostic met the target in six depth-ten cases. Its two additional passes were the 2° disks through the square at $(35°, 0)$ and $(30°, 9°)$. These are outcomes for a reassignment of the outgoing density. Table 7 reports the four XMLs built under the original 0.005 rule. Section 6.4 tests this diagnostic against newly emitted depth-ten XMLs.

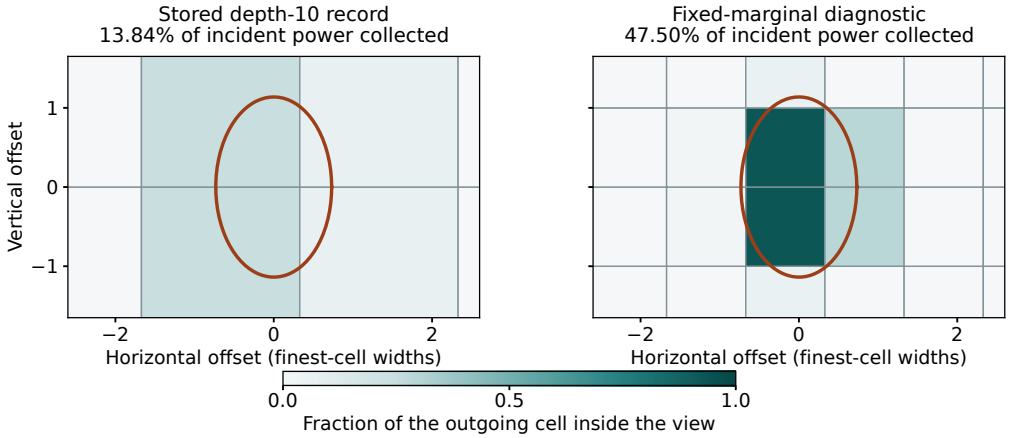


Figure 5: One matched 0.1° source and view through the rectangle at $(35°, 0)$, in local coordinates on the depth-ten angular lattice. The orange curve is the disk boundary. Fill shows the fraction of each outgoing cell lying inside it. The left panel uses the level-nine cells actually retained by the depth-ten XML. The right panel uses finest cells only in the fixed-marginal diagnostic. Percentages above the panels are collected power as a fraction of the same unblocked incident power.

6.4 Matched-view accuracy and storage at depth ten

Of the 24 matched views in Section 6.3, 20 missed the one-percentage-point target with the original depth-ten trees. We rebuilt both recesses at the same maximum depth under three policies. A parent cell may be retained when the range of the mean visibilities of its finest descendants is at most τ. We lowered τ from 0.005 to 0.0025 and 0.001. For the third policy we split every diagonal branch to depth ten, including branches of zero transmission. Each unmerged XML contains 4^{10} diagonal leaves. Off-diagonal entries remain zero. The two original XMLs were reproduced byte for byte before changing the rule. All six new files use the same cell integration and binary32 density encoding.

Which misses cross the threshold? All three policies recovered the same two square cases. For a 2° disk at $(35°, 0°)$, tightening the rule to 0.0025 reduced the absolute error from 1.096 to 0.530 percentage points. At $(30°, 9°)$ it reduced the error from 1.489 to 0.736 points for the same radius. The four original passes at the square's 45° cutoff remained within tolerance. The six smaller disks on the square at the other two axes and all twelve rectangle cases still missed. At the solar-sized radius of 0.266°, the square at 45° remained the only pass. Table 9 gives the full unmerged result. The closest remaining case was the 2° rectangle at $(45°, 0°)$. Its estimated error was 1.042 points, with its numerical error screen inside $[1.037, 1.054]$ under each new policy.

Table 8: Storage for the two depth-ten XMLs together and outcomes for the same 24 matched views. Recovered counts are relative to the 20 baseline misses. File sizes are in MiB (2^{20} bytes). Gzip uses level 6.

Merge rule	Diagonal leaves	XML (MiB)	Gzip (MiB)	Within / 24	Recovered / 20
0.005	348,764	9.82	1.87	4	0
0.0025	992,732	27.95	5.05	6	2
0.001	1,033,352	29.08	5.24	6	2
Disabled	2,097,152	50.07	5.58	6	2

Table 9: Absolute matched-view error with merging disabled, in percentage points of unblocked incident disk power. Columns give the common source and view half-angle. Bold values meet the target. A star marks a baseline miss recovered by all three new policies. Values use the ballistic quadrature estimate for display. Decisions use the unrounded XML integration screen and independently enclosed reference.

(θ, ϕ)	Recess	0.1°	0.266°	1°	2°
$(35°, 0°)$	Square	7.986	4.020	1.150	**0.530***
	Rectangle	17.334	8.702	2.455	1.117
$(30°, 9°)$	Square	15.920	4.295	1.532	**0.736***
	Rectangle	23.851	6.437	2.294	1.102
$(45°, 0°)$	Square	**0.009**	**0.040**	**0.045**	**0.049**
	Rectangle	24.255	8.376	2.009	1.042

What the extra storage changes. At $\tau = 0.0025$, each of the 20 previously failing disks already illuminates only depth-ten leaves. Decoding the two tighter adaptive XMLs for each recess also gives identical leaf depths and densities throughout the bounding box of every tested disk. The extra storage at 0.001 therefore changes cells outside the directions queried here. Across the 24 views the largest change in matched flux from $\tau = 0.0025$ to no merging was below 0.00014 percentage points. Disabling merging added 555,020 compressed bytes relative to 0.0025 (Table 8). Many additional unmerged leaves have zero transmission and repeat the same text in the XML.

The full-lattice test also checks the fixed-marginal diagnostic in Section 6.3. The emitted unmerged file assigns each finest cell its own computed marginal. The diagnostic assigns it the former parent's marginal. Their matched returns differed by at most 0.086 points on this panel and gave the same six passes. For the 0.1° rectangle disk at $(35°, 0°)$, the actual unmerged XML collected 47.54% of incident power inside the view and 64.85% over the whole exit. The ideal ballistic return was 64.87%. Even at full depth ten, power amounting to 17.31% of the incident flux was transmitted outside the matched view. At this fixed depth, tightening the merge rule cannot concentrate the outgoing power any further.

Numerical checks. We decoded each emitted XML and compared the matched-view flux with the independently enclosed ballistic reference, normalizing both to unblocked projected incident

disk power. Refinement reduced every new matched-flux integration screen to a width below 0.029 points. All 72 whole-exit error screens were below 0.064 points. In a separate check using Radiance 6.0.2, paired conditional-scattering and geometric samples gave 69 decisions consistent with the primary calculation under simultaneous 99% sampling intervals. The three intervals that straddled one point belonged to the same 2° rectangle at 45°, once per policy. Widening each deterministic XML flux screen by the native evaluator's tested relative density allowance of 0.0201% preserved all 72 primary decisions. The supplement gives the unrounded screens, sampling intervals and exact file sizes for each recess.

6.5 A diffuse-wall calculation for one specified beam

Let T_0 denote absorbing-wall transmission of an open, passive recess, with spatially uniform entrance illumination and no exterior return. If hemispherical wall reflectance at every hit is at most ρ, the total rear-exit fraction obeys $T_0 \le T \le T_0 + \rho(1-T_0)$. All power that reaches the rear after reflection must survive the first wall hit. The fraction reaching the rear after at least two hits is at most $\rho^2(1-T_0)$ by applying the same argument at the second hit. The normalization throughout is incident entrance power.

These inequalities also apply to the finite disks. Replace the upper bound u by $\rho + (1-\rho)u$ and keep the absorbing-wall lower bound l. For $\rho \le 0.2$ the resulting source-record minimax bounds at 26° and 34° are approximately 1.905 and 0.628 percentage points for the square and rectangle. At fixed source, comparing the rectangle lower bound against the square upper bound gives guaranteed transmission gaps exceeding 13.10, 15.75 and 18.61 percentage points at $26°, 30°, 34°$, respectively. The bounds allow different wall responses in the two geometries.

The same calculation applies to the earlier 2° cap centered at $t = \arctan(1/2)$. The absorbing-wall gap formula above gives 0.23343384, and $l_{1,1}(t)$ for that cap is approximately 0.43747048. Even assigning the square the maximum additional reflected power and the rectangle none gives $T_{\mathrm{rect}} - T_{\mathrm{sq}} \ge \Delta T_0 - 0.2(1-l_{1,1}) > 0.1209$. This conclusion requires a bound of 0.2 on hemispherical reflectance and collection of the full rear-exit power.

For the following ray calculation z increases from the illuminated entrance toward the rear exit. We give the magnitude of the transverse shift per unit inward depth. Reversing that sign exchanges the struck sidewall in these centered rectangles and leaves the transmission unchanged. The congruent parallel mouths and the spatially uniform angular field give equal unblocked powers at the two reference planes. Reflected contributions here use whole rear-exit flux.

A more specific calculation gives the once-reflected contribution. Use a collimated meridional beam of slope $p = 1/2$ incident uniformly over the entrance, the same two rectangles, and homogeneous Lambertian walls of albedo ρ. Since $pd \le a$, the fraction $h = pd/a$ that misses the exit strikes only the $x = a$ sidewall and illuminates it uniformly. The relation $x_0 = a - pz$ maps entrance position to first-hit depth with constant Jacobian, and the transverse coordinate y is unchanged. The sidewall-to-exit diffuse view factor F is the perpendicular common-edge rectangle configuration [31]. Inward diffuse emission sends fraction F to the rear, fraction F to the entrance by symmetry, and $1-2F$ to the remaining walls. Hence

$$T_1 = \rho h F, \qquad 1 - h + T_1 \le T \le 1 - h + T_1 + \rho^2 h(1-2F).$$

Here T_1 is the once-reflected rear-exit fraction. At $\rho = 0.2$, $F = 0.200044$ for the square and 0.167309 for the rectangle. The respective once-reflected fractions are 0.020004 and 0.008365. Including all remaining orders gives the model bounds $[0.52000, 0.53201]$ and $[0.75836, 0.76503]$, rounded outwards. The rectangle exceeds the square by at least 22.63 percentage points for this beam and wall model. An independent check sampled 100,000 uniform wall origins and cosine-weighted directions per geometry and counted rear-plane hits. View-factor estimates were 0.19809 and 0.16796, with binomial standard errors approximately 0.00127 and 0.00118. This beam has

zero angular extent. The finite caps in Table 5 use absorbing walls. The selected albedo is an input to this calculation.

7 Convex mouths, interior stops and measured tapers

A model built from the two openings can account for all potential straight rays between them. Its accuracy depends on whether the intervening wall clears those rays. We begin with extrusions of a convex mouth, then introduce an internal stop, and finally vary a circular radius continuously along the axis. The last case allows us to test what a finite set of wall measurements can certify.

7.1 From the covariogram to a displacement integral

Take parallel reference planes, a spatially uniform angular field, straight rays and perfectly absorbing boundaries. For a convex mouth M of area $A > 0$, the set of starting positions admitting a ray has area $c_M(z) = |M \cap (M + z)|$ for a translated copy displaced by z. The complete-response result for a known normal extrusion already follows from planar covariogram uniqueness (Appendix C, [13]). For an extrusion of positive depth d and tangential offset t, use orthonormal plane coordinates, $z = dp - t$ and $r(z) = (z + t, d)$. The cosine is $d/\|r\|$ and the solid-angle element is $d\,\mathrm{d}z/\|r\|^3$. Equation (9) thus has the following physical-coordinate form for any such convex mouth:

$$\mathcal{R} = d^2 \int_{M-M} \frac{L(r/\|r\|)\, c_M(z)}{\|r\|^4}\, \mathrm{d}z. \tag{27}$$

Here $M - M = \{x - y : x, y \in M\}$ is the closed displacement support. The common area factor enters through c_M. For equal coaxial disks of radius a, isotropic transmission follows independently from Sullivan's two-disk geometrical factor divided by $\pi^2 a^2$ [24, Eq. 8]. With $u = d/a$, this gives $2/(2 + u^2 + u\sqrt{4 + u^2})$. Radial integration of Eq. (27) with 256 Simpson subintervals at $u = 1/4, 1, 4$ gave maximum absolute discrepancy 7.85×10^{-7}. This compares two calculations on three specified geometries.

The concavity relevant here is that of $\sqrt{c_M}$ on $M - M$ [27, Prop. 2.5]. Indeed, put $K_z = M \cap (M + z)$. Convexity gives $(1 - \lambda)K_x + \lambda K_y \subset K_{(1-\lambda)x+\lambda y}$ for $0 \le \lambda \le 1$. The planar Brunn–Minkowski inequality gives the asserted square-root concavity. The unit square gives $c_M(t, t) = (1 - t)^2$ for $0 \le t \le 1$, with midpoint value $1/4$ below the endpoint average $1/2$.

7.2 Certifying a rounded map for a fixed polygon

Let M now denote a fixed convex polygonal template in reference coordinates. An affine image of the template evaluates $F_M(q) = c_M(q)/|M|$ at $q = Wp + \beta$ as in Eq. (22). Define $b_M(h) = \max_{x,y \in M} |\det(h, x - y)|$. This equals the length of h times the width of M perpendicular to h.

Proposition 7.1 (Fixed polygonal template)**.** *Let both records use the same template, tangent basis and normalized flux weights. Suppose W and $\widetilde{W}$ are invertible. Let $\mathcal{U}$ be the vertices of both $W^{-1}((M - M) - \beta)$ and $\widetilde{W}^{-1}((M - M) - \widetilde{\beta})$. With $\Delta W = W - \widetilde{W}$ and $\Delta\beta = \beta - \widetilde{\beta}$,*

$$|T - \widetilde{T}| \le \min\left\{1, \max_{p \in \mathcal{U}} \frac{b_M(\Delta W p + \Delta\beta)}{|M|}\right\}.$$

Proof. Translating a copy of M by h changes its overlap with a fixed copy by at most the area gained or lost by the translated copy. These two areas are equal. Along each line parallel to nonzero h, the convex section is an interval, and the one-sided loss is at most $\|h\|$. Integrating over the perpendicular projection gives the bound $b_M(h)$, also valid at $h = 0$. The bound is convex in h, being a maximum of absolute linear forms. On each closed slope support its composition with

the affine error map therefore attains a maximum at a vertex. Outside their union both responses vanish. Averaging the pointwise bound with the shared flux weights proves the claim. □

Since $M - M$ is a polygon, $b_M(h)$ is the maximum of $|\det(h, z)|$ over its finitely many extreme vertices. Thus neither the width functional nor the support test requires continuous optimization. For the unit square, $b_M(h) = |h_x| + |h_y|$, recovering Proposition 4.1. For polygons with more edge directions the number of support vertices can increase. The result covers rounding the affine map while retaining the template exactly. Perturbing its vertices changes the covariogram and requires a separate error bound.

Three constructed controls used the same affine record, rounded componentwise to the nearest multiple of 1/64. The unit square, a right triangle and an irregular pentagon had respectively 4, 6 and 10 vertices on each closed support. The exact rational certificates were 33/4000, 13/2000 and 21/4000. The largest absolute visibility changes on separate 21×21 grids were 0.0030534, 0.0020687 and 0.0016491. The grids illustrate slack at sampled points. The proof supplies each all-direction bound. Independent interval-product and triangle formulas agreed with polygon clipping at each of 289 rational shifts per shape. The supplement specifies the vertices, maps and grids.

7.3 Which stage information changes a finite-source prediction?

Sullivan derives a directional response from projected detector overlap and treats a three-element telescope [24, Sects. 3.2–3.3]. Li et al. project aperture functions for a convergent slit-slat SPECT collimator [3, Sect. 2]. For ideal thin parallel masks M_j at heights z_j above a common reference mouth M_0 ($z_0 = 0$), tracing from $x \in M_0$ to $x + z_j p$ gives

$$V(p) = \frac{\left|\bigcap_{j=0}^{m} (M_j - z_j p)\right|}{|M_0|}.$$

This formula also covers segments whose clear volumes between specified sections are their convex hulls. For a general taper the ray must clear every intervening section. A reference point in each projected mask must be the *same* point to transmit through the stack.

Use unit squares at $z = 0, 1$ and insert a centered rectangle at $z = 1/2$. The tall version has width 1/2 and height 1, the wide version width 1 and height 1/2. Both stacks share every mask area and centroid, all depths, the same reference axes and both end masks. Put $X = [1 - |p_x|]_+$, $Y = [1 - |p_y|]_+$. When both intervals have positive lengths their midpoint is the projected middle-mask center. The two fractions, also valid when either interval vanishes, are

$$V_{\mathrm{tall}} = \min(X, 1/2)Y, \qquad V_{\mathrm{wide}} = X \min(Y, 1/2).$$

At $p = (1/2, 0)$ they give 1/2 and 1/4. Throughout the 2° cone of Eq. (14), $3/8 < X < 3/5$ and $Y > 153/160$ by the component bounds in Appendix B. If $X \le 1/2$ the gap is $X(Y - 1/2) > 219/1280$. If $X \ge 1/2$ it is $(Y - X)/2 > 57/320$. A positive flux-weighted average preserves the smaller strict bound. Any single prediction from the shared record therefore errs by more than 219/2560 (8.55 percentage points) on at least one of these two stacks. This example isolates the missing middle-mask dimensions within the ideal direct-transport model. Exact polygon clipping reproduced each stack's directional formula at 289 rational slopes.

For polygonal stacks the set of accepted directions from a fixed starting point can also be a spherical polygon. This identifies where the clipped spherical-harmonic integral of Belcour et al. [11] can enter an angular calculation. A finite source with a small-circle boundary requires a separate treatment of that boundary, and whole-mouth flux still averages over starting positions. The calculations above use planar overlap.

7.4 Which parts of a circular taper determine direct transmission?

Let $x \in [a, b]$ increase from the entrance to the exit. Every normal section is a disk on the same axis, with positive continuous radius $r(x)$. A ray entering at transverse position $u \in \mathbb{R}^2$ has position $u + (x - a)p$ at station x. Here p points along propagation. Reversing its sign has no effect on the directional area for coaxial disks. Write $D_R(c)$ for the closed disk of radius R centered at c. The set of admissible entrance positions and its directional fraction are

$$K_r(p) = \bigcap_{a \le x \le b} D_{r(x)}(-(x - a)p), \qquad V_r(p) = \frac{|K_r(p)|}{\pi r(a)^2}. \tag{28}$$

This is the fixed-direction version of intersecting optical stops in position and direction [43]. Computing lines through a finite family of planar segments has related precedents [47]. Any admissible ray grazing a disk has its entrance point on the boundary of $K_r(p)$. A planar convex set has boundary of area zero (and has zero area if its interior is empty), so closed disks give the direct-power fraction for absorbing walls. Equation (7) averages V_r with the entrance flux weights.

Proposition 7.2 (An exact sufficient radius profile)**.** *Let g be the greatest convex function below r on $[a, b]$. Replacing r by g preserves the entire direct ray set. The two end disks alone give the correct directional area for every p if and only if $r(x)$ lies on or above the straight chord $\ell(x)$ joining its two end radii.*

Proof. The distance $f(x) = \|u + (x - a)p\|$ is convex by the triangle inequality, including for rays outside a meridional plane. Consequently $f \le r$ implies $f \le g$, and $g \le r$ proves the converse. For strict clearance, compactness gives $f + \eta \le r$ for some $\eta > 0$, and $f + \eta$ is also convex. The envelope retains both end radii: raising a convex minorant to the actual values at the two endpoints preserves convexity.

Every ray joining the end disks satisfies $f \le \ell$ by convexity. Conversely, suppose $r(x_*) < \ell(x_*)$ at an interior station. Choose a transverse unit vector e and join $r(a)e$ to $r(b)e$. Translate both endpoints a small distance toward the axis along $-e$. The endpoints now lie strictly inside their disks, while the ray at x_* still lies strictly beyond the wall. The same inequalities hold on an open set of nearby entrance positions at this fixed slope. Its positive area is missing from the end-disk prediction. □

For a specified piecewise-linear profile, the vertices of its lower convex hull therefore suffice. An outward bulge above the end-to-end frustum can store extra volume without changing even a partially transmitted beam. This is a sufficiency result for direct rays. Minimality of the retained profile and reconstruction from measured flux would each require a separate argument.

A smooth obstruction with matching volume. On $[1, 6]$ compare the straight frustum $\ell(x) = (7 - x)/6$ and the truncated Gabriel horn $r_H(x) = 1/x$. Let $t = \|p\| = \tan\theta$. When $t \le 1/6$ the back-projected exit fits inside the entrance since $5t + 1/6 \le 1$. The frustum transmits $1/36$ of entering power. This is the classical full-illumination limit for two circular openings [24, Sect. 3.2][42, Fig. 1]. The horn blocks every ray when $(7/2)t > 2/5 + 1/6$, since a ray cannot visit both disks at $x = 5/2$ and 6. Thus at $t = 0.165$, and throughout $17/105 < t \le 1/6$, the fractions are $1/36$ and zero.

The literal horn and frustum have different volumes. To match volume and retain the obstructing section, work in squared radius. Fix $h_0 = 1/10$. Let q_+ equal 1, then $4/25$, then $1/36$, making smooth decreasing transitions on the intervals $[5/2 - h_0, 5/2]$ and $[6 - h_0, 6]$. For each transition use $S(v) = \eta(v)/(\eta(v) + \eta(1 - v))$ at the rescaled station $v \in [0, 1]$, where $\eta(v) = e^{-1/v}$ for $v > 0$

and is zero otherwise. The transition from A to B is $A+(B-A)S(v)$. Since $S(v)+S(1-v)=1$, each transition averages its endpoint values. Hence

$$\int_1^6 q_+ \,\mathrm{d}x = \frac{3}{2} + \frac{7}{2}\frac{4}{25} - \frac{h_0}{2}\Big(1-\frac{1}{36}\Big) = \frac{7241}{3600}, \qquad \int_1^6 r_H^2 \,\mathrm{d}x = \frac{5}{6}, \qquad \int_1^6 \ell^2 \,\mathrm{d}x = \frac{215}{108}.$$

The last value lies strictly between the first two. In particular

$$r_*(x) = \sqrt{(1-\lambda)x^{-2} + \lambda q_+(x)}, \qquad \lambda = \frac{215/108 - 5/6}{7241/3600 - 5/6} = \frac{12500}{12723}. \tag{29}$$

This profile is smooth and strictly decreasing. It retains the horn's radii at $1, 5/2, 6$ and has the frustum's enclosed volume $\pi(215/108)$. Both channels transmit $1/36$ on axis. For any positive incident flux supported within the indicated oblique band their transmissions are $1/36$ and zero. A circular source of half-angle 0.08° centered at $\arctan(0.165)$ lies wholly in the band. The common record of end disks, axis, depth, volume and on-axis transmission therefore forces worst-case error at least $1/72$ (about 1.39 percentage points) by Eq. (17). Section order also affects transport in piecewise cylindrical Knudsen pores [44, Sect. 3.3], where the particles interact with walls. Here the obstruction concerns direct passage alone.

7.5 How closely must the bore radius be measured?

Suppose the end disks and their separation are known exactly. Let $\widehat{r}$ interpolate the reported radii linearly on stations $a = x_0 < \cdots < x_N = b$, including the known endpoint radii. Suppose the errors at these stations are at most ϵ, the largest spacing is h, and a bound $|r''(x)| \le M$ holds throughout each interval. Then

$$|\widehat{r}(x) - r(x)| \le \delta, \qquad \delta = \epsilon + Mh^2/8. \tag{30}$$

Indeed the interpolation of the station errors is bounded by ϵ. On $[x_i, x_{i+1}]$ the linear interpolant of the exact values differs from $r(x)$ by at most $M(x-x_i)(x_{i+1}-x)/2$, by the linear interpolation remainder. The curvature bound is an input assumption. Appendix B.3 constructs two smooth decreasing walls that match three exact station radii and force a worst-case error greater than one percentage point across a source of positive solid angle.

Proposition 7.3 (Transmission from bounded radius samples)**.** *Construct $K_\pm(p)$ by replacing r in Eq.* (28) *by $\widehat{r} \pm \delta$ and intersecting both results with the two exact projected end disks. A negative lower radius gives the empty lower set. Write $R_a = r(a)$, $R = \min(r(a), r(b))$ and $V_\pm = |K_\pm|/(\pi R_a^2)$. Then for every slope p,*

$$V_- \le V_r \le V_+, \qquad V_+ - V_- \le \frac{\min\{R^2,\, 4R\delta\}}{R_a^2}. \tag{31}$$

The same bounds hold after averaging with any shared nonnegative entrance flux weights of positive total mass.

Proof. Pointwise radius ordering nests the allowed ray sets. For either piecewise-linear envelope it suffices to intersect its disks at the stations: between two consecutive stations the norm of a ray is bounded by the chord of its two endpoint norms. If any lower radius is negative, a station already witnesses this. Both sets use the same entrance-area denominator.

Every entrance point whose closed disk of radius 2δ lies inside K_+ belongs to K_-. Indeed that disk lies inside every outer projected disk, so its center clears each outer radius by 2δ and satisfies the two exact end constraints. The area difference is thus bounded by the area of an inner boundary strip of width 2δ in the convex set K_+. For any planar convex body K, the strip of width w has

area at most w $\mathrm{Per}(K)$: by the coarea formula its area is the integral of the perimeters of the inner parallel sets up to w or the inradius, and those perimeters are nonincreasing [48, Sect. 2.2]. The exact mouth constraints place K_+ in a disk of radius R, so its perimeter is at most $2\pi R$ by the planar width formula. Its area is at most πR^2. If K_+ is empty or has empty interior both fluxes already vanish. These bounds also cover an empty K_- and hold under the stated normalized average. □

For a target absolute error τ, the midpoint of the flux interval is within τ whenever the computed interval has width at most 2τ. A sufficient design rule independent of direction or source shape is $\epsilon + Mh^2/8 \le \tau R_a^2/(2R)$. With $M > 0$ and positive error allowance after measurement noise, this gives

$$h \le \sqrt{\frac{8}{M}\left(\frac{\tau R_a^2}{2R} - \epsilon\right)}.$$

A case-specific disk intersection can give a tighter bound. A disjoint or externally tangent pair of projected disks certifies zero power at one slope. Certifying partial transmission instead requires bounding the common *area*. These estimates address radial uncertainty at registered stations with known centers and illumination. Numerical errors in the area and angular integrals must also be included when reporting a computed interval.

7.6 Partial transmission from synthetic bore measurements

We tested the measurement rule where the visible part of the exit changes with direction. The exact tube profiles generate a set of synthetic measurements and provide a separate numerical comparison. The prediction uses the radii reported at the stations, the exact end disks, the curvature bound and the source. These are designed examples for testing the bounds.

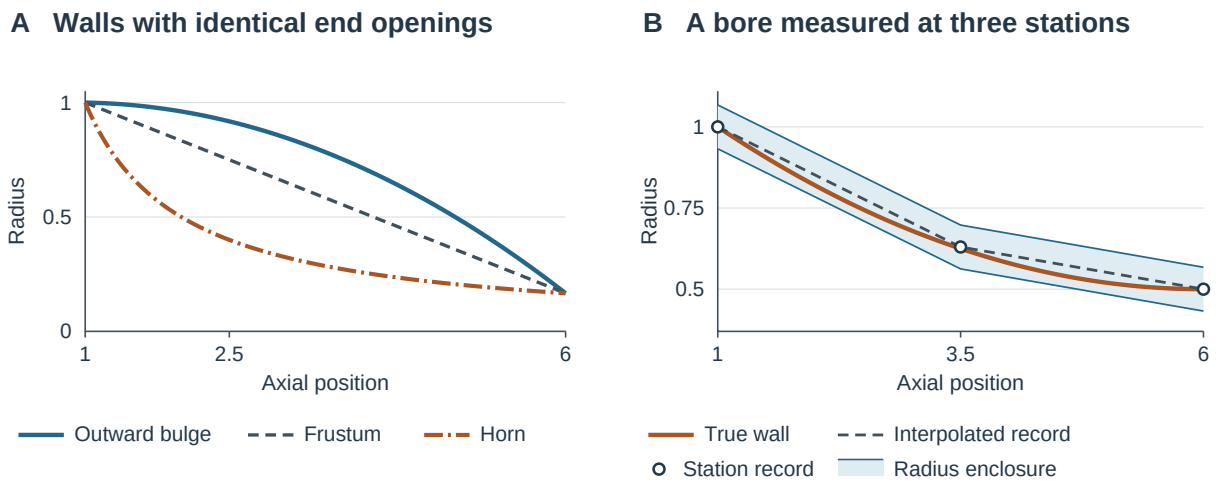


Figure 6: The left panel illustrates the end-disk criterion for exit radius 1/6. An outward bulge preserves the end-to-end rays, while the $1/x$ horn dips below the chord. On the right, three recorded radii of the broad inward taper are $1, 0.630, 0.5$. The generating middle radius is 0.625. The shaded band is the interpolated survey radius plus or minus 0.0675, using $\epsilon = 0.005$, $M = 0.08$ and station spacing 2.5. Both exact end disks are imposed separately. The vertical scales differ.

All fixtures span $1 \le x \le 6$ with entrance radius one. The three broad tubes share exit radius $1/2$ and the bound $M = 0.08$. The fourth tube is the literal $1/x$ horn with exit radius $1/6$ and $M = 2$. At each interior station the reported radius differs from the true radius by at most $\epsilon = 0.005$. The experiment refines the same measurements on nested uniformly spaced grids. Each source is a constant-radiance annulus of directions, uniform in azimuth, with slope magnitude $s = \tan\theta$

in $[0.16, 0.20]$ for the broad tubes and $[0.10, 0.13]$ for the horn. The annulus has nonzero solid angle. These trials include partial throughput instead of relying on the complete extinction in Section 7.4.

For a disk of projected center $-(x_i - a)s$ the horizontal chord at height y has half-width $\sqrt{r_i^2 - y^2}$ when $|y| \le r_i$. Intersect the chord intervals and integrate their common width over y to obtain the accepted area. This width is nonincreasing in $|y|$. At a fixed y, two chords of half-widths h_i, h_j give a common-width constraint $h_i + h_j - |x_i - x_j|s$, including the self-pairs $i = j$. Taking the positive part of the minimum over pairs gives the intersection length, so the area also decreases with s. These two monotonicities allow lower and upper sums in both variables. The projected-flux weight of a slope bin $[s_0, s_1]$ is proportional to $(1 + s_0^2)^{-1} - (1 + s_1^2)^{-1}$. The normalization uses the exact entrance disk and total incident flux.

For a one-percentage-point target, the analytic rule alone gives nine stations for all three broad profiles and seventeen for the horn on the tested dyadic grids. For example, the radius bound for the broad tubes with N equal intervals is $0.005 + 0.25/N^2$. Setting $N = 8$ limits the midpoint's ideal geometric error to 0.00890625 of entering power. The theorem applies in every illumination direction. The tighter intervals below use the specified source and include numerical integration error.

Table 10: First tested grid whose interval midpoint is certified within one percentage point. Tested station counts were $3, 5, 9, 17, 33, 65$, including the two mouths. Power intervals are percentages of entering power, rounded outward. Error bounds use exact endpoints before display rounding. The reference column is a separate numerical estimate using the generating wall.

	Stations required on tested grids		First source-specific certificate		Reference
Profile	Data and source	Uniform rule	Power interval (%)	Error (pp)	Power (%)
Broad frustum	5	9	$[12.520, 14.275]$	≤ 0.877	14.207
Outward bulge	5	9	$[12.520, 14.275]$	≤ 0.877	14.207
Inward quadratic	5	9	$[12.280, 14.230]$	≤ 0.975	13.748
$1/x$ horn	17	17	$[0.465, 1.373]$	≤ 0.454	1.102

With $z = (x - 1)/5$ the broad frustum has radius $1 - z/2$. The other broad profiles add $(2/5)z(1 - z)$ or subtract $z(1 - z)/2$. The frustum and outward bulge give identical true responses by Proposition 7.2. The inward quadratic reduces transmission under the same source. For that tube three stations gave a midpoint bound of 2.574 percentage points and five reduced it to 0.975. For the horn nine stations still allowed 1.080 points. Seventeen reduced the bound to 0.454. Exceeding the target in a coarse bound leaves that grid undecided.

The data-specific intervals use 512 transverse strips, 64 angular bins, directed integer square roots and rational flux weights. A separate calculation uses the two-disk lens formula for the frustum and bulge. For the inward taper and horn it minimizes the squared transverse clearance over the continuous profile, then integrates chords in the other transverse coordinate. Every reference value lies inside its reported interval at all six resolutions. A separate replay integrated active circle arcs on a 4,096-interval wall with 48 angular Gauss–Legendre nodes. Its four nominal values agreed with the continuous-profile or lens calculations to within 0.0000033 percentage points. The reference calculations are numerical cross-checks. The error guarantee follows from the inner and outer inequalities and directed quadrature. The station counts are conditional on the declared shape class, measurement tolerance and source.

8 An audit of isotropic quadrature on the archived panels

Under isotropic illumination, transmission is the diffuse form factor between the two mouths [4]. How far did the failure in Table 4 extend? We took the isotropic query for every geometry on both archived 72-prism panels. For each prism we enclosed the ideal transmission twice, once for the original geometry and once for its decoded binary16 record. This retrospective calculation uses the same archived fixtures, including the known failure. It supplies an error test for each returned value independently of agreement between quadrature estimates.

8.1 A separate bound for each output

Jiao et al. [26] derive compression bounds for specified quantities of interest. Here the quantity is a ratio of transmitted to unblocked flux. For any scalar quantity of interest J, choose the admissible inputs and a common error unit. Trace the observable from the intended input to the specified input, its decoded representation, its quadrature and its actual floating-point evaluation. Inserting these intermediate values gives Eq. (26) by the triangle inequality, provided that all four bounds hold on the same admissible case in the same output units. If normalization depends on an uncertain denominator, its joint bounds must exclude zero. For example, if a target ratio N/D and its approximation use $D, \widehat{D} \ge d_* > 0$, $|N - \widehat{N}| \le e_N$, $|D - \widehat{D}| \le e_D$ and $|\widehat{N}/\widehat{D}| \le B$, then

$$|N/D - \widehat{N}/\widehat{D}| \le (e_N + B e_D)/d_*.$$

This follows by adding and subtracting $\widehat{N}/D$. It explains why a certified numerator alone cannot settle a normalized daylight query near zero incident power. A narrowing sequence of estimates alone supplies no bound on the remaining integral error. The Community Codec preview verifies complete restored STEP bytes [33]. Here distinct byte streams may be admissible whenever the change in the specified observable is bounded.

For this test the original geometry is the exact binary rational represented by the archived input. For the decoded test, the half-precision map and offset are read from the archived bytes. Write G for the two-dimensional mouth matrix, t for its transverse offset and $d > 0$ for its depth. For a decoded map we may set $d = 1$, $G = W^{-1}$ and $t = -G\beta$. Put $K = |\det G|/(\pi d^2)$. The isotropic integral is a tent-weighted average of $(1 + \|(Gq + t)/d\|^2)^{-2}$ multiplied by K. After subdivision into n_1 by n_2 cells per quadrant, use the exact tent-weighted centroid in each cell. Appendix I bounds the Hessian norm by four and the variance in each coordinate by $1/(12n_j^2)$. The absolute Taylor remainder is bounded by

$$E = \frac{K}{6} \sum_{j=1}^{2} \frac{\|G_{:j}/d\|^2}{n_j^2}.$$

Starting at $n_1 = n_2 = 1$, double the count in the coordinate with the larger contribution until $E \le 1/4000$ or the grid reaches 2^{24} cells. Rational input arithmetic, bounds on π, outward-rounded centroid evaluations and summations give a closed interval $[l, u]$ for each ideal result. The floating-point assumptions are stated in Appendix I. For a saved output z its error is bounded below by $\max(0, l - z, z - u)$ and above by $\max(|z - l|, |z - u|)$. Both endpoints matter when deciding whether a tolerance has been met.

All 288 enclosures met the remainder target in 163,488,128 cells altogether. For the analytic calculation we used the archived R10 primitive, integrated the coordinate corresponding to the larger column norm of G/d analytically, and applied sixteen Gauss nodes on each half of the other axis. Each analytic estimate was evaluated on the same geometry used for its enclosure. The largest certified error upper bounds across the panels were 0.000306 for both original and decoded records (rounded up). Replication g66 was the only archived isotropic failure at either 0.001 or 0.01. For its archived return of 0.16881765, the original-geometry absolute error exceeded 0.02165, and the error against the decoded geometry exceeded 0.02169. The latter comparison

Table 11: Retrospective isotropic decisions at absolute transmission tolerance 0.001. Each row covers the same 144 prisms. The original-geometry comparison includes representation error in the archived binary16 calculation. The decoded-geometry comparison isolates the solver on its stored input.

Returned value	Enclosed geometry	Within	Outside	Unresolved
Archived adaptive binary16 calculation	Original	143	1	0
Archived adaptive binary16 calculation	Decoded	143	1	0
Analytic inner coordinate, order 16	Original	144	0	0
Analytic inner coordinate, order 16	Decoded	144	0	0

isolates integration and evaluation on the stored record. The earlier mixture failure follows from the same broad component. This audit covers only the isotropic queries.

8.2 A shallow square exposes the remaining boundary layer

For an aligned square of side a and depth d, put $\eta = d/a$ and use isotropic front-hemisphere illumination. Since $V \geq 1 - \eta(|p_x| + |p_y|)$ and each absolute slope has expectation one under $\mu\,\mathrm{d}\Omega/\pi$,

$$\max(0, 1 - 2\eta) \leq T_{\mathrm{sq}}(\eta) \leq 1.$$

The primitive in Appendix J.2 gives the exact reduction

$$T_{\mathrm{sq}}(\eta) = \frac{2}{\pi}\int_0^1 (1-x)\frac{\eta^2}{(\eta^2+x^2)^{3/2}}\arctan\frac{1}{\sqrt{\eta^2+x^2}}\,\mathrm{d}x. \tag{32}$$

Its closed form, consistent with the classical coaxial-square diffuse form factor [4], is

$$\begin{aligned} T_{\mathrm{sq}}(\eta) &= \frac{4}{\pi}\left[\sqrt{1+\eta^2}\arctan\frac{1}{\sqrt{1+\eta^2}} - \eta\arctan\frac{1}{\eta} + \frac{\eta^2}{4}\log\frac{(1+\eta^2)^2}{\eta^2(2+\eta^2)}\right] \\ &= 1 - \frac{4}{\pi}\eta + O\left(\eta^2\log\frac{1}{\eta}\right), \qquad \eta \to 0^+. \end{aligned} \tag{33}$$

The outer integrand concentrates unit mass at $x = O(\eta)$, with peak asymptotic to $1/\eta$. A 16-point Gauss–Legendre rule on $[0, 1]$ has first node about 0.0052995. In the constructed sweep $a/d = 4, 8, \ldots, 1024$, at $a/d = 128$ the rule returns 1.060545 versus exact 0.984639. At $a/d = 512$ the peak lies inside the first node: the rule returns 0.290870, below the rigorous lower bound 0.99609375 and exact value 0.99611361. Its error exceeds 0.70522.

Thus an exact inner integral does not ensure a resolved outer integral. For this isotropic square, the midpoint of $[1 - 2\eta, 1]$ gives certified error at most η when $\eta \leq 1/2$; other shapes and sources need their own bounds. R13 checks compare the marginal and closed form at selected scales, supplementing the archived sweep without claiming a general solver guarantee.

9 Scientific interpretation and scope

9.1 Relation to existing methods

Projected overlap in collimators and particle telescopes [3, 24], diffuse form factors [4], response moments [7], low-order irradiance [6], wavelet transfer [9] and deformable transfer [8] are established. Our test isolates an explicitly omitted coordinate, without claiming prior systems generally omit geometry. EnergyPlus retains dimensions and vertices [1]; Radiance can assign reveal information to different transport matrices [2]. Such operators count as retained inputs even behind a short interface.

The contribution combines finite-source error floors, fixed-pose rounding certificates and tests separating representation from evaluation. Quantity-controlled compression and error-driven refinement have broader precedents [26, 25]. Clipped-domain SH [11, 23] illustrates why retaining a boundary differs from truncating a kernel. Fenestration BSDFs describe a broader operator than scalar total flux, with application-dependent angular resolution [12, 5, 17, 18]; separating a direct peak has practical precedent [14].

For circular channels, the full-illumination and extinction angles of a straight frustum are established collimator results [42, 24]. Solar radiometry weights the source by its angular penumbra function [45], and baffle design distinguishes the unobscured field, direct-light exclusion and scattered light [46]. Phase-space stop intersection [43] and planar line-transversal algorithms [47] also precede the present work. Our measured-bore result specifies which shape data, interpolation assumption and numerical enclosure jointly support an absolute error bound for the whole exit. The one-percentage-point criterion and the uniform-over-entrance illumination are part of that claim.

9.2 Which angular fields can an aperture array distinguish?

A shape-specific response can also define a measurement operator. Let $D = [-1, 1]^2$ in the slope coordinates of Eq. (22). For channels sharing a common mouth normal, set $V_k(p) = F(W_k p + \beta_k)$ and $w(p) = (1 + \|p\|^2)^{-2}$, the density of projected solid angle. A detector calibrated to read one under unit radiance returns the linear quantity

$$y_k(L) = \frac{\int_D L(p) V_k(p) w(p)\, dp}{\int_D V_k(p) w(p)\, dp}.$$

Calibration fixes the denominator independently of the unknown field. When $\beta_k = 0$, both w and V_k are even under $p \mapsto -p$, a half-turn in azimuth about the common mouth normal. Changing aspect ratio, rotating mouths in plane, or shearing them preserves this symmetry. Every odd component cancels by substitution. For example, the positive fields $1 + 0.4p_1$ and $1 - 0.4p_1$ produce identical readings. Angularly diverse detector responses have also been constructed with diffraction gratings [28]. The calculation here uses geometric overlap.

In a synthetic test, twenty seeded arrays each comprised 32 rotated rectangular channels. Their acceptance centers were $c_k = -W_k^{-1}\beta_k = \lambda d_k$, with d_k drawn uniformly from a slope disk of radius 0.57. The unknown fields used the fifteen orthonormal product Legendre modes of total degree through four on D (uniform square measure). Six modes have odd total degree. We estimate all coefficients with unconstrained truncated-SVD least squares, retaining singular values above 10^{-10} times the largest. Fits have no positivity constraint or additional regularization. Every centered design had numerical rank nine. At every nonzero scale tested ($\lambda = 0.01, 0.03, 0.1, 0.3, 0.6, 1$), each design had rank fifteen at a relative singular-value cutoff of 10^{-10}.

Offset scale	Rank	κ_{retained}	Field error	Held-out RMS
0	9	73,001	5.31	0.249
0.01	15	107,635	10.76	0.835
0.60	15	3,045	0.253	0.0102
1.00	15	334	0.0364	0.00117

Table 12: Synthetic aperture inversion with known channel calibration and additive Gaussian noise of standard deviation 0.001 in unit-radiance response units. Conditioning is the median across designs on the retained subspace. Relative RMS field error includes the unit background. The last column is RMS prediction error on sixteen independently drawn full-offset channels, in unit-radiance response units. Errors are medians across 1,280 fields per scale.

At offset 0.01 retained rank rises from nine to fifteen, but median field error increases from 5.31 to 10.76 and held-out RMS from 0.249 to 0.835. Larger tested offsets improve conditioning and prediction. These are finite-scale results for the stated noise and estimator, not an asymptotic noise law. An upper bound on small singular values alone cannot upper-bound reciprocal variance; that requires lower bounds and nondegeneracy.

At full offsets median field error is still 43.6% of the true field's RMS angular variation (10th–90th percentiles 22.1–82.7%). Known gains, ideal geometry and the polynomial class are assumptions. All kernels vanish outside radius $0.57 + \sqrt{0.34^2 + 0.24^2} < 0.987$; a positive bump in a corner of D preserves every reading. Arbitrary continuous skies remain unidentifiable.

9.3 Daylight and compact-record implications

Whole-exit agreement cannot establish restricted-view accuracy: modeled power may spread beyond the receiver. The area bound is necessary, not sufficient; merged leaves and the finest lattice both matter. Here 2.69 times the compressed storage recovers two matched-view cases, while eighteen still miss on the unmerged depth-ten lattice. These are designed-panel outcomes, not building-wide error rates.

A Community Codec 0.3.0-preview replay [40] treats six EnergyPlus example openings as straight black recesses. The host supplies pose, units and opening association. Six GX01 records use 49 bytes under a 0.010 directional representation budget. Across 11,993 eligible Chicago typical-year opening–hour queries with front-facing 0.266° sun disks, a conservative support test skips 475 integrations; the largest sampled original–decoded difference is 0.29534 percentage points. Storage excludes building geometry, weather, pose and scattering. This is a representation replay, not an enclosure of absolute solar-integral accuracy or measured building performance.

9.4 Physical interpretation and next evidence

Verification and physical validation differ [19]; Radiance comparisons with calculations and measurements illustrate both [15, 16]. The open bench protocol [41] proposes a separate square–rectangle flux comparison. No physical readings are included. As-built geometry, source distribution, reflectance, detector response, collection losses and joint uncertainties must be characterized before testing a physical prediction.

For a tapered bore, axial station locations, the two mouth radii and the curvature bound enter the guarantee explicitly. A finite set of readings cannot itself verify a curvature limit between readings. That input would have to come from an independent inspection, a fabrication specification with justified margins, or a declared profile family. Source misalignment and residual wall reflection would enter a separate physical uncertainty budget.

Radiometry applications require their own field, spectrum and noise models. Moonlight studies [34, 35] and night-sky mapping [36] have different targets. A cutoff luminaire additionally needs source position and outgoing scattering; no biological response such as that studied by Fabian et al. [37] is inferred here.

9.5 Conclusion

An aperture description can omit substantial parts of a wall and still predict every direct ray. For coaxial circular sections the lower convex envelope is one such description. End openings and volume alone can miss complete extinction over a finite source. With bounded curvature and radius error, finite station readings bound the intervening wall, then the partly visible exit and its transmitted power. A calculation can return the width of that interval along with the estimate. The affine, quadrature and restricted-view examples show why a geometric guarantee must be

carried through to the particular optical quantity being reported. All performance claims here refer to the declared models, sources and computational checks.

Author contributions and AI assistance

Mostafa Nasr set the research questions, direction and revision objectives and is responsible for the claims. Through R10, he supplied Eqs. (6)–(7) and (8)–(9) for argument development, reviewed drafts and checked other mathematics, mostly through individual SymPy or MATLAB calculations. This historical account does not establish his completed review of every later addition.

AI tools generated most prose, computational and figure code, and many derivation steps. Tools through R10 included OpenAI GPT-5 variants, ChatGPT/Codex, Claude variants, Mistral, DeepSeek and Cursor; exact allocation cannot be reconstructed. The author's estimate of roughly a thousand read–review–revise cycles is retrospective, not logged. Complete transcripts were not preserved. Codex prepared later text, derivations and computations under his direction. For R13 it reconstructed source, condensed exposition and ran the stated checks. For R14, Codex carried out the literature comparison, drafted the measured-taper arguments and implemented and checked the reported synthetic fixtures. Author review of the new material is pending. This assistance is distinct from independent peer review. The supplement supports archived computations; not every analytical statement has a separate executable verifier.

Data availability and revision history

The R8 deposit preserves benchmarks and diagnostics; R9 preserves records, finite-source and F1 planning studies; R10 preserves further representation/integration analysis. The R11 deposit supplies its supplement and source. The corrected R12 deposit carries those archives unchanged, without editable R12 source. Appendix K maps the evidence and R13 reconstruction. R14 uses the published R13 PDF and source package as its editorial and editable parent. The R14 review package contains the editable manuscript, new computational checks and their results; the series record tracks public editions. No physical measurements are included. Original prose, figures, data and documentation are CC BY 4.0; original code is MIT; third-party terms remain in force.

A Details behind the information test

A.1 Equal records, equal kernels

The forward implication of Proposition 3.2 follows from identical inputs. For the reverse implication, assign $F(D, L)$ the common response of the geometries sharing D. The criterion ensures that the assignment is well defined.

If integrable kernels agree almost everywhere, their integrals agree for every bounded nonnegative field. Conversely, suppose two kernels differ on a set of positive solid-angle measure. At least one of the sets on which their difference is strictly positive or strictly negative has positive measure. Its indicator is an allowed nonnegative illumination and separates the two responses. This proves kernel equality almost everywhere when the lighting family includes all such fields.

For a restricted lighting family spanning a q-dimensional subspace of $L^2(S^2)$, with kernels in L^2, equality of all allowed responses extends by linearity to the span. It is therefore equivalent to equality of the q moments in Eq. (21). This gives an existence criterion, not a lower bound on the size or computational cost of every possible encoding.

A.2 Why scalar exactness is a separate condition

For a fixed geometry, the scalar rule

$$\widehat{\mathcal{R}}_{\mathrm{scalar}} = \kappa_0 P_0 \tag{A.1}$$

is exact for all bounded nonnegative fields if and only if

$$H_g(s) = A\kappa_0[n \cdot s]_+ \quad \text{almost everywhere.} \tag{A.2}$$

Apply the preceding indicator-field argument to the difference of these kernels. On the front hemisphere this says $V_g = \kappa_0$ almost everywhere.

To see why this is not a test of geometric sufficiency, consider the fixed known law $H = A[n \cdot s]_+^2$. Its geometry is completely specified by (n, A), although it is not a scalar multiple of the cosine kernel. At $A = 1, n = \mathbf{e}_z$, isotropic calibration gives $\kappa_0 = 2/3$. Under $L = 1 + 0.8s_z$, the true response is $16\pi/15$ while the scalar gives $46\pi/45$, a relative error of $-1/24$. The scalar formula fails, but a calculation evaluating the known law still has sufficient geometry inputs.

A.3 Error floors and missing coordinates

For two responses $r_1 < r_2$, the triangle inequality gives

$$r_2 - r_1 \le |y - r_1| + |y - r_2|,$$

which proves Eq. (17); the midpoint attains equality. For positive responses, an optimum for the worst relative error lies between them. Equating $(y - r_1)/r_1$ and $(r_2 - y)/r_2$ gives the harmonic-mean optimum and Eq. (18).

For the ordered frame of Eq. (2), fixed (Q, A, d) and normal extrusion leave two continuous shape coordinates in Eq. (3): a_0 and c_0, with $b_0 = A/a_0$. Rectangles set $c_0 = 0$, leaving one aspect-ratio coordinate. An oblique extrusion $v = dn + t_1x + t_2y$ adds two offset coordinates. These are coordinate counts for the chosen family, not minimum descriptor-size theorems. Mouth translation is undetectable by the stated total-power observable under spatially uniform illumination.

B Finite-source separation without numerical quadrature

These calculations establish the information limit using bounded sources, independently of numerical integration accuracy.

B.1 Square–rectangle pair

On a 2° cone about $s_* = (1, 0, 2)/\sqrt{5}$,

$$\|s - s_*\|_2 \le 2\sin(\pi/180) < \pi/90 < 7/200.$$

Using $4/9 < 1/\sqrt{5} < 9/20$ gives

$$2/5 < s_x < 1/2, \qquad 4/5 < s_z < 1, \qquad |s_y| < 7/200.$$

Put $p_x = s_x/s_z$ and $p_y = |s_y|/s_z$. Then $2/5 < p_x < 5/8$ and $0 \le p_y < 7/160$. All overlap factors are positive, so

$$\begin{aligned} V_{\text{rect}} - V_{\text{sq}} &= (1 - p_x/2)(1 - 2p_y) - (1 - p_x)(1 - p_y) \\ &= p_x/2 - p_y > 5/32. \end{aligned} \tag{B.1}$$

A positive cosine-weighted average preserves the strict bound, proving Eq. (15). In particular, a bound on s_x alone must not be mistaken for the same bound on its slope s_x/s_z.

For the exact gap, multiply the visibility difference by s_z. The numerator is the cap integral of $s_x/2 - |s_y|$, and the denominator is the cap integral of s_z. Axial symmetry gives

$$\int_{\mathcal{C}_\alpha} s \, \mathrm{d}\Omega = \pi \sin^2 \alpha \, s_*.$$

Because the y axis is perpendicular to s_*, cap polar coordinates give

$$\int_{\mathcal{C}_\alpha} |s_y| \, \mathrm{d}\Omega = 4 \int_0^\alpha \sin^2 \gamma \, \mathrm{d}\gamma = 2\alpha - \sin(2\alpha).$$

Substitution yields Eq. (16). The source normalization cancels from the ratio.

The explanatory companion calculation, chosen after the original benchmark, estimates transmissions 0.49171954 and 0.72515358 at polar order 24 and 384 azimuths. The largest separate polar and azimuth changes are 2.28×10^{-7} and 1.67×10^{-6}. These are empirical refinement changes, not enclosures or proof of every reported digit. The companion is not a new held-out benchmark.

B.2 Sheared pair

Set $\ell = 1$ in Eq. (19); normalized transmission is scale invariant. The source axis is $s_0 = (4, 2, 5)/(3\sqrt{5})$. Throughout its 2° cone the overlap factors stay positive, and

$$H_{0.3}(s) - H_0(s) = 0.3 s_y \left(1 - \frac{s_y}{s_z}\right). \tag{B.2}$$

The cone radius is less than 0.035. It gives $s_y > 0.26$, $s_z - s_y > 0.39$, $s_z < 0.79$, and also $s_x - 0.3 s_y > 0$ and $s_z - s_x > 0$. For example, $\sqrt{5} < 9/4$, $\sqrt{2} < 10/7$ and $2 \sin 1^\circ < 7/200$ give the conservative rational margins used here.

For a uniform cap of unit angular integral, the unblocked response is

$$P_0 = (s_0)_z \frac{1 + \cos\alpha}{2} < 0.75. \tag{B.3}$$

Averaging Eq. (B.2) therefore gives

$$T_{0.3} - T_0 > \frac{0.3(0.26)(0.39)}{(0.79)(0.75)} > 0.05. \tag{B.4}$$

This establishes a worst common-prediction error above 0.025.

Numerical source-centred integration gives powers 0.0895463863 and 0.1430533968 with $P_0 = 0.7451289671$, hence the transmissions quoted in the main text. Polar orders 12, 24 and 48 and azimuth counts 48, 96 and 192 were varied separately. The largest change was 1.39×10^{-17} in angular-response units. The resulting absolute minimax estimate is 3.59045 percentage points and the relative floor estimate is 23.0039%; these numerical estimates are distinct from the analytic inequalities above.

B.3 A finite-source obstruction between three measured stations

On $[0, 1]$ let $f(x) = 1 - x/2$ and use the smooth step S from Section 7.4, extended by 0 and 1 outside the unit interval. Put $k = 1/32$, $A = 15/64$, and

$$h(x) = 1 - kx - AS(x/k) - AS((x - 1/2)/k).$$

The two radii are positive and strictly decreasing and agree at $x = 0, 1/2, 1$. At $x = k$ the second wall has radius $783/1024$.

Take any slope within Euclidean distance $1/128$ of $p_0 = (-1/2, 0)$ in the transverse slope plane. Consider exit coordinates in the rectangle $19/64 \le y \le 7/16$, $|z| \le 15/64$. Its corners lie inside the exit disk. Back-projection at the central slope gives an entrance radius at most $15\sqrt{17}/64$. Adding $1/128$ for the worst slope perturbation leaves it strictly below one. Hence the whole rectangle passes the frustum at every slope in this disk. At $x = k$, the smallest possible y is

at least $19/64 + 31/64 - (31/32)/128$, strictly larger than $783/1024$, so none of these rays pass the second tube. All rays admitted by the second tube are also admitted by the frustum by the end-disk criterion.

The rectangle has area $135/2048$. Dividing by the unit-radius entrance area gives a gap of at least $135/(2048\pi)$ throughout the slope disk. Any positive incident flux supported there preserves this gap. An estimate based only on the three shared station radii, depths and axis therefore has worst-case error at least $135/(4096\pi)$, about 1.05 percentage points. This example assumes qualitative smoothness and monotonicity. The contrasting upper bound in Proposition 7.3 uses quantitative curvature information.

C The complete response and shape identifiability

For a compact convex planar mouth M with nonempty interior, define its *covariogram*

$$c_M(z) = |M \cap (M + z)|. \tag{C.1}$$

This established overlap function measures how much of a shape overlaps a translated copy [13]. Under the same straight absorbing-channel assumptions,

$$V_M(s) = \frac{c_M(\delta(s))}{|M|}, \qquad H_M(s) = \mu c_M(\delta(s)), \qquad \mu > 0. \tag{C.2}$$

Convexity is essential: for a nonconvex channel, visible endpoints do not necessarily imply a clear intervening ray.

Corollary C.1 (Application of planar covariogram uniqueness)**.** *Fix a common normal and a known normal extrusion $v = dn$, with $d > 0$. Among compact convex planar mouths with nonempty interior, equal ideal total-power responses for all bounded nonnegative angular fields imply equality up to translation or reflection in a point. The same conclusion holds for equality of normalized transmissions under every such field with positive incident front flux, even if the mouth areas are initially unknown.*

Proof. Equality of all total-power responses gives equal kernels almost everywhere. Covariograms are continuous, so the kernels are continuous on the open front hemisphere. Any pointwise difference there would persist on a neighborhood of positive measure; thus they agree everywhere on that hemisphere.

The displacement map is onto the entire mouth plane: for any $z \in n^{\perp}$,

$$s(z) = \frac{z + dn}{\sqrt{\|z\|_2^2 + d^2}}, \qquad \mu = \frac{d}{\sqrt{\|z\|_2^2 + d^2}} > 0. \tag{C.3}$$

Dividing the equal kernels by this common positive cosine gives equal covariograms. The mouths are planar convex bodies, so Theorem 1.1 of Averkov and Bianchi [13] gives equality up to translation or reflection in a point.

For normalized transmission, write $A_M = |M|$. The common incident-flux denominator leaves

$$T(M, L) = \frac{\int_{\mu>0} L(s)\mu\, c_M(\delta(s))/A_M \,\mathrm{d}\Omega}{\int_{\mu>0} L(s)\mu \,\mathrm{d}\Omega}.$$

The same indicator-field, continuity and inverse-map arguments give equality of c_M/A_M everywhere. By Tonelli's theorem,

$$\int_{\mathbb{R}^2} c_M(z)\,\mathrm{d}z = A_M^2, \qquad \int_{\mathbb{R}^2} \frac{c_M(z)}{A_M}\,\mathrm{d}z = A_M.$$

Thus the normalized covariogram recovers area, after which the covariograms themselves agree and the same uniqueness theorem applies. Conversely, translation and reflection in a point preserve the covariogram and both response types. □

This is an application of an established theorem, not a new uniqueness result. It needs the entire illumination family and known positive depth. Joint scaling of mouth and depth preserves all normalized transmissions. Finite noisy data do not establish this identifiability, stable inversion or a minimum encoding.

D Representation bounds, normalization and grazing directions

D.1 Proof of the support-corner bound

Let $f_i(z) = [1 - |z|]_+$. Each f_i is 1-Lipschitz and lies in $[0, 1]$. Subtracting two products one factor at a time gives

$$\begin{aligned} |F(q) - F(\widetilde{q})| &\le |f_1(q_1) - f_1(\widetilde{q}_1)| + |f_2(q_2) - f_2(\widetilde{q}_2)| \\ &\le |q_1 - \widetilde{q}_1| + |q_2 - \widetilde{q}_2|. \end{aligned} \tag{D.1}$$

With $q = Wp + \beta$ and $\widetilde{q} = \widetilde{W}p + \widetilde{\beta}$, this is the affine-error norm in Eq. (25). Outside the union of the two support parallelograms, both visibilities vanish. On each parallelogram, the norm is convex. Every interior point is a convex combination of the vertices, so its value is no greater than the largest vertex value. The pointwise difference is also at most one. Averaging under the probability weights of Eq. (23) proves the bound.

The basis E, normal and illumination must be shared by both records. Singular decoded maps and overflows are rejected before the test. The exact-rational computation in the study compares archived binary input values, not uncertain manufactured dimensions. The supplement names the planar matrix G and offset β as `B` and `b`; the paper reserves B for the three-dimensional edge matrix.

D.2 Why visibility error is the natural normalized error

Suppose an approximation $\widehat{V}$ retains exact cosine weighting and the exact unblocked denominator. Then

$$\widehat{T} - T = \int (\widehat{V} - V)\, \mathrm{d}\nu_L. \tag{D.2}$$

For bounded measurable visibilities and all bounded nonnegative fields with positive front flux,

$$\sup_L |\widehat{T} - T| = \operatorname*{ess\,sup}_{\mu > 0} |\widehat{V} - V|. \tag{D.3}$$

The upper bound follows from averaging. For the reverse bound, choose a number below the essential supremum and an indicator field on a positive-measure set where the signed error exceeds it in magnitude with one sign. That set has positive cosine-weighted measure because $\mu > 0$. Its weighted average exceeds the chosen number. Letting the number approach the essential supremum proves equality. Restricting the lighting family can reduce the worst error.

The exact domain is $\mu > 0$, whereas the implementation returns zero for $\mu \le 10^{-12}$. To compare it with the exact observable, include this cutoff in the effective approximate visibility but retain the uncut P_0. Cutting the denominator too defines a different observable. For a flat opening, light supported entirely in $0 < \mu \le 10^{-12}$ has true $T = 1$ and implemented $T = 0$; such fields are outside the saved source matrix.

For a positive-depth mouth of diameter D and tangential offset $t = v - dn$, visibility vanishes if $\|\delta\| > D$. Since

$$\|\delta\|_2 \geq \frac{d\sqrt{1-\mu^2}}{\mu} - \|t\|_2 ,$$

it is zero throughout

$$0 < \mu < \frac{d}{\sqrt{d^2 + (D + \|t\|_2)^2}}. \tag{D.4}$$

A cutoff strictly below this limit removes no transmitted power. The smallest such limit among the eighteen positive-depth shapes in the original matrix exceeds 0.17. There is no uniform collar for a family with arbitrarily small depth-to-diameter ratio, and the two flat controls have none.

If a method approximates $K = H/A = \mu V$ directly, its corresponding front-hemisphere normalized error involves $|\widehat{K} - K|/\mu$. This can grow near grazing. On front support with $\mu \geq \mu_{\min} > 0$, a uniform kernel error ϵ instead gives transmission error at most $\epsilon/\mu_{\min}$. Raw SH can also leak onto the rear hemisphere; arbitrary rear illumination can then cause unbounded normalized error while front flux stays fixed. The tested finite cones do not cross the horizon. Positive interpolation of visibility preserves $[0, 1]$ before integration error; raw SH predictions are left unclipped.

D.3 What a harmonic truncation bound says

Let P_b be orthogonal projection onto real SH through degree b. For $L, K_g \in L^2(S^2)$, an exact coefficient contraction is

$$\mathcal{R}_b = A\langle P_b L, P_b K_g\rangle = A \sum_{j=1}^{(b+1)^2} l_j k_{g,j}. \tag{D.5}$$

Orthogonality removes the cross terms, and Cauchy–Schwarz bounds the remaining product:

$$|\mathcal{R} - \mathcal{R}_b| \leq A \, \|L - P_b L\|_2 \, \|K_g - P_b K_g\|_2 . \tag{D.6}$$

The norms here are L^2 norms. Rigid rotation preserves the kernel tail norm; shape changes need not. Numerical projection and evaluation add errors beyond this truncation statement. The analytic degree-two toy contracts to about 2.44×10^{-15} error because its illumination is exactly bandlimited. That check establishes a basis convention, not a coefficient count for finite cones.

E Near-equal input records and derivative bounds

For correct responses r_i and supplied records x_i, insert both predictions into the triangle inequality:

$$|r_2 - r_1| \leq |r_2 - F(x_2)| + |F(x_2) - F(x_1)| + |F(x_1) - r_1|.$$

If the middle term is at most $\Lambda\|x_2 - x_1\|_2$, then

$$\max_i |F(x_i) - r_i| \geq \tfrac{1}{2} \left[|r_2 - r_1| - \Lambda \, \|x_2 - x_1\|_2 \right]_+ . \tag{E.1}$$

Replacing the true separation by the covered lower bound S and the input distance by a valid upper bound ρ gives the lower bound $\frac{1}{2}[S - \Lambda\rho]_+$.

A derivative describes how sensitive a prediction is to a small change in input. For a differentiable scalar calculation, a bound

$$\sup_{0 \leq \lambda \leq 1} \|\nabla F(x_1 + \lambda(x_2 - x_1))\|_2 \leq \Lambda \tag{E.2}$$

over the whole connecting segment implies the required output bound by integrating that derivative along the segment. A derivative at one point, or the largest slope observed at a finite sample, does not establish the same bound between samples. For uncertain input sets, the chosen Λ and ρ must cover all admissible records and relevant connecting segments, with scales and illumination differences treated consistently. This is a conditional robustness result, not an impossibility theorem for arbitrary functions of unequal manufactured inputs.

E.1 Scale-consistent inputs and an ordering alternative

Choose fixed physical reference scales $A_*, d_* > 0$ if area and depth enter an input metric separately, and use $A/A_*, d/d_*$, together with dimensionless pose/source coordinates. A change from millimetres to metres changes the coordinate units and the reference units together, not the resulting distance. More strongly, if a normalized predictor under the angular model obeys

$$F(Q, \lambda^2 A, \lambda d, L) = F(Q, A, d, L) \quad (\lambda > 0),$$

then at fixed Q, L it factors through $\xi = d/\sqrt{A}$. To see this without a dimensional-analysis shortcut, compare any two positive-area pairs with equal ξ: setting $\lambda = \sqrt{A_2/A_1}$ maps the first pair to the second, so their predictions agree. Thus a function $f(Q, \xi, L)$ is well defined. A finite source distance, fixed optical length or wavelength-dependent law supplies extra dimensionless ratios and is outside this reduction.

Scale invariance alone does not make f monotone. For a fixed centered, normally extruded template under the present angular law, increasing normalized depth reduces each clipped tent factor, so that template's T is nonincreasing in ξ. A predictor required to respect this property forms a restricted class. If the two supplied poses and illuminations agree and uncertain records satisfy $\inf \xi_{\mathrm{rect}} \geq \sup \xi_{\mathrm{sq}}$, every predictor in that class has $y_{\mathrm{rect}} \leq y_{\mathrm{sq}}$. For true transmissions with $r_{\mathrm{rect}} - r_{\mathrm{sq}} \geq S > 0$,

$$S \leq r_{\mathrm{rect}} - r_{\mathrm{sq}} \leq |r_{\mathrm{rect}} - y_{\mathrm{rect}}| + |y_{\mathrm{sq}} - r_{\mathrm{sq}}|, \qquad \max_i |y_i - r_i| \geq S/2. \tag{E.3}$$

This recovers the $S > 0.02$ exclusion without estimating a derivative, but only for that declared scale-invariant, monotone class and covered ordering. If the ordering, shared field or pose condition is unresolved, this shortcut is unavailable. Manufacturing a depth difference after observing a comparison would not validate it retrospectively; any modified geometry and prediction must be declared and qualified in advance.

F Encoding formats and geometry panel

Panel G uses seed 592601105, $w = 2^{a/2}$, $h = 2^{-a/2}$, $d = 2^u$, and $G = R_\theta[(w, 0)^\mathsf{T}, (kh, h)^\mathsf{T}]$. Roll spans a full turn and $t = -G\beta$. Certificates interpret each generated binary64 input as its exact rational value. The first three nonflat types use shear ranges $\pm 0.6, \pm 1, \pm 2$ respectively; near-parallel rows use signed shear of order 2^a. Fixed-point candidates use signed nearest rounding on $[-2, 2)$; floating candidates use IEEE binary16/32. Centered maps omit β. The fallback preserves five binary64 values for centered geometry and seven for offset geometry; flat records have a separate mode.

Streams contain a 9-byte header, a 4-bit mode per shape and whole-byte payloads. At tolerance 0.01 the 2,000 modes are 627 fixed-10, 256 fixed-12, 586 binary16, 426 binary32, 89 exact-input and 16 flat, totaling 28,951 bytes. Without fixed-point modes the total is 31,797; exact-input storage is 103,921. Pose, indexing and transfer costs are excluded.

There were 29,872 planned direction probes, including 12,462 at polar angles at least 85°. Probes witness errors; they do not certify continuous coverage. Invalid candidates are rejected first. Fixed-12 rejects 680 out-of-range and 217 rank-lost records. The post-hoc primitive control rounds

Table 13: Designed strata; a, u are aspect and depth exponents. The last column bounds each offset coordinate.

Stratum	Count	a	u	$\max \lvert\beta_i\rvert$
Moderate, centered	512	$[-2, 2]$	$[-2, 1]$	0
Moderate, offset	512	$[-2, 2]$	$[-2, 1]$	0.8
Shallow, grazing	384	$[-3, 3]$	$[-16, -5]$	2
Slender, offset	384	$\pm[4, 16]$	$[-7, 7]$	3
Nearly parallel inverse rows	192	$[6, 27]$	$[-5, 5]$	2
Flat controls	16	$[-2, 2]$	$d = 0$	0

G, d, β, rejects orientation reversal, rank loss, overflow or collapsed depth, and uses the same certificate. Its binary16/32 qualification counts at budgets $(0.001, 0.005, 0.01)$ are $(925, 1466, 1494)$ and $(1846, 1900, 1908)$. Neither chart dominates; adding primitive modes projects 28,547 bytes but no combined stream was emitted. The earlier twenty-shape, 107-byte pilot is a separate, less demanding panel.

G Computational protocols and independent checks

The deposited protocols and machine-readable configurations are the reproducibility authority; Appendix K locates them. This summary retains the distinctions needed to interpret the comparisons.

G.1 Geometry and ray checks

The domain requires finite rank-two edges with tangent singular-value ratio above 10^{-12}, finite nonzero directions, and positive normal separation unless $v = 0$. A nonzero tangential offset at zero depth is excluded. Overlap and triangle paths share input checks and the front cutoff but use independent visibility formulas.

The 93-case gate covers shear, roll, out-of-plane deformation, both extrusion conventions, scale, rigid rotation, flat, grazing and rear cases; eight further fixtures test the domain. Two seeds and three prefixes yield 558 estimates. At 65,536 origins the largest discrepancy is 0.00563936, inside the predeclared 99% simultaneous Hoeffding radius 0.00941669 [21]. Common-origin comparisons retain pairing. For the finite shear cone, twelve marginal estimates across seeds 58201 and 94733 have simultaneous response radius $\sqrt{\log(24/0.01)/(2N)}$. At the largest prefix $N = 65,536$, its normalized radius is 0.0103417: this check alone does not establish one-point accuracy. The finite-cone proof does.

The integrated-ray study fixes nine geometry indices and 32 source–geometry cases before errors are inspected. It samples directions proportional to $L\mu$ and mouth origins uniformly. All 224 prefix/pooled comparisons meet simultaneous 99% bounds over 14,680,064 positive-depth rays. A repaired open-segment endpoint convention leaves these outcomes unchanged. A separate 192-integral isotropic sweep records small independent-axis refinement changes; those are diagnostics, not enclosures.

G.2 Original benchmark and later source panels

The 352-case matrix uses eight calibration geometries and twelve held-out geometries, giving 112/240 cases. It varies shear, width, depth and roll, with normal extrusion. Broader affine cases belong to the geometry gate. Sources are isotropic light, the nonnegative degree-two field $1 + 0.6\widehat{a} \cdot s + 0.25(s_x^2 - s_y^2)$ with $\widehat{a} = (0.4, -0.2, 1)/\left\|(0.4, -0.2, 1)\right\|_2$, and uniform $1^\circ, 2^\circ, 5^\circ$ cones

on archived axes, including two held-out axes. SH degrees are $2, 4, 8, 12, 16, 24$; cap coefficients use exact Legendre moments, avoiding global-grid aliasing of the source while retaining kernel truncation.

Calibration selects 8×32 cone and 8×8 per-quadrant direct rules, a 33×128 visibility table with 0.0025 allocation, and 96×384 SH projection. Reference refinements vary coordinates independently and jointly. The maximum empirical allowance is 2.0871×10^{-5}. A later five-case audit finds shifts exceeding some allowances and changes the worst SH error from 0.0449837 to 0.0449028. It neither replaces the original results nor certifies the matrix. On positive-depth held-out cases, updated zonal/full SH pass 49/181 of 200, so the findings do not depend on flat controls. Rear-only sources have $P_0 = 0$ and are raw-power diagnostics.

The two 72-prism panels use seeds 127194 and 67319052, each with 24 ordinary centered, 24 ordinary offset and 24 stress cases. Normals incline by 0.03–0.44 radians. Ordinary aspect ratios are $\exp(z)$, $z \in [-1.2, 1.2]$; stress ratios are 7–24 or reciprocals. The initial protocol's near-square stress description was inaccurate; replication retains the executed distribution. Each prism receives two broad fields, thirteen cones and two mixtures, including grazing directions and shape-specific support features. Caps remain front-facing.

Reference refinement starts at 32×128 versus 64×256 for caps and 32^2 versus 64^2 per quadrant for broad fields, and continues when twice the largest specified change exceeds 0.0003. Two elementary replication queries reach the ceiling above that trigger; the maximum allowance is 0.000779. Polygon clipping agrees at common nodes on 64 selected cones to 8.89×10^{-16}; alternate angular integration checks the first panel's 144 broad responses. These comparisons are not universal error guarantees.

For source routing, let $r = L - L_q$, $D = \int \mu_+ L \, \mathrm{d}\Omega > 0$ and $N_q = \int \mu_+ V L_q \, \mathrm{d}\Omega$. Then

$$\left| T - \frac{N_q + \frac{1}{2} \int \mu_+ r \, \mathrm{d}\Omega}{D} \right| \leq \frac{\int \mu_+ |r| \, \mathrm{d}\Omega}{2D},$$

since $|V - 1/2| \leq 1/2$. Error in evaluating N_q is additional. The revised broad rule doubles to order 128; g66 stops incorrectly at 32. Range alarms cannot replace accuracy checks: on the original updated-full evaluation they miss all nineteen accuracy failures and raise fifteen false alarms.

G.3 Workload and memory accounting

Original batches use $N = 1, 10, 100, 1000$ and three repeats; component timings use seven. Setup recurs with shape changes. The efficiency follow-up chooses rules using calibration only and verifies 48 fresh cases on eight shapes; direct and lookup maximum discrepancies are 9.0674×10^{-6} and 7.9518×10^{-4}. The cache amendment preserves predictions to 1.11×10^{-16}. For broad fields alone at 1,000 fixed-shape queries, degree-two SH takes 0.005036 seconds versus 0.009262 for cached direct; with changing shape the respective medians are 1.852 and 0.06167 seconds. These are measured batch results, not exact crossover locations.

Payloads are 72 bytes for B, v, 160 for prepared geometry, 200/5,000 for degree-24 zonal/full SH, and 33,792 for the lookup. Degree-two response and source each need 72 bytes. The direct broad cache adds 8,192 bytes. Shared rules, query arrays and Python overhead are separate; per-method transient peaks were not instrumented. Later warmed-panel timing omits moment construction, transfer, pose and several setup costs, and its ranges overlap. It establishes no end-to-end speed or energy advantage.

H Clipped-domain integration: formulas and limits

For an inner-mouth point x_m, let P_{x_m} be the spherical projection of the outer mouth. The adapter computes

$$\mathcal{R} = \int_{BU} \int_{P_{x_m}} L(s)(n \cdot s)\, \mathrm{d}\Omega\, \mathrm{d}A_{x_m}. \tag{H.1}$$

For broad fields, $L(s)(n \cdot s)$ has degree at most three. Sixteen SH coefficients, obtained by a full-sphere rule exact for that polynomial space in real arithmetic, are transformed into zonal factors. The implementation follows Belcour et al., Eqs. (6), (9)–(11) and Algorithm 1 [11, 10], with axial moments related to Arvo's treatment [23]. Seven fixed asymmetric lobe directions support the factorization through degree three; the largest measured basis condition number is 18.719. The local routine evaluates factors separately rather than reproducing the published preassembled implementation.

For a uniform circular cone, the adapter uses an inscribed great-circle polygon $\mathcal{C}_k$. Intersecting it with the projected mouth gives the integration region. Inside that region the integrand is $L_0(n \cdot s)$, a degree-one function, so the angular step has no harmonic truncation in exact arithmetic. The original circular-cone radiance and denominator are retained; the polygon is not renormalized. The omitted source region contributes at most

$$b_{\mathrm{cap}} = \frac{A(|\mathcal{C}| - |\mathcal{C}_k|)}{|\mathcal{C}| P_0}. \tag{H.2}$$

Indeed, the exact difference is $A/(|\mathcal{C}|P_0)$ times the integral of μV on the omitted region, and $0 \le \mu V \le 1$. For a fully front cap, multiplication by $\mu_{\max} = \cos(\max\{0, \theta_0 - \alpha\})$ tightens the bound. The recorded comparison uses the original bound. It controls cap omission only; spatial quadrature and arithmetic remain separate. A floating-point evaluation of this bound is not automatically an outward-certified value. The comparator does not call the overlap law.

H.1 Contour recurrence

For a lobe axis a, set $x_a = a \cdot s$. With inward unit edge normals h_j and $E_{r,j} = \int_j x_a^r \, \mathrm{d}t$, the spherical divergence theorem gives

$$M_0 = |P|, \qquad M_k = \frac{k-1}{k+1} M_{k-2} + \frac{1}{k+1} \sum_j (a \cdot h_j) E_{k-1,j}, \tag{H.3}$$

omitting the first term for $k = 1$. This follows from $\Delta_{S^2} x_a^k = k(k-1)x_a^{k-2} - k(k+1)x_a^k$. On an edge starting at p with tangent $h \times p$, write $x_a(t) = A\cos t + B \sin t$ for $0 \le t \le \theta$, where A, B here are scalar edge coefficients. Then

$$E_0 = \theta, \qquad E_1 = A \sin\theta + B(1 - \cos\theta), \tag{H.4}$$

$$E_r = \frac{(r-1)(A^2 + B^2)E_{r-2} - [x_a^{r-1} x_a']_0^\theta}{r}, \qquad r \ge 2. \tag{H.5}$$

The derivative x_a' describes change along the edge. These recurrences replace an angular area integration by edge integrals. In particular,

$$\int_P s \, \mathrm{d}\Omega = \tfrac{1}{2} \sum_j \theta_j h_j. \tag{H.6}$$

All uses have $k \le 3$. The routine supports consistently wound convex minor-arc polygons in an open hemisphere; the geometry adapter is restricted to the study's fixed-normal positive-depth and coincident-mouth cases.

Controls use analytic octant monomials, rotation and winding tests, reconstruction of all sixteen SH modes, and independent radial-triangle integration over arbitrary test polygons. They also test illumination reconstruction, linear moments, zero depth, empty domains, scale invariance, rejection and cache equivalence. They do not use overlap as the angular oracle. A singular symmetric lobe choice and a winding error were corrected before the saved comparison protocol. The earlier five-case prototype used angular triangle quadrature for smooth light; it remains a feasibility record, not evidence that the later published equations had already been run.

H.2 Settings, decisions and local costs

The reused integrated-ray subset provides twelve calibration and twenty evaluation cases. Candidate mouth orders are $m = 4, 8, 16, 32$ and cap-polygon sizes $k = 32, 64, 128, 256$. The first lexicographic choice meeting the 0.005 calibration indicator is frozen at $m = 32, k = 64$. The indicator is

$$I = e + 2|T_{2m,k} - T_{m,k}| + b_{\mathrm{cap}} + u_{\mathrm{ref}}. \tag{H.7}$$

The cap-size change is recorded separately, not double-counted. Reference coordinates are refined independently. Decisions use $u = I - e$ in the same pass/fail/unresolved rule. This remains an empirical criterion.

For direct integration, the maximum discrepancy on the twenty reused cases is 0.00001626; the maximum indicator is 0.00004875. The corresponding values are 0.0006448 and 0.001371 for lookup, and 0.005524 and 0.01809 for the adapter. Its two unresolved cases share the shallow 1.2×0.5, depth-0.35, roll-23° rectangle under two 2° cones. Mouth-order changes are 0.005203 and 0.004272; even dropping the cap allowance leaves indicators 0.01593 and 0.01120. Tightening the cap bound alone cannot resolve them.

All twelve calibration cases pass. All eighteen broad-field cases, including calibration and flat controls, differ by less than 4.0×10^{-15}. The extra stress case `v_rect80a__normal_01` has $e = 0.001947, I = 0.004331$, compared with the archived whole-kernel SH error near 0.0450. No selected adapter prediction leaves $[0, 1]$.

Fixed-shape broad queries cache integrated axial moments; cached and uncached outputs agree to 1.17×10^{-15} over all 33 cases. Cone queries still construct and clip their source polygon. At selected order, numerical payload is 205,104 bytes including 72 bytes of geometry input, excluding object overhead and temporaries. Cold quadrature and basis construction take about 0.00207 and 0.000421 seconds. Median totals for 100 local mixed queries are: The serial NumPy adapter uses

Table 14: Local clipped-domain workload, three repeats. The adapter has two unresolved evaluation cases, so these are not a matched-accuracy ranking across the full set or timings of the published optimized implementation.

Implementation	Fixed shape	Changing shape
Cached direct	0.007284	0.012022
Visibility lookup	0.007723	0.047376
Local clipped-SH adapter	1.40935	14.52798

Python loops and a uniform mouth rule. The run took 66.46 seconds within a 600-second ceiling and recorded 177 distinct comparison responses plus cost workloads. A later audit of the seven timing pairs finds all four fixed-shape pairs passing; the changing-shape cycle includes an already unresolved case. It does not enlarge the evaluation set.

I Enclosing the two replication failures

The g66 calculation was constructed after the numerical failures were identified. It encloses the ideal isotropic integral using the archived binary64 inputs as exact rationals, without changing the saved prediction or classification.

Define

$$K = \frac{|\det G|}{\pi d^2}, \quad x(q) = \frac{Gq + t}{d}, \quad f(x) = (1 + \|x\|_2^2)^{-2}, \quad w(q) = (1 - |q_1|)(1 - |q_2|).$$

Then $T = K \int_{[-1,1]^2} w(q) f(x(q))\, \mathrm{d}q$ and $\int w = 1$. Partition each quadrant into 128 by 4,096 cells. In every cell, evaluate the smooth factor at the exact tent-weighted centroid and multiply by the cell's exact weight.

The first-order Taylor term integrates to zero at that centroid. At radius r, the tangential and radial Hessian eigenvalues of f are $-4/(1 + r^2)^3$ and $(20r^2 - 4)/(1 + r^2)^4$, both bounded in magnitude by four. With half-widths h_1, h_2, weighted coordinate variances are at most $h_j^2/3$ and the product weight makes the coordinates independent. If G_j is column j, the total remainder is bounded by

$$2K \sum_{j=1}^{2} \frac{\|G_j/d\|_2^2 h_j^2}{3} < 0.001092389. \tag{I.1}$$

The verifier brackets π by rational alternating series, computes masses and centroids rationally, and rounds kernel evaluations and summations outward. It assumes ordinary correctly rounded IEEE-754 binary64 operations as checked in the verifier. Combining 2,097,152 weighted evaluations with the remainder gives the intervals in Table 4.

For the mixture, $q_1 > 1$ becomes a linear directional half-space condition after multiplication by $\mu > 0$. An exact rational separating-vector test places the entire recorded cone in this blocked region. Its transmitted contribution is zero, so the open-flux mixture inherits 65% of the isotropic response. The verifier covers both the exact decimal weight and its saved floating-point value. These bounds establish the two errors independently of agreement between high-order quadratures.

J Geometry-aware quadrature and sharper representation tests

This follow-up was made after the original results and the review questions were known. It explains a recorded failure and examines tighter certificates on specified examples; it does not replace the original calibration, evaluation sets or encoding decisions. The revision 10 supplement retains executable calculations, exact rational witnesses, complete numerical sequences and the unchanged error verifiers.

J.1 The resolution scale hidden in the displacement integral

For isotropic illumination, define $\mathsf{A} = G/d = W^{-1}$, $\beta = -G^{-1}t$ and $H = \mathsf{A}^{\mathsf{T}}\mathsf{A}$. Here H is a matrix, distinct from the directional kernel $H_g(s)$. The normalized integral is

$$T = \frac{|\det G|}{\pi d^2} \int_{[-1,1]^2} \frac{(1 - |q_1|)(1 - |q_2|)}{[1 + (q - \beta)^{\mathsf{T}} H (q - \beta)]^2}\, \mathrm{d}q_1\, \mathrm{d}q_2. \tag{J.1}$$

On each quadrant the numerator is polynomial. The analytic continuation of the smooth denominator uses $1 + z^{\mathsf{T}} H z$, $z = q - \beta \in \mathbb{C}^2$, not a Hermitian complex norm. Write $z = x + iy$. A singularity satisfies

$$x^{\mathsf{T}} H y = 0, \qquad y^{\mathsf{T}} H y = 1 + x^{\mathsf{T}} H x. \tag{J.2}$$

Consequently $\|y\|_2 \geq 1/\sigma_{\max}(\mathsf{A})$. This lower bound on the distance from the whole real plane $\mathbb{R}^2$ is attained at $q = \beta \pm i v_{\max}/\sigma_{\max}$, where $v_{\max}$ is a unit right singular vector. For a simple largest singular value these are the nearest points of the singular set to β; a repeated value permits every unit vector in its largest right singular subspace. The complete singular set is a complex hypersurface, not just those points. Attainment of this distance relative to the integration square also requires $\beta \in [-1, 1]^2$, as holds for g66; an offset centre outside the square has a larger distance to that square.

Along a real line $q = q_0 + uv$, put $r = q_0 - \beta$, $a = v^{\mathsf{T}} H v$, $b = v^{\mathsf{T}} H r$ and $c = 1 + r^{\mathsf{T}} H r$. The slice poles are

$$u_{\pm} = -\frac{b}{a} \pm i \frac{\sqrt{ac - b^2}}{a}. \tag{J.3}$$

For unit v through β, their imaginary distance is $1/\|\mathsf{A}v\|_2$. These are double poles of the rational factor unless a numerator cancels part of the singularity. The clipped tent as a whole is not analytic across its creases or support boundary.

On the real domain, the same scale describes a narrow concentration of integration weight. Through the centre, $1/\sigma_{\max}$ is the *quarter-height halfwidth* of the smooth factor; its half-height halfwidth is $\sqrt{\sqrt{2} - 1}/\sigma_{\max}$. A condition number alone does not specify the needed resolution: $\mathsf{A} = cI$ has condition number one even for arbitrarily narrow peaks as c grows. Nor does the singular-value scale alone determine the error on a bounded domain: offset, source boundary and node positions matter.

For the original g66 inputs, the executed values are

$$\beta \simeq (0.77342478, -0.41712898)^{\mathsf{T}}, \quad (\sigma_{\max}, \sigma_{\min}) \simeq (54.70212987, 0.91420723).$$

A narrow normal is $(-0.03162019, 0.99949996)^{\mathsf{T}}$. The long axis is inclined by about 1.81201° to q_1. The narrow quarter-height halfwidth is 0.01828082 and the long one 1.09384390. Thus splitting at zero aligns the tent creases but not the off-centre ridge. The two column norms of A are about 1.95621410 and 54.67478408: the second coordinate is the difficult direction in this example.

Table 15: Post-hoc replay of the original quadrant tensor rule on g66. Order n means n nodes in each coordinate of each of four quadrants, not a global node count. All entries are numerical estimates.

Order n	Original-input $Q_{n,n}$	Decoded-record $Q_{n,n}$
8	0.188612469695	0.186826947542
16	0.167780439485	0.168638319614
24	0.090094587587	0.089943433828
32	0.169004239382	0.168817649590
64	0.146861783620	0.146813390022
128	0.147003143831	0.146966694119
256	0.147004102181	0.146967585768

The archived stopping rule acts on the *decoded* record. Its four relevant values are

$$Q_{16,16} = 0.16863831961440257, \qquad Q_{16,32} = 0.16881764958981135,$$
$$Q_{32,16} = 0.16863831961440254, \qquad Q_{32,32} = 0.16881764959045284.$$

Twice the largest specified change is 0.000358659952100604, agreeing with the saved indicator to rounding accuracy. Refining q_1 at fixed second-coordinate order 32 changes only 6.42×10^{-13}. The next joint refinement instead changes the result by 0.02200426. For the original-input record, the 16-to-32 indicator is 0.00244760 and the same rule continues to order 128. A small coefficient change altered the coarse-grid integration error much more than the underlying isotropic response.

The ridge was not wholly unsampled. For decoded orders 16 and 32, the second-coordinate gaps bracketing its centre are about 0.09330 and 0.04741, but respectively ten and 29 tensor nodes lie in $\|\mathsf{A}(q-\beta)\|_2 \le 1$. Different grids assign incorrect aggregate weight to different parts of the strip. Two overestimates nearly agree; the order-24 dip makes the nonmonotone behaviour particularly clear. Nearby poles explain the resolution hazard. They do not by themselves establish an optimal coordinate system or a certified stopping criterion.

J.2 Analytically removing the difficult coordinate

For fixed $x = q_1$, write $\mathsf{A}(q-\beta) = ay + b(x)$ with $y = q_2$. Define

$$\alpha = a^{\mathsf{T}}a, \qquad m = \frac{a^{\mathsf{T}}b}{\alpha}, \qquad \delta = 1 + \frac{\det(a,b)^2}{\alpha} > 0. \tag{J.4}$$

The determinant expression is the two-dimensional identity for $1 + \|b\|_2^2 - (a^{\mathsf{T}}b)^2/\alpha$ and avoids that subtraction. With $u = y + m$ and $D_u = \alpha u^2 + \delta$, primitives for D_u^{-2} and uD_u^{-2} are

$$J_0(u) = \frac{u}{2\delta D_u} + \frac{\arctan(u\sqrt{\alpha/\delta})}{2\delta\sqrt{\alpha\delta}}, \tag{J.5}$$

$$J_1(u) = -\frac{1}{2\alpha D_u}. \tag{J.6}$$

On a tent half-interval, with factor $1 - sy$ and $s \in \{-1, 1\}$, the primitive is $(1 + sm)J_0 - sJ_1$. Evaluate it on $[-1, 0]$ and $[0, 1]$, then integrate the remaining factor $1 - |x|$ in one dimension, again split at zero. The implementation evaluates arctangent differences with `atan2` to retain the branch across the peak.

At outer orders eight and sixteen per half-interval, the original-input estimates are respectively 0.147004102202962 and 0.147004102182854; the decoded estimates are 0.146967585794443 and 0.146967585774395. Higher-order tensor and opposite-direction analytic checks agree within 6.1×10^{-12} on these records. Integrating the broad first coordinate analytically, but leaving the narrow second coordinate at order 32, still gives the bad decoded estimate. Removing one dimension is useful only if the unresolved structure is actually addressed.

A second post-hoc control adds splits at β while retaining the zero splits. At order 32 per subinterval it gives 0.146967671854 for the decoded record, using 9,216 tensor evaluations rather than the original 4,096. This is not an equal-work timing comparison. Neither control is a general accuracy guarantee: other sources and geometries need their own analysis, and floating-point primitive evaluation can suffer cancellation in other regimes.

Both unchanged outward verifiers reproduce their saved endpoints. The original enclosure in Table 4 therefore remains the rigorous evidence, not the agreement between the new estimates. Comparing the four displayed analytic-second-coordinate estimates at outer orders eight and sixteen with that enclosure gives an a posteriori absolute-error bound below 0.00113 against the original ideal integral. It does not certify all displayed digits, a 0.001 allocation, or the one-dimensional algorithm on other cases. The two saved failures remain failures.

J.3 Exact maxima need more than partition vertices

The support-corner test maximizes a convex *affine-error envelope*. Keeping the cofactor weights instead leads to a different function. In common slope coordinates, each product of two affine tent factors is quadratic, and the difference of two such products is generally quadratic too. A polygon's vertices need not maximize its absolute value.

An exact example uses $f(u) = [1 - |u|]_+$ and

$$W = I, \qquad \widetilde{W} = \begin{pmatrix} 1 & 1/2 \\ 0 & 1 \end{pmatrix}, \qquad \beta = \widetilde{\beta} = 0. \tag{J.7}$$

Then $V = f(x)f(y)$, $\widetilde{V} = f(x + y/2)f(y)$, and

$$|V - \widetilde{V}| \le \tfrac{1}{2}|y|(1 - |y|) \le \tfrac{1}{8} \quad (|y| \le 1), \tag{J.8}$$

with equality at $(0, 1/2)$ or $(1, -1/2)$; outside that strip both vanish. Partition by all lines $x = -1, 0, 1$, $x + y/2 = -1, 0, 1$ and $y = -1, 0, 1$. Every partition vertex in the support union has $y = 0$ or $y = \pm 1$, where the difference is zero. The vertex-only result is therefore zero, the true supremum $1/8$, and the support-corner bound $1/2$. Tightness cannot be inferred by looking only at support corners, nor summarized by a universal factor.

The exact-rational algorithm splits a bounding rectangle of the two supports by all their support edges and zero creases, discards zero-response cells and forms the signed quadratic Q in each remaining polygon. It tests vertices, stationary parameters on each edge, and interior solutions of $\nabla Q = 0$. For a singular Hessian, any interior stationary affine set has constant value and meets the boundary of the bounded polygon; edge and vertex checks cover it. Maximizing $|Q|$ over these candidates is exact for rational records. No subdivision of the conic $Q = 0$ is needed: its zero values cannot create a positive maximum.

Table 16: Four post-hoc binary16 record checks. Displayed decimals approximate exact rational values retained with their witnesses in the supplement; they are not outward decimal enclosures.

Fixture	Vertex-only maximum	Exact supremum	Corner bound	Ratio
g00	0.0000376476	0.0000403878	0.0001783231	4.4153
g24	0.0002744115	0.0002744115	0.0003376658	1.2305
g48	0.0004641438	0.0005063614	0.0013257842	2.6183
g66	0.0038894204	0.0043346797	0.0114918467	2.6511

The g00, g48 and g66 maxima occur inside an edge; g24 attains its maximum at a vertex. For g66 the exact result is below 0.004335, hence below the 0.01 representation target. This tighter record guarantee does not cover the solver's failed integration. No revised coverage or storage rate is claimed for the earlier 2,000-geometry panel.

J.4 A flux-weighted certificate for a restricted source family

Let $D_L = \int_{\mu>0} L\mu \, \mathrm{d}\Omega > 0$ and compare visibilities at the same exact pose, with exact cosine weighting and denominator. If the source family satisfies $0 \le L/D_L \le \kappa$ almost everywhere, then

$$|T - \widetilde{T}| \le \kappa I, \qquad I = \int_{\mu>0} |V - \widetilde{V}|\mu \, \mathrm{d}\Omega = \int_{\mathbb{R}^2} \frac{|V(p) - \widetilde{V}(p)|}{(1 + \|p\|_2^2)^2} \, \mathrm{d}p. \tag{J.9}$$

The second equality follows from $\mu = (1 + \|p\|_2^2)^{-1/2}$ and $\mathrm{d}\Omega = (1 + \|p\|_2^2)^{-3/2} \, \mathrm{d}p$. Necessarily $\kappa \ge 1/\pi$ for a nonempty normalized family. The minimum of this bound and any all-source bound is also valid. A common radiance upper bound and a strictly positive common lower bound on D_L are one way to justify κ.

For the half-shear above, the full support is contained in $[-3/2, 3/2] \times [-1, 1]$. Interval bounds on each tent factor and on the slope weight, accumulated as outward Darboux sums over 1024×1024 cells, give the outward decimal enclosure

$$I \in [0.1277706842949135,\ 0.1390760178538125]. \tag{J.10}$$

No tail is omitted. The resulting representation-error upper bounds, rounded upward, are 0.044269272 for $L/D_L \le 1/\pi$ and 0.088538543 for $L/D_L \le 2/\pi$, compared with the all-source supremum 0.125. These bounds assume correctly rounded IEEE-754 binary64 basic operations expanded outward, exact rational input conversion and exact rational block aggregation. The smaller numerical isotropic difference, about 0.00480132, uses signed cancellation and is not the basis of this certificate.

If a different approximation has noncompact support, a bounded-region calculation must include the tail. A valid bound $|V - \widetilde{V}| \le M_R$ outside a radius R gives

$$\int_{\|p\|_2 > R} \frac{|V - \widetilde{V}|}{(1 + \|p\|_2^2)^2} \, \mathrm{d}p \le \frac{M_R \pi}{1 + R^2}. \tag{J.11}$$

Two visibilities in $[0, 1]$ permit $M_R = 1$; an unclipped approximation outside that range needs a separate bound.

Merely requiring bounded radiance does not remove worst-direction sensitivity. For example, take $L = 1$ on a slope square of halfwidth $\epsilon \le 1/16$ about $(1/4, 1/2)$ in the half-shear example, and zero elsewhere. Its signed error is $y(1 - y)/2$, so

$$\tfrac{1}{8} - \tfrac{1}{2}\epsilon^2 \le |T - \widetilde{T}| \le \tfrac{1}{8}, \qquad 0 < D_L \le 4\epsilon^2 \longrightarrow 0. \tag{J.12}$$

The radiance cap stays fixed while the normalized contrast grows. More generally, bounded indicator illuminations on sets approaching the essential supremum establish Eq. (D.3). The broad-family certificate is therefore a different guarantee, not a correction to that equality. Neither guarantee includes arithmetic, a particular source integral, denominator uncertainty or physical-input error.

K Reproducibility and study chronology

The earlier numerical studies ran with Python 3.12.14, NumPy 2.3.5 and an Apple M3 on macOS arm64, requesting one numerical-library thread, including the Accelerate setting. Archived geometry and primitive-control generator replays used Python 3.9.6. Source snapshots, configurations, outputs and hashes identify the code and inputs behind each result; host timings can vary.

The original focused suite records 33 passing checks: ten helper, seven workspace, nine domain and seven representation checks. Later diagnostics and the clipped comparison expanded the earlier suite to fifty passing checks. These are historical test results associated with the retained runs, not independent human peer review or physical validation. In those earlier studies, the optional Radiance `rtrace` comparison [22] was not executed because no usable executable was found. The revision 11 Radiance check covers the six finite-cap cases in Section 6.

K.1 What was fixed before execution?

"Pre-declared" means saved in a local protocol before the specified execution or selection stage. It does not mean public preregistration or independent timestamping. A later study may have its own frozen protocol while remaining post hoc relative to earlier findings.

K.2 Review-response analysis

The revision 10 analysis in Appendix J was executed with Python 3.9.6 and NumPy 2.0.2. Its 33 focused tests cover the analytic primitive, quadrature replay, exact rational extrema, degeneracies and interval bounds. The two original g66 verification programs were rerun unchanged and reproduce their saved endpoints exactly. These tests are distinct from the historical 33-check

Table 17: Evidence chronology retained in the rewrite. Later studies do not retroactively enlarge the original held-out set.

Study or check	Status and decision boundary
Main 352-case study	Case matrix, calibration/evaluation split, candidates and selection rules fixed before the main run. The 240 evaluation cases were held out from parameter selection; degree 24 was the specified diagnostic if no SH candidate passed.
Integrated wall rays	Geometry indices, sources, seeds and comparison rule fixed before model errors were inspected.
Lookup efficiency	Motivated by the original timing result; calibration-only rule selection followed by 48 fresh confirmation cases. The direct-cache amendment changes cost accounting, not accuracy settings.
Five-case audit	Cases selected after errors were known. Alternate rules check sensitivity; they neither supply fresh validation nor turn allowances into enclosures.
Aspect-ratio companion	Later explanatory choice for simple algebra, not held-out validation.
Clipped-SH comparison	Post-hoc adapter study with settings frozen before evaluating twenty reused cases. The known worst SH case is an additional stress test; the later seven-pair workload audit is not a larger evaluation set.
Geometry panel G	Strata and support certificate fixed before the 2,000-case study. Primitive-coordinate candidates added after inspecting that study.
Finite-source panels	Initial routing rule followed by a revision on panel 1, then a separately seeded replication. The g66 enclosures were constructed after the failures were identified.
F1 planning and intake	Simulated sensitivities, a coarse digital control and prospective uncertainty rules; no comparative physical readings.

suite above. The four certificate fixtures and quadrature variants were selected after the original results were known; they do not constitute a new held-out study or revise the earlier stream sizes.

The revision 10 supplement preserves a small workspace-shaped tree so the follow-up can run without changing its paths. Under `research/path-a-review-20261108/numerics/`, `reproduce.py` writes to a new output directory, `results/analysis.json` holds the reported calculations and `results/manifest.json` records commands, software versions and matching before/after hashes for 24 preserved inputs. The previous revision 9 supplement and the initial revision 10 sources are retained as inputs under `Aperture-R10-author-review/`. The revised manuscript is a separate source package. The bench notes distinguish nominal design inspection from physical characterization; no new as-built observations are supplied.

K.3 Locating the evidence

The revision 8 deposit cited in Data availability contains `aperture_v6_licensed_archive.zip` for the earlier benchmark and solver archive, and `aperture_revision_8_diagnostics.zip` for the later numerical audit. Within those retained materials, the principal run directories are:

- Geometry gate: `results/geometry_gate_20261004T062554625239Z/`.
- Main matrix: `results/bounded_study_20261004T063142675548Z/`.
- Integrated wall rays: `results/integrated_mesh_20261004T062801246681Z/`.
- Lookup efficiency and cached-direct costs: `results/lookup_efficiency_audit_20261004T063638749320Z/`.
- Aspect-ratio companion: `results/aspect_ratio_companion_20261010T163751447400Z/`.

- Clipped comparison: `results/clipped_sh_comparison_20261010T172049740951Z/`.

The revision 9 scientific supplement supplies its own checksums and replay instructions. Its smoke check decodes the three geometry streams and recomputes their exact support bounds, reproduces both g66 enclosures, checks the 66 simulated bench conditions and the coarse-control calculation, and exercises the intake with synthetic and unfilled records. The pilot has a separate replay. Complete studies write to new output directories so the archived outputs remain unchanged.

K.4 Revision 11 computational evidence

The polygon, interior-stop and Klems witnesses are designed examples. Boundary depths 4, 6 and 8 are primary; depths 5 and 7 and 51 finer-offset center-ray checks are later examinations of the same cases. The restricted-view protocol fixed its 1,890 configurations, normalization and initial policy before its main run; two screens were then refined, followed by a centered-radius search. Deeper trees reuse the same 24 matched configurations. Fixed-marginal redistribution follows XML generation; tighter merging and unmerged policies follow that diagnostic. None enlarges the original held-out benchmark. The isotropic audit is retrospective; the shallow-square control deliberately lies outside its panels. The inverse study is synthetic with a fixed field basis, noise model and estimator.

The R11 supplement contains `daylight/`, `boundary/`, `viewcone/`, `deepview/`, `mergeview/`, `geometry/`, `audit/`, `inverse/` and `gx01-replay/`. These directories supply protocols, code, inputs or seeds, tables, checksums and replay instructions; XML archives preserve actual tree bytes. Large sweeps, enclosures, inversions, tree generation and native readback ran on Modal; direct-ray and small analytical checks and the tagged building replay ran locally. Each result retains its original denominator, chronology and numerical assumptions. No physical readings are included.

K.5 Revision 13 source and verification boundary

R13 reconstructed the manuscript from the issued R11 source and the published corrected R12 PDF. It preserved the forward-hemisphere condition on the outgoing bound, the shallow-square reduction and asymptotics, and the finite-offset interpretation of the inverse results. Its archive records the editorial comparison and selected analytical checks. The earlier computational studies retain their original datasets and validation status.

K.6 Revision 14 profile and measurement checks

R14 uses the published R13 source. The new examples in Sections 7.4–7.6 followed a targeted literature comparison. All four profiles, station errors and illumination ranges are designed fixtures for this study.

In the package, `checks/check_taper_geometry.py` checks the obstructions, volume and lower hull. The routine in `checks/taper_bounds.py` computes the enclosures from the reported station data, exact end disks, curvature assumption and source. Separate routines use the generating walls for verification. A second implementation in `checks/reference_crosscheck.py` compares nominal power using circle arcs and Gaussian angular quadrature. Machine-readable results and replay instructions accompany the source.

References

[1] EnergyPlus. *Engineering Reference*, version 26.2.0, “Window Heat Balance Calculation,” subsection “Beam Solar Reflection from Window Reveal Surfaces.” Official documentation. Accessed 10 October 2026.

[2] A. McNeil. *The Five-Phase Method for Simulating Complex Fenestration with Radiance*, version 2.2, 11 September 2013, section 6.1, p. 12. Official tutorial. Accessed 10 October 2026.

[3] Y. Li, J. Oldendick, C. E. Ordonez and W. Chang. The geometric response function for convergent slit-slat collimators. *Physics in Medicine and Biology*, 54(6), pp. 1469–1482, 2009. doi:10.1088/0031-9155/54/6/006.

[4] P. Schröder and P. Hanrahan. A closed form expression for the form factor between two polygons. Princeton University, Research Report CS-TR-404-93, January 1993. Institutional report record.

[5] G. J. Ward, R. G. Mistrick, E. S. Lee, A. McNeil and J. C. Jonsson. Simulating the daylight performance of complex fenestration systems using bidirectional scattering distribution functions within Radiance. *LEUKOS*, 7(4), pp. 241–261, 2011. doi:10.1080/15502724.2011.10732150.

[6] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. *Proceedings of SIGGRAPH*, pp. 497–500, 2001. doi:10.1145/383259.383317.

[7] P.-P. Sloan, J. Kautz and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics*, 21(3), pp. 527–536, 2002. doi:10.1145/566654.566612.

[8] P.-P. Sloan, B. Luna and J. Snyder. Local, deformable precomputed radiance transfer. *ACM Transactions on Graphics*, 24(3), pp. 1216–1224, 2005. doi:10.1145/1073204.1073335.

[9] R. Ng, R. Ramamoorthi and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics*, 22(3), pp. 376–381, 2003. doi:10.1145/882262.882280.

[10] L. Belcour. IntegralSH: Analytical Integration of Spherical Harmonics. Author-maintained implementation documentation. `https://github.com/belcour/IntegralSH`. Accessed 10 October 2026. Used to check method identity; this software was not executed.

[11] L. Belcour, G. Xie, C. Hery, M. Meyer, W. Jarosz and D. Nowrouzezahrai. Integrating clipped spherical harmonics expansions. *ACM Transactions on Graphics*, 37(2), Article 19, pp. 1–12, 2018. doi:10.1145/3015459.

[12] J. H. Klems. A new method for predicting the solar heat gain of complex fenestration systems: I. Overview and derivation of the matrix layer calculation. *ASHRAE Transactions*, 100(1), pp. 1065–1072, 1994. Author manuscript: LBL-34715, October 1993.

[13] G. Averkov and G. Bianchi. Confirmation of Matheron’s conjecture on the covariogram of a planar convex body. *Journal of the European Mathematical Society*, 11(6), pp. 1187–1202, 2009. doi:10.4171/JEMS/179.

[14] D. Geisler-Moroder, G. J. Ward, T. Wang and E. S. Lee. Peak extraction in daylight simulations using BSDF data. *Proceedings of Building Simulation 2021: 17th Conference of IBPSA*, pp. 1506–1512, 2021. doi:10.26868/25222708.2021.30374.

[15] A. McNeil, J. C. Jonsson, D. Appelfeld, G. Ward and E. S. Lee. A validation of a ray-tracing tool used to generate bi-directional scattering distribution functions for complex fenestration systems. *Solar Energy*, 98, pp. 404–414, 2013. doi:10.1016/j.solener.2013.09.032.

[16] A. McNeil and E. S. Lee. A validation of the Radiance three-phase simulation method for modelling annual daylight performance of optically complex fenestration systems. *Journal of Building Performance Simulation*, 6(1), pp. 24–37, 2013. doi:10.1080/19401493.2012.671852.

[17] P. Apian-Bennewitz et al.; D. Geisler-Moroder (ed.). Analysis and evaluation of BSDF characterization of daylighting systems. IEA SHC Task 61 / EBC Annex 77, technical report, 2021. doi:10.18777/ieashc-task61-2021-0012.

[18] D. Geisler-Moroder, E. S. Lee, P. Apian-Bennewitz et al. BSDF data generation for daylight applications: A call for international standardization. *Lighting Research & Technology*, 57(6–7), pp. 579–603, 2025. doi:10.1177/14771535241291838.

[19] W. L. Oberkampf and T. G. Trucano. Verification and validation in computational fluid dynamics. *Progress in Aerospace Sciences*, 38(3), pp. 209–272, 2002. doi:10.1016/S0376-0421(02)00005-2.

[20] T. Möller and B. Trumbore. Fast, minimum storage ray-triangle intersection. *Journal of Graphics Tools*, 2(1), pp. 21–28, 1997. doi:10.1080/10867651.1997.10487468.

[21] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301), pp. 13–30, 1963. doi:10.1080/01621459.1963.10500830.

[22] Radiance. *rtrace manual*. `https://radsite.lbl.gov/radiance/man_html/rtrace.1.html`. Accessed 4 October 2026.

[23] J. Arvo. Applications of irradiance tensors to the simulation of non-Lambertian phenomena. *Proceedings of SIGGRAPH*, pp. 335–342, 1995. doi:10.1145/218380.218467.

[24] J. D. Sullivan. Geometrical factor and directional response of single and multi-element particle telescopes. Nucl. Instrum. Methods 95 (1971), 5–11. DOI.

[25] B. Walter et al. Lightcuts: a scalable approach to illumination. SIGGRAPH (2005). Author page.

[26] P. Jiao et al. Toward quantity-of-interest preserving lossy compression for scientific data. Proc. VLDB Endow. 16 (2022), 697–710. Publisher PDF.

[27] D. Langharst, M. Roysdon and A. Zvavitch. General measure extensions of projection bodies. *Proceedings of the London Mathematical Society*, 125 (2022), 1083–1129. doi:10.1112/plms.12477.

[28] A. Wang, P. Gill and A. Molnar. Light field image sensors based on the Talbot effect. *Applied Optics* 48 (2009), 5897–5905. doi:10.1364/AO.48.005897.

[29] J. Mardaljevic, Aperture-Based Daylight Modelling: Introducing the 'View Lumen', Building Simulation 2019, 1129–1136. doi:10.26868/25222708.2019.210810.

[30] G. Ward, `klems_ang.cal`, Radiance source, revision 1.2 (2016), full Klems ring and sector definitions.

[31] J. R. Howell, *A Catalog of Radiation Heat Transfer Configuration Factors*, configuration C-14, perpendicular rectangles sharing an edge.

[32] Lawrence Berkeley National Laboratory, Radiance 6.0.2, release `rad6R0P2`, binary distribution and source revision.

[33] M. Nasr. Community Codec 0.2.0 preview. Software release, 2026. Versioned release and replication kit.

[34] K. Krisciunas and B. E. Schaefer. A model of the brightness of moonlight. *PASP* 103 (1991), 1033–1039. doi:10.1086/132921.

[35] M. W. Coughlin, C. W. Stubbs and C. Claver. A daytime measurement of the lunar contribution to the night sky brightness in LSST's ugrizy bands—initial results. Author preprint, 2015.

[36] F. Falchi et al. The new world atlas of artificial night sky brightness. *Science Advances* 2 (2016), e1600377. doi:10.1126/sciadv.1600377.

[37] S. T. Fabian et al. Why flying insects gather at artificial light. *Nature Communications* 15 (2024), 689. doi:10.1038/s41467-024-44785-3.

[38] P. Shirley and K. Chiu. A low distortion map between disk and square. *Journal of Graphics Tools* 2(3), 45–52 (1997). doi:10.1080/10867651.1997.10487479.

[39] A. McNeil and E. S. Lee. On the benefits of a variable-resolution bidirectional scattering distribution data format. Lawrence Berkeley National Laboratory, 2012. LBNL report.

[40] M. Nasr. Community Codec 0.3.0-preview, 2026. Versioned release, GX01 format and provenance, and building geometry and weather replay.

[41] Corridor Numerics. Research programme and open optical bench protocol. Accessed 10 November 2026.

[42] A.-T. Chai. Some basic considerations of measurements involving collimated direct sunlight. NASA CP-2010, NASA Lewis Research Center, 1977. NASA report, p. 3 and Fig. 1.

[43] A. V. Gitin. Phase space description of vignetting in optical systems. *Optics Communications*, 285(10–11), pp. 2485–2491, 2012. doi:10.1016/j.optcom.2012.01.048.

[44] S. E. Albo, L. J. Broadbelt and R. Q. Snurr. Transmission probabilities and particle–wall contact for Knudsen diffusion in pores of variable diameter. *Chemical Engineering Science*, 62(23), pp. 6843–6850, 2007. doi:10.1016/j.ces.2007.08.041.

[45] P. Blanc, B. Espinar, N. Geuder, C. Gueymard, R. Meyer, R. Pitz-Paal, B. Reinhardt, D. Renné, M. Sengupta, L. Wald and S. Wilbert. Direct normal irradiance related definitions and applications: The circumsolar issue. *Solar Energy*, 110, pp. 561–577, 2014. doi:10.1016/j.solener.2014.10.001.

[46] S. O. Rohrbach. Corrected recursive equations for vane aperture sizes and locations in one- and two-stage stray light baffle systems. *Earth Observing Systems XXVIII*, Proceedings of SPIE, 2023. doi:10.1117/12.2677700. NASA manuscript.

[47] H. Edelsbrunner, H. Maurer, F. Preparata, A. Rosenberg, E. Welzl and D. Wood. Stabbing line segments. *BIT Numerical Mathematics*, 22(3), pp. 274–281, 1982. doi:10.1007/BF01934440.

[48] V. Amato, A. L. Masiello, G. Paoli and R. Sannipoli. Sharp and quantitative estimates for the p-Torsion of convex sets. *Nonlinear Differential Equations and Applications NoDEA*, 30(1), Article 12, 2023. doi:10.1007/s00030-022-00822-y.