

Physicist verification handbook

Equations, reference planes, executable calculations, exact inputs, source implementations and numerical limits for the 337.1 / 777.4 nm CubeSat study.

Verification scope. Independent mathematical checks and native-array reconciliation are distinct from numerical convergence and experimental validation. GHM point-nadir rare-path convergence remains incomplete. Unbundled physical inputs are identified.

IMPLEMENTATION_CHECKS_PASS_WITH_EXPLICIT_PHYSICAL_LIMITS

The editable notebooks, source code and full numerical tables accompany this methods handbook. Open index.html in the portable package, or use the verification section of the scientific site. The traceability page maps each calculation to its equations, notebook, source functions and numerical checks. Requirements, exact input checksums and execution receipts are supplied.

After installing requirements.txt, run python verify.py for lightweight checks or python scripts/execute_notebooks.py for fresh notebook execution. See README.md for opt-in production jobs.

Worked-example parameter deck

Receiver choices below are illustrative assumptions. Received fluence is already at the aperture plane. Values are provided in SI; they do not constitute a target event population or a flight design.

Input	Value / convention
aperture_diameter_m	0.005
full_circular_fov_rad	1.0471975511965976
end_to_end_primary_efficiency	0.2
wavelength_m	3.371e-7
equivalent_rectangular_bandwidth_nm	4
on_window_s	0.003
continuum_W_m2_sr_nm	1.77e-9
admitted_auroral_brightness_R	0
projected_aurora_filling	1
dark_primary_rate_s_inverse	50
received_source_fluence_ph_m2	100000000
status	HYPOTHETICAL_WORKED_EXAMPLE_NOT_FLIGHT_REQUIREMENTS
background_scope	Ground continuum component only; not total polar cloud radiance
photon_plane	Fluence at aperture after propagation; apply efficiency once

Review contract and photon reference planes

What can be verified

The package exposes equations, parameter values, units, data origins, implementation locations, residuals and tolerances. Original r4 inputs are read-only references. Changed assumptions define a new conditional scenario.

The evidence levels are independent mathematical checks, reconciliation of the same simulation, numerical convergence, and experimental validation. The last two may remain incomplete while the first two pass. This is consequential for the tabulated-ice point-nadir estimator.

Photon reference planes

For an isotropic point source and a small planar nadir aperture:

$$K_{\text{vac}} = \frac{1}{4\pi R_0^2}, \quad f_{\text{ap}} = AK, \quad N_{\text{PE}} = \eta A \Phi_{\text{ap}}. \quad (\text{V1})$$

Here K is delivered photons per square metre per emitted photon. It is not an escape probability. The directional gain K/K_{vac} can exceed one. The transport chapter uses K_0 for its attenuated, unscattered direct term; K_{vac} always means the vacuum reference here. The aperture-interception fraction is a different quantity. When aperture-plane fluence has already been measured, applying atmospheric attenuation again counts that loss twice.

For a uniform extended radiance field and a circular cone in air:

$$G = A \int_{\Omega} \cos\theta \, d\Omega = A\pi \sin^2\theta_{\text{max}}, \quad \Omega = 2\pi(1 - \cos\theta_{\text{max}}). \quad (\text{V2})$$

Using A times geometrical solid angle for a wide filled cone omits the cosine projection. A nonuniform angular optical efficiency requires its own integral.

A band radiance in rayleigh and an energy spectral radiance are different inputs:

$$1 \text{ R} = \frac{10^{10}}{4\pi} \text{ photons m}^{-2} \text{ sr}^{-1} \text{ s}^{-1}, \quad E_{\gamma} = \frac{hc}{\lambda}. \quad (\text{V3})$$

Dividing a band-integrated auroral brightness by 4 nm gives an average spectral comparison only under that declared convention. Its photon-count contribution is integrated once; it must not also be multiplied by the filter bandwidth.

Numerical acceptance and uncertainty

Algebra checks state absolute and relative tolerances. Quadrature checks compare separately formulated integrals with closed forms. Monte Carlo checks disclose seeds, batch counts and standard errors. A native-array identity does not validate the distribution that generated those arrays.

Adjacent response bins share photon histories. Their error bars must not be added as independent variances. Finite-batch standard errors can miss rare large contributions. Small estimated errors do not remove the demonstrated GHM convergence gap.

Conditions outside the model

The source is assumed isotropic. The new cylinder is homogeneous and conservative with vacuum outside. The molecular column is a separate direct-beam calculation. Scalar transport does not model polarization, event-specific geometry or an observed ice-habit distribution.

Ideal Poisson waveforms omit analog correlation, afterpulsing, particle bursts and operational event-selection effects. There is no empirical Kp/lunar narrowband radiance distribution, universal minimum thermal 777/337 ratio, selected-device orbit-degradation model or target fluence/exposure population in this package. These are missing inputs, not numerical zeros.

Design algebra and analytic diagnostics: equations, assumptions and verification

This is a reproducibility companion, not a new atmospheric observation or a flight qualification. The independent implementation is `modules/design_review.py`. It imports no original calculation script. It reads the declared parameter grids and comparison targets from `data/design_results.json` and `data/analytic_results.json`. Fresh values and comparison receipts are written to `checks/design_*` and `checks/diagnostics_*`.

The stable equation identifiers D01-D14 and P01-P09 connect this account to the code, notebooks and machine-readable method maps. Exact Poisson detection thresholds and their coverage are verified elsewhere: in the FOV and auroral conversion checks here, the original critical count and required detected signal are explicitly inputs. No method below turns selected bright events into a complete source population.

1. Reference planes, units and constants

The reference planes are: emitted source photons; photons delivered through the atmosphere; photons crossing the specified entrance plane; and registered detections. A measured spacecraft irradiance or fluence already contains the atmospheric delivery. Applying cloud transmission again would attenuate it twice.

Geometry, time and photon wavelength arguments to the independent functions use metres, seconds and metres. The tabulated radiance convention per nanometre is retained only in the paired product `L_per_nm * band_nm`; equivalently multiply the former by 10^9 to obtain per metre and the latter by 10^{-9} to obtain metres. Those factors cancel. Rows labeled mm, nm, microseconds or degrees are converted explicitly at the function boundary. Steradians are dimensionless in SI but are retained in the bookkeeping to prevent confusing radiance with irradiance.

The SI defining values used are $h = 6.62607015 \times 10^{-34}$ J s, $c = 299792458$ m/s, $e = 1.602176634 \times 10^{-19}$ C, and $k_B = 1.380649 \times 10^{-23}$ J/K. They are not fitted to any event. The official primary values are in the [BIPM defining-constants table](#), rows c, h, e and k (inspected 22 October 2026). The historical US Standard Atmosphere constants below are retained separately rather than silently modernized.

D01: photon energy

$$E_\gamma = hc/\lambda, \quad N_\gamma = E_{\text{pulse}}/E_\gamma.$$

The energy relation uses vacuum wavelength. The legacy numerical examples insert nominal 337.1 nm and 777.4 nm: $5.8927495021 \times 10^{-19}$ J and $2.5552429343 \times 10^{-19}$ J per photon. This does not establish whether a particular spectroscopic bandhead or filter certificate was quoted in air or vacuum. A roughly 0.03% photon-energy change is small, while the roughly 0.1 nm line-position shift at 337 nm can matter to narrowband registration. Such metadata must be retained before a real spectral fold.

D02-D03: ordinary versus projected solid angle

For a circular cone of half-angle θ ,

$$\Omega = \int d\Omega = 2\pi(1 - \cos\theta), \quad \Omega_p = \int \cos\vartheta d\Omega = \pi\sin^2\theta.$$

Multiplying a flat physical aperture by ordinary solid angle instead of projected solid angle overcounts oblique uniform radiance. For a hemisphere, $\Omega = 2\pi$ sr and $\Omega_p = \pi$ sr. At 60 degrees full cone the values are about 0.842 sr and 0.785 sr.

A square field cannot be silently replaced by a circular cone with the same diagonal. For square tangent-plane half-width $u = \tan(\theta_{\text{diagonal}})/\sqrt{2}$,

$$\Omega_{\square} = 4\tan^{-1}\left(\frac{u^2}{\sqrt{1+2u^2}}\right), \quad \Omega_{p,\square} = \int_{-u}^u \int_{-u}^u \frac{dx dy}{(1+x^2+y^2)^2}.$$

The second integral is independently evaluated with 96 by 96 Gauss-Legendre nodes. The MMIA heritage field is described by its square diagonal; its nominal 1 mm ultraviolet entrance aperture must not be confused with its 25 mm photocathode. A proposed 25 mm entrance is a new design.

D04-D05: received fluence and background counts

For a small flat aperture and a point source at angle θ to its normal,

$$S = \Phi_{\text{ap}} A \eta \cos\theta.$$

Φ_{ap} is received photon fluence per area normal to propagation; if it was already defined per the physical tilted aperture, the projection is already included and must not be repeated. η is the end-to-end detection probability, not automatically bare quantum efficiency.

For uniform spectral energy radiance L_{λ} , constant response, a rectangular/equivalent band and a circular field,

$$b_{\text{sky}} = \frac{L_{\lambda} \Delta\lambda A \Omega_p \eta}{E_{\gamma}}, \quad B = (b_{\text{sky}} + d)T.$$

The general expression is an integral of projected effective area times radiance times $\lambda/(hc)$ over wavelength and solid angle. The simplified expression holds the photon energy and efficiency constant across the band. A Gaussian FWHM is not a rectangular width; a real source/filter/detector needs the full integral. A band-integrated Rayleigh radiance is converted by $1 \text{ R} = 10^{10}/(4\pi)$ photons $\text{m}^{-2} \text{sr}^{-1} \text{s}^{-1}$. At nominal 337.1 nm, $1 \text{ R} = 4.6893010583\text{e-}10 \text{ W m}^{-2} \text{sr}^{-1}$, integrated over the emission feature. Dividing by 4 nm is merely an equivalent band mean, not the line profile.

The background RMS $\sqrt{B}$ presumes ideal independent Poisson detections and known mean background. It is not a detection threshold, nor the uncertainty after fitting a finite baseline. The exact-count notebook supplies the inference procedure.

2. Optical and electrical design algebra

D06: etendue and filter size

Lossless geometrical optics preserves phase-space throughput,

$$G = n^2 \int_A \int_{\Omega} \cos\vartheta \, d\Omega \, dA.$$

For uniform air-to-air circular cones, $G = A_{\text{ent}} \pi \sin^2(\theta_{\text{field}})$. If the detector can accept an entire air hemisphere and the filter has internal half-cone beta,

$$A_{\text{det}} \geq A_{\text{ent}} \sin^2 \theta_{\text{field}}, \quad A_{\text{filter}} \geq \frac{G}{\pi \sin^2 \beta}, \quad d_{\text{filter}} \geq d_{\text{ent}} \frac{\sin \theta_{\text{field}}}{\sin \beta}.$$

A 25 mm entrance and 60 degree full circular cone already require at least 122.7 mm² detector area in this ideal limit. Restricting a filter to a 10 degree internal half-cone requires at least about 71.98 mm clear diameter. These are necessary bounds, not a lens prescription, a usable filter drawing or a 1U packaging proof. Refractive coupling changes the n^2 factors; real windows, numerical aperture, pupil geometry, vignetting and detector packaging require explicit treatment. Telecentric chief rays do not remove the marginal-ray cone.

D07: approximate interference-filter angle shift

Using a scalar effective index n_{eff} and incidence from air,

$$\lambda_c(\theta) = \lambda_0 \sqrt{1 - \sin^2 \theta / n_{\text{eff}}^2}.$$

This follows an ideal optical-thickness/cavity argument plus Snell refraction. The table assumes $n_{\text{eff}} = 1.8$; it is not a measured coating parameter. Real multilayer filters require angle-, polarization-, temperature- and environment-dependent measured transmission. A cone must be integrated over its ray distribution; evaluating only a chief ray is insufficient. A radially graded coating can compensate a position-dependent mean incidence but does not eliminate the spread of angles at a point.

D08-D09: a deliberately hypothetical Gaussian spectrum

Let the source be a normalized photon spectrum with standard deviation σ_s and let the unit-peak filter have standard deviation σ_f and centre offset δ . With $\text{FWHM} = \sqrt{8 \ln 2} \, \sigma$,

$$C = \frac{\sigma_f}{\sqrt{\sigma_s^2 + \sigma_f^2}} \exp \left[-\frac{\delta^2}{2(\sigma_s^2 + \sigma_f^2)} \right], \quad W_{\text{eq}} = \int T(\lambda) d\lambda = \sqrt{2\pi} \, \sigma_f.$$

These results come from completing the square in the product of two Gaussians. The verification additionally integrates that product with 256 Gauss-Legendre nodes in source-standard-normal coordinates z in $[-12, 12]$. Exactly 48 selected cases cover 3 source widths, 4 filter widths and 4 incidence angles (0,5,10,20 degrees). All 372 table rows are independently recomputed algebraically. The deterministic quadrature checks are not 48 observations.

A 4 nm unit-peak Gaussian FWHM has $W_{\text{eq}} = 4.2578680777$ nm. A background-dominated comparison proportional to $C/\sqrt{W_{\text{eq}}}$ assumes a fixed detected photon-background density per wavelength (or an explicit approximation holding $\lambda/(hc)$ constant). With exactly flat energy radiance instead, the wavelength factor should remain inside the integral. Signal shot noise, detector noise, real rotational populations, the asymmetric N2 band, field/cone mixing and actual coating curves are not represented. These curves do not establish the optimal nitrogen filter.

D10: electrical noise bandwidth is different

For a one-sided power spectral density and a unit-gain linear estimator,

$$\sigma_P^2 = \int_0^\infty \text{NEP}(f)^2 |H_{\text{est}}(f)|^2 df.$$

For a first-order low-pass, the noise-equivalent bandwidth is $\pi f_c/2$, not f_c . For a rectangular average of duration T , $H(f)$ is a phase factor times $\text{sinc}(\pi f T)$, giving one-sided ENBW = $1/(2T)$. Here ENBW has units Hz; spectral equivalent band W_{eq} above has units wavelength. They cannot be substituted for each other.

D11: out-of-band rejection

Let R be the unblocked out-of-band detected-photon-rate integral divided by the accepted in-band integral. If a uniform blocking factor at most 10^{OD} applies over that entire leak range,

$$b_{\text{leak}}/b_{\text{in}} \leq R 10^{-\text{OD}}, \quad \text{OD} \geq \max[0, \log_{10}(R/f_{\text{leak}})].$$

R must include the actual source spectrum, detector sensitivity, projected angle weighting and all optics. A narrow local OD6 certificate does not imply OD6 over the full visible/near-infrared sensitivity range. The 58 rows use assumed R and allowed fractions; they do not measure R for a polar cloud scene.

D12: dark-only NEP and multiplication

For primary dark rate d , mean multiplication M and second-moment excess-noise factor $F = E[G^2]/E[G]^2$, the ideal one-sided current PSD is $2 e^2 M^2 F d$. Dividing by responsivity $e M \eta/E_{\gamma}$ gives

$$\text{NEP}_{\text{dark}} = \frac{E_{\gamma}}{\eta} \sqrt{2 F d}.$$

Its units are $\text{W}/\sqrt{\text{Hz}}$. The gain cancels; gain helps overcome following amplifier noise but does not improve primary photon statistics. This model assumes all relevant dark pulses share the stated multiplication statistics. APD surface/bulk dark components, SiPM correlations, frequency-dependent noise, afterpulses, CMOS readout, amplifier voltage/current noise and background add terms with their own response. F and d here are explicit hypothetical inputs, not device-class specifications or end-of-life forecasts.

For an ideal average over T , this white-noise model gives $\sigma_P = \text{NEP}/\sqrt{2T}$, hence incident photon-equivalent RMS $\sqrt{(F d T)/\eta}$. Its agreement with direct count variance checks the one-sided convention. A finite off-background estimate adds uncertainty. In particular, a ten-to-fifty kilohertz sampling rate is not the estimator ENBW.

D13: high-pass droop

A single-pole AC-coupled voltage step decays as $\exp(-t/\tau)$. Requiring fractional plateau loss at T no greater than δ gives

$$\tau \geq \frac{-T}{\ln(1 - \delta)}, \quad f_c \leq \frac{1}{2\pi\tau}.$$

For 3 ms and 1% droop this is $\tau \geq 0.298497$ s and $f_c \leq 0.533187$ Hz. This describes an analog step plateau. AC shaping of individual photon pulses can still preserve a millisecond count-rate envelope if discrimination, count losses and recovery are characterized. The two architectures must not be confused.

D14: sample integration and phase

The test signal is a unit linear ramp from zero to one, with full ramp duration $D = T_{10-90}/0.8$, followed by an infinite plateau. Every sample is the exact average over its nonoverlapping exposure interval. The independent implementation

separates the ramp triangle from the post-ramp rectangle. It assigns the sample to the exposure midpoint and linearly interpolates the 10% and 90% crossings.

The 900 cases vary 3 true rises, 3 sample rates and 100 relative phases. They are deterministic estimator examples, not a physical classifier. At 10 kS/s, a 30 microsecond rise can produce an interpolated apparent rise around 134 microseconds even without noise. The interpolation can place a threshold before source onset; this is a biased reconstruction from averages, not acausal photons. At 50 kS/s the spacing is 20 microseconds. A 50 kHz analog passband exceeds the 25 kHz Nyquist frequency of that baseband sampling rate; the actual anti-alias filter and integration must be specified.

3. Clear-air reference calculation

P01: pressure-coordinate standard-atmosphere column

The retained USSA1976 model uses geopotential Earth radius $R = 6356766$ m, $g_0 = 9.80665$ m/s², $R^* = 8314.32$ J/(kmol K), $M = 28.9644$ kg/kmol; layer bases are $H = 0, 11, 20, 32, 47, 51, 71, 84.852$ km with lapse rates $-6.5, 0, 1, 2.8, 0, -2.8, -2$ K/km. Sea-level pressure and temperature are 101325 Pa and 288.15 K. These are a dry reference atmosphere, not polar weather or an invented vertical cloud profile.

For geometric altitude z , $H = R z / (R + z)$. Within a layer $T = T_b + L(H - H_b)$. Hydrostatic pressure obeys

$$p/p_b = (T/T_b)^{-g_0 M / (R^* L)} \quad (L \neq 0), \quad p/p_b = \exp[-g_0 M (H - H_b) / (R^* T_b)] \quad (L = 0).$$

Rather than repeat the original 1 m density trapezoid, the independent method uses hydrostatic balance to cancel temperature:

$$N_{\text{col}} = \int_z^\infty \frac{p}{k_B T} dz = \frac{R^*}{g_0 M k_B} \int_{p_{\text{top}}}^{p(z)} \frac{dp}{[1 - H(p)/R]^2} + N_{\text{tail}}.$$

$H(p)$ is analytically inverted separately in each lapse/isothermal layer. Integration uses 96-point pressure-coordinate quadrature per layer and is repeated with 192 points. The same small pressure-tail approximation above roughly 86 km is retained and exposed: $N_{\text{tail}} = R^* p_{\text{top}} / [g_0 M k_B (1 - H_{\text{top}}/R)^2]$. It is not a detailed thermosphere, nor an exact 550 km gas composition model. It contributes only a few parts per million of sea-level optical depth in this reference calculation.

Primary locators: USSA1976 Part1 tables2-4 and Part4 tables18-19 (curated CT041/S034); NASA/TM-2005-213659 AppendixA, equationsA1-A5 and pp34-35, <https://ntrs.nasa.gov/api/citations/20050207438/downloads/20050207438.pdf>. The legacy constants are retained so a reviewer can reproduce the original reference before exploring other atmospheres.

P02-P03: molecular scattering and ballistic dilution

Bodhaine et al.1999 equations29-30 provide the dry-air rational fit used in the original code. Writing $u = \lambda / (1 \text{ micrometre})$,

$$r(u) = \frac{1.0455996 - 341.29061/u^2 - 0.90230850u^2}{1 + 0.0027059889/u^2 - 85.968563u^2}, \quad \sigma_R = 10^{-32} r(u) \text{ m}^2.$$

The simple sea-level/reference-pressure fit is $\tau_R = 0.0021520 r(u) p/p_0$. The standard-density integral instead uses $\tau_R = \sigma_R N_{\text{col}}$ and direct no-collision transmission $\exp(-\tau_R)$. The source cross-section fit is separately traceable to CT029/S032, https://web.gps.caltech.edu/~vijay/Papers/Rayleigh_Scattering/Bodhaine-et-al-99.pdf. The

separate frozen CloudScat arithmetic uses its own refractivity/King-factor mixture, $N_0 = 2.6867811e25 \text{ m}^{-3}$ and 7200 m exponential scale height; the independent function retains those inputs rather than treating it as the same atmosphere.

The molecular column has units m^{-2} and sigma has m^2 , so tau is dimensionless. The direct aperture fraction for an isotropic point source is

$$f_{\text{direct}} = \frac{A}{4\pi r^2} \exp[-\tau_R - \tau_{\text{cloud}} - \tau_{\text{other}}].$$

Its fraction refers to all emitted photons, before detector efficiency. For a source at 16 km and satellite at 550 km, $r = 534 \text{ km}$ at nadir. The exponential concerns uncollided photons only. Scattered photons may return to the accepted direction; this calculation is not the total cloudy transfer, and $\tau_{\text{cloud_above}} = 0$ does not remove reflection/scattering from the cloud beneath the source.

The initial diagnostic comparison retained four failures on the small *difference* between standard and frozen transmissions: discrepancies were only about $1e-8$ percentage points, from the original density-grid discretization. Relative error on a nearly cancelled difference was an unsuitable test. The final check propagates the original recorded $1 \text{ m}^2/\text{m}$ column-refinement envelope to that difference and records both it and the 96/192 pressure-quadrature agreement. This is a numerical comparison tolerance, not a rigorous atmospheric uncertainty interval. The original failure receipt remains in `checks/diagnostics_initial_attempt.json`.

4. Ideal propagation and detector diagnostics

P04-P05: an above-plane reflector

A point source is height h above an infinite horizontal plane; the observer is H above that plane, with $0 \leq h < H$. The plane instantly reflects fraction ρ with BRDF ρ/π . For radial coordinate r , the incident photon fluence is $N h/[4 \pi (r^2 + h^2)^{3/2}]$. Including both plane and aperture projection factors gives reflected-to-direct directional ratio

$$B = \int_0^{H \tan \theta_{\text{FOV}}} \frac{2 \rho h H^2 (H - h)^2 r dr}{(r^2 + h^2)^{3/2} (r^2 + H^2)^2}.$$

The new implementation changes to $q = \ln(1+r/h)$, a log radial coordinate distinct from the original $x = h/\sqrt{r^2 + h^2}$. In the $h \rightarrow 0$ limit, $B = 2 \rho$ and total nadir gain is $1 + 2 \rho$; the all-direction upward escape fraction is only $(1 + \rho)/2$. Directional gain above one does not create photons or mean a photon probability exceeds one.

The extra arrival delay after subtracting common source-observer flight is

$$\delta(r) = \frac{\sqrt{h^2 + r^2} + \sqrt{H^2 + r^2} - (H - h)}{c}, \quad \delta_{\min} = 2h/c.$$

The code evaluates an algebraically stable positive version. A linear-ramp source with infinite plateau is convolved with this delayed reflection and normalized by its final plateau $1+B$. The reported rise is the difference of the *amplitude* crossings, not delay quantiles or onset-to-peak. All 108 rows are independently recomputed, including the rise roots. No cloud residence time, lateral transport, real BRDF, Earth curvature, gas extinction, finite 3 ms pulse end or noise is included. ρ is not inferred from overhead cloud optical depth. The $h=0$ surface limit is not a real embedded cloud-top boundary condition.

BRDF source: R4-LAMBERT/S112, https://www.pbr-book.org/4ed/Reflection_Models/Diffuse_Reflection, Lambert model and radiance normalization. All plane geometry/convolution equations above are explicit derivations for this diagnostic.

P06: diffusion first passage, absorption and degeneracy

For an instantaneous source at depth L in the half-space $z > 0$ with absorbing plane $z=0$, homogeneous diffusion D and uniform killing rate γ ,

$$\partial_t u = D \partial_z^2 u - \gamma u, \quad u(0, t) = 0, \quad u(z, 0) = \delta(z - L).$$

Subtracting the image Gaussian and taking outward flux at the boundary yields

$$h(t) = \frac{L}{\sqrt{4\pi D}} t^{-3/2} \exp[-L^2/(4Dt) - \gamma t], \quad t > 0.$$

h has units s^{-1} . Its total integral is $\exp[-L \sqrt{\gamma/D}]$, not generally one. For $\gamma=0$, its CDF is $\operatorname{erfc}[L/\sqrt{4Dt}]$. For $\gamma > 0$ the finite-window CDF is

$$F(T) = \frac{1}{2} \left[e^{-b} \operatorname{erfc}(a - q) + e^b \operatorname{erfc}(a + q) \right], \quad a = \frac{L}{\sqrt{4DT}}, \quad q = \sqrt{\gamma T}, \quad b = L\sqrt{\gamma/D}.$$

The code independently integrates h using $u = L/\sqrt{4Dt}$, then $v = \ln u$, and checks the all-time and finite-time closed forms for all 48 rows. The quoted simple exponential tail formula is not used to infer real ice or ozone absorption.

The peak solves $\gamma t_{\text{peak}}^2 + (3/2)t_{\text{peak}} - L^2/(4D) = 0$. For no absorption $t_{\text{peak}} = L^2/(6D)$. With $x = t/t_{\text{peak}}$,

$$h/h_{\text{peak}} = x^{-3/2} \exp[-1.5(1/x - 1)], \quad T_{10-90} = 0.07432773536 L^2/D.$$

The coefficient comes from the two *leading amplitude* roots, not cumulative probability. Bai et al.2023 (10.1029/2022JD037460, sections 6-7; curated bai2023-bath-waveform/S100) quotes a rounded coefficient about 0.0742 and context for D . The present method-of-images derivation is explicit and does not claim an uninspected equation bitmap was read.

Exactly the same L^2/D and γ give the same whole $h(t)$, including amplitude. For example, $(L, D) = (1 \text{ km}, 3e9 \text{ m}^2/\text{s})$, $(2 \text{ km}, 1.2e10)$ and $(3 \text{ km}, 2.7e10)$ are indistinguishable in this ideal kernel. External transport/source information is necessary to break that degeneracy.

The diffusion PDE is parabolic and has nonzero response for all positive t , including $t < L/c$. It is not a causal early-time radiative-transfer benchmark. Rows flag cases where even the 10% threshold precedes straight light flight. Equivalent transport mean free path $3D/c$ and depth ratio $L/(3D/c)$ are diagnostics, not measured cloud microphysics. A real finite cloud, ballistic light, phase function, boundary, finite source duration and noise require other calculations.

P07: counter dead time can reshape an uncorrected rise

For a stationary nonparalyzable counter, one renewal cycle has mean live waiting $1/r$ plus dead time τ , giving

$$m(r) = \frac{r}{1 + r\tau}, \quad r = \frac{m}{1 - m\tau}, \quad \frac{dr}{dm} = \frac{1}{(1 - m\tau)^2}.$$

With true rate $b+S f(t)$, subtract the registered baseline and normalize to registered peak. For a linear input ramp and $x = S\tau/(1+b\tau)$,

$$y = \frac{f(1+x)}{1+xf}, \quad f(y) = \frac{y}{1+x(1-y)}, \quad \frac{T_{10-90,\text{out}}}{T_{10-90,\text{in}}} = \frac{1+x}{1+x+0.09x^2}.$$

The independent code takes each measured amplitude level and inverts the mean rate directly instead of reusing that closed shape expression. All 84 rows are checked. The example $\tau=17 \text{ ns}$ is retained as a specific study parameter; it is

not a universal detector dead time. This quasi-static mean approximation assumes rise scales much longer than tau. It does not simulate event likelihood, paralyzable hold-off, afterpulses, threshold walk, gain sag or electronics. Inversion does not recover lost photons or restore Poisson output statistics. It is not evidence that published *corrected* MMIA waveforms suffer the raw bias.

Architecture context: Chanrion et al.2019 MMIA, DOI10.1007/s11214-019-0593-y, section6.3; curated MMIA architecture source S066 and the retained code/parameter audit.

P08-P09: finite spherical field and conditional source weights

The original toy uses Earth radius 6371 km, source shell at17 km, satellite at550 km. For observer off-axis angle theta, the cloud-cap central angle is psi = asin[(r_s/r_c)sin theta]-theta, provided the ray intersects the near shell. In independent shell coordinates,

$$r^2 = (r_s - r_c)^2 + 4r_s r_c \sin^2(\psi/2), \quad \cos\theta = \frac{r_s - r_c \cos\psi}{r}, \quad dA_c = 2\pi r_c^2 \sin\psi \, d\psi.$$

The cap area is $4\pi r_c^2 \sin^2(\psi_{\max}/2)$. The equivalent observer-ray Jacobian is $dA_c = r^2 d\Omega_{\mu_c}$; the independent implementation does not use that integration coordinate.

For an isotropic point source with photons N in the accepted band/window and no cloud or gas attenuation,

$$S = \frac{\eta A \cos\theta}{4\pi r^2} N, \quad N_{\text{cut}}(\theta) = \frac{4\pi r^2 S_{\text{required}}}{\eta A \cos\theta}.$$

With an *assumed* cumulative source law $C N^{-\alpha}$, the tabulated weight is

$$W = \int_{\text{cap}} [N_{\text{cut}}(\theta)/N_{\text{pivot}}]^{-\alpha} dA_c.$$

W has area units. The pivot 10^{20} photons is an arbitrary normalization, not an observed source yield. S_{required} is the supplied ideal-count 95%-power cut; this module does not fit or independently verify it. W is a bright-enough population weight, not the full integral of detection probability. Below-cut detections are omitted. Conditional on that threshold and ideal monotone detection model, the above-cut contribution alone lies between 0.95 and 1 times W times its unknown normalization. Actual rates require source geography, brightness law, occurrence, live exposure, SAA/attitude masks, background distribution, association and useful-science acceptance.

Only in a local small-field, uniform, background-limited approximation does W scale as diameter^{alpha} times Omega^(1-alpha/2). The alpha=2 crossover is not a universal wide-field optimum. All 216 table rows are checked in shell-central-angle coordinates, including background counts, cap area, on-/off-axis source thresholds, weights and their ratios. Exact count-threshold columns are explicitly excluded from this module's independent scope and cross-linked to the statistics notebook.

5. Reproducibility and what these checks establish

check_design() recomputes all 1,639 design rows, constants, 48 independent Gaussian quadratures and hand/unit limits, including explicit SI-exponent bookkeeping. That bookkeeping checks declared equations, not arbitrary runtime quantities supplied with undeclared units. check_diagnostics() independently checks 5 Rayleigh altitude rows, 108 plane rows, 48 diffusion rows, 84 dead-time rows, 216 field/population rows and 192 auroral count-conversion rows. It preserves the source units and declared hypothetical parameters. Comparison outputs record row/column errors, tolerances, input hashes and the module hash. New numbers are calculated by the code; comparison arrays are not copied as results.

The auroral conversion checks use the supplied detected S as an input; they recompute radiance-to-rate, Poisson RMS *under its assumptions*, photon/current/fluence conversions and occupancy. They do not establish a moderate/disturbed

polar radiance distribution. Likewise, the Noll/Bianchi actual raw spectral inputs are absent from this curated verification bundle; their original integrations are not newly reproduced here. The formulas and interpretation limits are provided, while true source re-integration requires the calibrated arrays and wavelength convention.

Passing proves conditional algebra/numerical agreement within declared tests. It does not validate a real anvil, a source luminosity function, a rise-to-depth classifier, actual receiver NEP, out-of-band rejection, end-of-life radiation behavior, a current calibration or useful polar event rate. The cloud and statistics notebooks provide separate algorithmic tests and retained convergence/selection limits.

6. Comparison strength and deliberately wrong controls

Tiny positive SI outputs (photon energies, NEP, aperture fractions, photocurrents and etendue) are checked with relative tolerances and zero generic absolute floor. A generic $1\text{e-}12$ absolute floor would make a zero answer pass for many such values and is not used in the final receipts. Examples of explicit absolute tolerances include near-zero wavelength shifts ($1\text{e-}12$ nm), negligible Gaussian capture fractions ($2\text{e-}14$), sample-crossing roundoff ($2\text{e-}7$ microseconds), original reflection-root resolution ($5\text{e-}6$ microseconds), and the documented Rayleigh grid-refinement residual. Their units and ranges are recorded per comparison field. These are algorithm-comparison tolerances, not physical uncertainty statements.

Six negative controls replace selected comparison targets in memory with zero: photon energy, dark-only NEP, Gaussian capture, reflected aperture fraction, projected etendue and photocurrent. The checks must FAIL on those selected quantities; the control receipt is PASS only when the fault is rejected. No source files are altered. Reproduce with `python checks/design_negative_controls.py` and `python checks/diagnostics_negative_controls.py`. Initial comparison-policy and initial Rayleigh-failure receipts are preserved with their superseded status; final checks and negative controls are authoritative for this verification module.

Nonfinite fresh or reference values are rejected explicitly. Relative-error summaries exclude a zero reference and instead report its absolute error and a zero-reference counter; this avoids meaningless relative errors from division by a tiny artificial denominator. A deliberately nonfinite guard test must be rejected.

Scalar time-dependent transport: specialist derivation and verification

This note describes the executed R4 finite-cylinder model and what the light review independently recomputes. It is not a new physical validation. The tabulated-GHM point-nadir estimator remains nonuniformly converged. This is an expected warning, not a pass condition to suppress. The raw ice LUT is deliberately absent.

Symbols and units

Symbol	Meaning	SI units
$L(r, \hat{\Omega}, t)$	photon radiance propagating along $\hat{\Omega}$	photons m ⁻² s ⁻¹ sr ⁻¹
$\Phi_s(t), N_s$	emitted photon rate and event photon count	photons s ⁻¹ ; photons
$\beta_e, \beta_s, \beta_a$	extinction, scattering, absorption coefficients	m ⁻¹
$\omega = \beta_s / \beta_e$	single-scattering albedo	dimensionless
$p(\mu), f(\mu)$	phase density per solid angle; density per cosine	sr ⁻¹ ; dimensionless per unit μ
g	mean scattering cosine	dimensionless
A, n_A	small planar collecting area and normal toward scene	m ² ; unit vector
R_0, r_j, ℓ_j	direct range, collision-observer range, accumulated source-collision path	m
$K(\Delta), h(\Delta)$	aperture-area transfer response and vacuum-normalized response	m ⁻² s ⁻¹ ; s ⁻¹
m_{bi}	relative response MASS in time bin i of batch b	dimensionless
$\delta t, B, N_b$	native bin width, independent batches, histories per batch	s; counts; counts

RTE and boundary/source statement

For scalar, elastic, stationary-medium transport:

$$\left(c^{-1} \partial_t + \hat{\Omega} \cdot \nabla \right) L = -\beta_e L + \beta_s \int_{4\pi} p(\hat{\Omega}' \cdot \hat{\Omega}) L(r, \hat{\Omega}', t) d\Omega' + \frac{\Phi_s(t)}{4\pi} \delta^3(r - r_s). \quad (\text{R1})$$

The isotropic point source emits into all 4π. Coefficients are constant inside the cylinder and zero outside. There is no incoming exterior illumination, ground reflection or reflecting cloud boundary. Exterior photons move in vacuum. The cylinder is convex, so a ray that has left cannot re-enter without a change of direction outside, which this model excludes. An above-top downward launch may enter the cloud before leaving again.

The main grid assumes $\omega = 1$, hence $\beta_a = 0$. This is a model input, not measured UV ice absorption. There is no molecular scattering, ozone, aerosol, refraction or Earth curvature. A full physical atmosphere would need additional spatially varying scattering/absorption terms and their path delays.

Cloud base 14 km, top 17 km and radius 30 km are example geometry. On-axis sources are 15/16/17/18 km; the observer is 550 km and has a 60 degree full circular FOV. The area response uses a flat horizontal aperture. A finite source, real cloud structure and off-axis population are not represented.

Phase conventions and actual input

$$\int_{4\pi} p \, d\Omega = 1, \quad f(\mu) = 2\pi p(\mu), \quad \int_{-1}^1 f(\mu) \, d\mu = 1, \quad g = \int_{-1}^1 \mu f(\mu) \, d\mu. \quad (\text{P1})$$

Azimuth is uniform for the randomly oriented scalar model. The supplied P11 convention has integral2 in cosine, so normalizing its sampled representation and dividing f by 2π supplies the probability per steradian. P11 itself must not be used as p without normalization.

Native330/340 nm P11 at effective radius 10 μm is linearly interpolated in wavelength to the numerical coordinate 337.1 nm. The resulting density is represented piecewise linearly in μ , normalized, and its inverse CDF solved segment by segment. Sampling and next-event evaluation use this identical representation. Its computed $g = 0.7709446878541932$. The HG comparison is

$$f_{HG}(\mu) = \frac{1 - g^2}{2(1 + g^2 - 2g\mu)^{3/2}}, \quad \langle P_2(\mu) \rangle_{HG} = g^2. \quad (\text{P2})$$

The review independently integrates HG normalization/moments, checks its inverse CDF, and draws seeded angles. It does **not** pretend to recompute GHM normalization without the raw LUT. The processed ice-habit mixture, roughness, size-distribution applicability, interpolation and omitted polarization remain physical input uncertainties. The molecular band and air/vacuum wavelength registration are not folded here.

Raw dataset: https://docs.hygeos.com/public.php/webdav/ic_baum_ghm_sol.nc ; developer information: https://www.ssec.wisc.edu/ice_models/polarization.html . Original full NetCDF SHA256 is 2b37e27cfff0dd7e048df16c5fb612cd7a83e3c657c851bb50e758e11e870e; executed private bracket SHA256 is 2cb7bca802632479473a0812912ec8007fc69087ade2767d3c06cc3e7cb81a02. Rights to redistribute the raw LUT were not established. A reproducible private extraction recipe is provided in the separate R4 code bundle; obtaining that input is an optional prerequisite for a full RT rerun, not for these light checks.

Direct contribution: deterministic and independently testable

For the on-axis vertical source-observer segment, define $H = z_t - z_b$:

$$\beta_e = \tau_{total}/H, \quad \ell_{above} = \max[0, \min(z_t - z_s, H)], \quad \tau_{above} = \beta_e \ell_{above}. \quad (\text{D1})$$

$$K_{vac} = \frac{1}{4\pi R_0^2}, \quad K_0 = K_{vac} e^{-\tau_{above}}, \quad R_0 = z_o - z_s. \quad (\text{D2})$$

K_0 has units m^{-2} per emitted photon; the exact unscattered delay response is $K_0 \delta(\Delta)$. For an off-axis source the receiver cosine and actual finite cloud-segment length must be included; D2 is this grid's on-axis specialization. The review recomputes D1/D2 for all 56 cases from geometry and total tau, not by copying the stored direct value.

At 15 km overlying tau is 2/3 of total tau; at 16 km it is 1/3; at/above 17 km it is 0. Varying total tau changes cloud below an internal source as well as above it. Thus this sweep is not a universal relation in overlying tau alone. The total directional response may exceed vacuum because downward emission is redirected upward; it is not an escape probability and need not be at most 1.

Next-event estimator and source weights

The physical source angular density is $q_0 = 1/(4\pi)$. If an initial direction is sampled from another normalized density q, use

$$w_{launch} = q_0/q. \quad (N1)$$

The production internal-source proposal mixes isotropic and zenith-centered phase densities with equal weights. Its likelihood weight is bounded above by 2. This affects launch efficiency only; physical scatter directions retain the chosen phase law. At/above the cloud top the production launch is isotropic. The true source remains isotropic in every case.

For collision j , $\hat{u}_j = (r_o - r_j)/r_j$ points **toward** the observer and $\mu_j = \hat{\Omega}_{in,j} \cdot \hat{u}_j$. With a receiver normal toward the scene,

$$c_{A,j} = \max(0, -n_A \cdot \hat{u}_j). \quad (N2)$$

A small receiver subtends $Ac_{A,j}/r_j^2$. Let χ_{FOV} be its circular acceptance. If w_j is weight **before** implicit absorption at collision j ,

$$\hat{K}_{scat} = \frac{1}{N_b} \sum_{history} \sum_j w_j \omega p(\mu_j) \frac{c_{A,j}}{r_j^2} \exp[-\beta_e \ell(r_j \rightarrow r_o)] \chi_{FOV,j}. \quad (N3)$$

This is per unit area. Multiply by A once to obtain a collected fraction. The segment length in the exponential is the intersection of the **finite collision-observer segment** with the cloud, not an infinite ray or a vertical shortcut. In the main conservative case omega 1 contributes no extra factor. In an implicit-absorption implementation that has already multiplied w by omega before evaluating N3, omit the displayed additional omega. Double counting it would be wrong.

The direct contribution D2 is added exactly once per emitted history ensemble, not at every collision. A source-only proposal needs no extra 4π multiplier beyond N1. A phase function per steradian and its cosine density are not interchangeable. Aperture cosine, area and QE/PDE also must not be counted twice; QE/PDE is absent from this transport calculation.

Physical delays, impulse atom and native histogram

$$\Delta_j = (\ell_j + r_j - R_0)/c \geq 0, \quad t_{direct} = R_0/c. \quad (T1)$$

The accumulated physical path includes an external source's vacuum travel to cloud entry. Boundary numerical tolerances must not create physical path or delay. The triangle inequality supplies causality; a diffusion PDE is a different approximation and can have pre-flight tails.

Let $h = K/K_{vac}$. For native bin i , the stored value is

$$m_{bi} = \int_{i\delta t}^{(i+1)\delta t} h_{scat,b}(\Delta) d\Delta. \quad (T2)$$

The NPZ arrays contain scattered **mass**, not density, and a final overflow element for delays at/after 12 ms. Their ordinary bins are 2 μs wide. Mean density is mass divided by bin width in seconds. The direct atom has relative mass $e^{-\tau_{above}}$, supplied separately. In the reviewed convolution it is at 0, while ordinary scattered mass is placed at bin centers 1,3,5,... μs on a 1 μs grid. Plot bins later coarsen only the tail. The native arrays, rather than plotted/rebinned arrays, are used for this review's fluence reconciliation.

Zero observed overflow is an empirical statement about these histories, not proof that every possible path arrives within 12 ms. Truncated packet weight is separately recorded. No production packet reached the 20000 scatter safety cap.

Pulse convolution, retarded time and rise definitions

With retarded time $u = t_{absolute} - t_{source,onset} - R_0/c$,

$$\frac{N_{ap}(u)}{A} = \int_0^\infty K(\Delta) \Phi_s(u - \Delta) d\Delta. \quad (T3)$$

The assumed unit-peak source shape is

$$s(u) = C_s e^{-u/\tau_d} (1 - e^{-u/\tau_r}) H(u), \quad \tau_d = 3 \text{ ms}, \quad u_{peak} = \tau_r \ln(1 + \tau_d/\tau_r). \quad (T4)$$

Cs normalizes the source's own peak to 1. tau_r is calibrated to intrinsic rising 10–90% times 30/80/200 μs, relative to that peak. The 3 ms is an assumed exponential decay time; it is not a fitted 3 ms duration or a measured intrinsic NBE waveform. In a lossy/broadened pulse, renormalizing each plotted curve to its own peak is useful for shape only; compare the separately stored peak gain/fluence for photometry.

An optional ideal first-order response is $e(u) = 2\pi f_c e^{-2\pi f_c u} H(u)$, with f_c representing a delta. The mean waveform is $s * h * e$. Analog cutoff is not sample rate, and this does not simulate actual post-counting noise correlations or a complete anti-alias chain.

The reported metrics are separately $t_{90} - t_{10}$ and $t_{peak} - t_{onset}$. Onset is the known model onset 0. A vanishingly weak direct ray can define that mathematical onset even when it is unobservable. Heritage observed rise numbers cannot be inserted as intrinsic 10–90% parameters without a source/detector fit. No population classifier is validated here.

Independent batch uncertainty and covariance

For gain $G_b = G_0 + \sum_i m_{bi}$, with independent equal-size batches,

$$\tilde{G} = B^{-1} \sum_b G_b, \quad SE(\tilde{G}) = \sqrt{\frac{\sum_b (G_b - \tilde{G})^2}{B(B-1)}}. \quad (U1)$$

$$\text{Var}(\tilde{G}) = \sum_i \sum_j \text{Cov}(\tilde{m}_i, \tilde{m}_j). \quad (U2)$$

The direct term is deterministic, so it adds no MC variance. A history can contribute to several bins/collisions. Adding only squared bin errors omits covariance and is generally wrong. The review computes the sum-covariance identity from the $B \times J$ centered matrix without allocating a $J \times J$ covariance matrix, and compares it with the stored fluence SE. It also displays the incorrect independent-bin estimate for contrast.

Finite-batch SE is not a guaranteed confidence interval for an underresolved heavy-tail estimator. For the 15 km/tau 10 GHM control, increasing from 300000 to 1200000 histories changes the estimated gain noticeably. A 1 degree cone proposal preserves isotropic source moments but does not remove multi-scatter variance; some batches concentrate a large share of their estimate in a few histories. These are retained convergence failures, not evidence of a physical GHM/HG timing difference. Timing SEM from individual batch waveforms is a nonlinear diagnostic, especially for an argmax; it is not automatically the sampling SE of the timing functional of the averaged kernel.

Executable equation-based first-scatter check

For an isotropic phase, axisymmetry lets azimuth integrate out. At initial cosine mu, let $se(\mu)$ be the distance to the cylinder boundary. For a first collision at distance s, let $r(s, \mu)$ and ℓ_{out} be its observer range and remaining cloud path. The relative first-order fluence is

$$G^{(1)} = \frac{R_0^2}{2} \int_{-1}^1 d\mu \int_0^{se(\mu)} \beta_e e^{-\beta_e [s + \ell_{out}(s, \mu)]} \frac{c_A(s, \mu)}{r(s, \mu)^2} ds. \quad (S1)$$

The review implements this integral independently with Gauss quadrature, splitting the cosine intervals where side and horizontal boundary paths exchange. Order 24/48/96 results and the retained isotropic MC first-collision estimate are compared. This is a first-collision equation check, not an independent full RT engine or a GHM check.

What the review checks and does not check

All 56 direct terms, native batch mass/gain/SE, nonnegative bins, overflow, packet-count accounting and weighted boundary/escape bookkeeping are recomputed from packaged inputs. The source-importance launch weights average 1 statistically and are not forcibly normalized. The review checks that escaped plus explicitly truncated weight equals launch weight, not that every finite importance sample sums exactly N.

Raw metadata reports no noncausal paths and no truncation. The individual three-dimensional histories are not retained, so those per-path geometry statements cannot be independently replayed from a histogram. The code, source input and seeds would be needed for replay. No original RT run is triggered by this notebook.

Two earlier parameter-matched ideal-HG CloudScat comparisons agreed within their combined MC errors; see `cloud_validation_summary.json`. They test these ideal HG geometries only, use a shared timing postprocessor, and do not validate GHM, real atmosphere or a source-depth classifier. CloudScat primary method:

<https://gmd.copernicus.org/articles/13/5549/2020/> .

The executable reviewer maps these equations to exact function names and source lines in `reference/cloud` and writes `checks/transport_code_map.json`. It records file hashes. The original executed engine/runner snapshots, when supplied, must be distinguished from later canonical files that added diagnostics. Do not infer byte-identical execution from a similar filename.

Optional full transport reruns require the lawful raw bracket and the separate reproduction launcher. They are deliberately outside the default light notebook. Empirical cloud microphysics/structure, absorption/gases, real source spectrum and waveform, noisy onset and labelled-event inference remain open.

Worked-history ledger

The notebook also shows a fresh HG history with fixed seed 2026102202, source 16km, total tau 10 and the same mean cosine as the main model. Six collisions occur before escape. Each row records launch/propagation weights, an exponential free optical path, the 3D collision position, incoming and outgoing unit vectors, the phase probability toward nadir, observer optical depth and survival, range, aperture cosine, delay, and raw next-event summand. The ledger is an explanatory example, not an original production packet or a reported event-fluence estimate. Its end status distinguishes physical escape from a deliberate illustration cap.

Statistical verification: detection, pulse shape and a red upper limit

This chapter specifies the probability experiment before giving numerical checks. The measurements below are synthetic receiver-input photon counts. A “30 μs ”, “80 μs ” or “200 μs ” label describes the expected received waveform. It is not an independently known cloud depth, discharge current or physical source class.

The notebook **03_statistics.ipynb** executes a new worked likelihood fit and rebuilds the reported decisions from all 81 saved trial summaries. It does not rerun the original 81,000-pulse generator by default. The original scripts are retained under reference/statistics for comparison; the verification module does not import their fitting or interval functions.

Reference planes and units

The optical entrance is the reference plane after atmospheric propagation. Let Φ_{ap} be event photon fluence at that plane, in photons m^{-2} , and A_{ent} the effective entrance area in m^2 . For a specified spectral, angular and temporal response η , the expected primary detected count is

$$E_{\gamma}(\lambda) = hc/\lambda, \quad s = \eta A_{\text{ent}} \Phi_{\text{ap}}. \quad (\text{S01})$$

Here s is an expectation in photoelectrons or registered primary photon events, before PMT/APD charge multiplication. It is not a realized count and does not include background. For a real broadband/extended scene, replace the scalar η by the appropriate spectral-angular-time integral. Internal charge gain is not another photon-detection efficiency.

The reverse conversion is

$$N_{\text{ap}} = s/\eta, \quad \Phi_{\text{ap}} = s/(\eta A_{\text{ent}}), \quad F_{E,\text{ap}} = E_{\gamma} \Phi_{\text{ap}}. \quad (\text{S02})$$

These equations do not infer total emitted photons. A spacecraft-plane observed fluence has already undergone cloud and atmospheric transport; do not attenuate it again. The module imports h , c and the two photon energies from modules/common.py.

State the on/off experiment

In this section b is the expected background **count in the on window**, not a count rate. Let $s \geq 0$ be the expected source count in that window and $r = T_{\text{off}}/T_{\text{on}} > 0$ the known relative exposure. A stationary background and disjoint, uncontaminated windows give

$$N \sim \text{Pois}(s + b), \quad M \sim \text{Pois}(rb), \quad N \perp M. \quad (\text{S03})$$

The joint likelihood, with integer observations n and m , is

$$L(s, b) = \frac{e^{-(s+b)} (s+b)^n}{n!} \frac{e^{-rb} (rb)^m}{m!}. \quad (\text{S04})$$

This model excludes deadtime, saturation, afterpulses, correlated avalanches, background drift and source contamination of the off window. An effective exposure ratio can include differing throughput or live time only when that correction is justified. A finite prewindow is not an externally known background.

Known-background detection and power

If b is known exactly and fixed before the test, choose the smallest integer $k \geq 1$ satisfying

$$P_b(N \geq k) \leq \alpha. \quad (\text{S05})$$

Declare a positive excess when $N \geq k$. The count discreteness generally makes the actual size smaller than α . At $b=0$, $k=1$ for every $0 < \alpha < 1$. This zero-background control is mathematical, not a claim that a spacecraft detector has no false counts.

For a requested detection probability p_{det} , solve for the expected source count:

$$P_{b + s_{p_{\text{det}}}}(N \geq k) = p_{\text{det}}. \quad (\text{S06})$$

The notebook uses $p_{\text{det}}=0.95$. It verifies the returned integer, the preceding integer and the achieved power. This is a specified-gate signal-count requirement, not a mission completeness curve.

Unknown background: exact conditional on/off test

Under the null $s=0$, condition on the observed total $t=N+M$:

$$N \mid (N + M = t) \sim \text{Binom}\left(t, \frac{1}{1+r}\right), \quad p_{\text{on/off}} = P(X \geq n \mid t). \quad (\text{S07})$$

For every fixed t the rejection probability is at most α ; averaging over t retains that bound for every constant unknown b . The condition is on **N+M**, not on M alone. A displayed minimum on-count at a fixed observed off-count is a convenient boundary representation, not a conditional false-alarm guarantee given M .

The review independently sums the finite binomial tail for selected examples and compares it with SciPy's survival function. It does not insert M/r into a known-background test.

Search trials

If K is a predetermined number of tested hypotheses, or a deterministic upper bound, the union bound and linearity of expectation give

$$P(\text{one or more false candidates}) \leq K\alpha, \quad E[N_{\text{false candidates}}] \leq K\alpha. \quad (\text{S08})$$

Independence between overlapping tests is unnecessary for this bound. A retrospectively observed stopping count is not automatically a valid K . The illustrative allocation $\alpha=0.1/(50,000 \times 86,400 \times 3)$ assumes that particular number of evaluated starts and window families. It is neither an empirically measured event rate nor a complete model of an adaptive spacecraft trigger, clustering or particle bursts.

A red nondetection is an upper confidence limit

For a Poisson total mean μ , the one-sided Garwood upper limit is

$$U(n; \alpha) = \frac{1}{2} \chi_{2(n+1), 1-\alpha}^2, \quad P_{\mu} = U(X \leq n) = \alpha. \quad (\text{S09})$$

Its repeated-sampling coverage is at least $1-\alpha$; discreteness can make it conservative. With externally known b ,

$$s_{\text{up}} = \max[0, U(n; \alpha) - b]. \quad (\text{S10})$$

For $s > 0$, this covers s whenever $U(N; \alpha)$ covers $s + b$. At $s = 0$ the floor always covers. A downward fluctuation relative to a known b can produce an upper bound of zero; it is not a posterior assertion that a source is impossible.

For a finite independent off window, choose $\alpha_t + \alpha_b = \alpha$ **before examining the data**. Define the lower bound

$$L(m; \alpha_b) = \begin{cases} 0, & m = 0, \\ \frac{1}{2} \chi^2_{2m, \alpha_b}, & m > 0. \end{cases} \quad (\text{S11})$$

Then

$$s_{\text{up, off}} = \max[0, U(n; \alpha_t) - L(m; \alpha_b)/r] \quad \text{covers with probability at least } 1 - \alpha_t - \alpha_b. \quad (\text{S12})$$

On the event that both component bounds cover, $U \geq s + b$ and $L/r \leq b$, so their difference is at least s . The probability either component fails is bounded by $\alpha_t + \alpha_b$. Independence of the two bound-failure events is unnecessary for this union-bound proof; the observation model in S03 still specifies independent windows.

At $n = 0$ and known $b = 0$, the 95% upper mean is $-\ln(0.05) = 2.995732$ counts. At $n = m = 0$, the equal-split conservative finite-baseline construction gives $-\ln(0.025) = 3.688879$ counts. These are different stated procedures, not conflicting physical measurements. The latter is a conservative bound, not an “exact joint 95%” interval or a Bayesian probability.

Unconditional coverage need not remain nominal after selecting only nondetections, accepted shapes or extreme ratios using the same data. An independent trigger can avoid some selection effects, but its independence must be demonstrated. A red upper limit is on the expected registered source count in the specified window. It requires calibrated response and live time to become an aperture-photon bound, and a separate physical lower red/UV bound to exclude a thermal alternative. No universal positive thermal 777/337 bound is supplied by these statistics.

Unknown amplitude, onset and background

For the waveform calculation, use β for a background **rate** in counts s^{-1} . For hardware bin j of duration d_j in seconds, let $q_j = \int_{\text{bin}} f(t - t_0; R) dt$, also in seconds, where f is a unit-peak pulse template. A is the peak source count rate in counts s^{-1} . The bin expectation is

$$\lambda_j = A q_j + \beta d_j, \quad n_j \sim \text{Pois}(\lambda_j) \quad \text{independently.} \quad (\text{S13})$$

The background prewindow is another bin with $q_{\text{pre}} = 0$. The unit-peak template rises linearly over $D_{\text{ramp}} = R_{10} - 90/0.8$, stays on a plateau, and ends 3 ms after onset. Exact primitive differences give q_j ; a midpoint sample is not substituted for the integral.

For each candidate onset and rise, define $Q = \sum q_j$, $D_{\text{tot}} = \sum d_j$, $h_j = q_j/Q$ and $w_j = d_j/D_{\text{tot}}$. With $\eta_{\text{tot}} = A Q + \beta D_{\text{tot}}$ and $\theta = A Q / \eta_{\text{tot}}$,

$$\ell = N \log \eta_{\text{tot}} - \eta_{\text{tot}} + \sum_j n_j \log[(1 - \theta) w_j + \theta h_j] - \sum_j \log(n_j!), \quad N = \sum_j n_j. \quad (\text{S14})$$

Both h and w sum to one. Therefore, for $N > 0$, the maximum in total scale is $\hat{\eta} = N$. The remaining one-dimensional objective is concave:

$$g'(\theta) = \sum_j \frac{n_j (h_j - w_j)}{(1 - \theta) w_j + \theta h_j}, \quad g''(\theta) = - \sum_j \frac{n_j (h_j - w_j)^2}{[(1 - \theta) w_j + \theta h_j]^2} \leq 0. \quad (\text{S15})$$

Check both endpoints and solve the interior derivative root where appropriate. At $N=0$ the mixture is unidentifiable and $A=\beta=0$ suffices. The fitted rates are

$$\hat{A} = \hat{\theta}N/Q, \quad \hat{\beta} = (1 - \hat{\theta})N/D_{\text{tot}}. \quad (\text{S16})$$

The new worked example saves every n_j , d_j , q_j and fitted λ_j . It compares this profile solution with a separate two-parameter SciPy optimization of the original $A q_j + \beta d_j$ likelihood. The original R4 grouped/vectorized implementation is not imported. Agreement checks implementation of the algebra; it cannot validate the physical pulse family.

The practical study searches a declared onset grid using early rise bins and a late plateau, with a 10 ms source-free prewindow. It excludes the final pulse cutoff from the shape fit. A coarse onset bracket is provided externally; blind search and its false-candidate rate are not established.

Detection, acceptance and shape error have different denominators

The rule first detects a fixed candidate gate. It then selects a shape and rejects insufficiently separated profile scores:

$$\text{accept} = \text{detected} \wedge [\ell_{\text{winning profile}} - \ell_{\text{competing profile}} \geq \ln 19]. \quad (\text{S17})$$

This margin was fixed in advance. It is not calibrated 95% confidence, a 95% correctness guarantee or a posterior probability. Maximizing over unknown parameters and unequal template families does not supply a normalized likelihood under a specified population prior.

The saved trials support three separate quantities:

- detected / all trials;
- accepted / all trials;
- wrong / accepted trials.

Always read error among accepted cases together with acceptance and rejection. An error fraction based on zero accepted cases is undefined, not zero. A sparse accepted subset has a wide uncertainty interval even if it contains no observed mistakes.

The review rebuilds all decisions, confusion columns and denominators from saved synthetic scores/count summaries. It recomputes 648 metric intervals across 81 scenarios. Those summaries do not retain every original bin count; rebuilding the confusion table is not an independent re-fit of all 81,000 pulses. The new explicit-bin worked example supplies a separate likelihood check.

Monte Carlo intervals

For k errors out of a fixed n relevant trials, the Wilson score interval uses $z=\Phi^{-1}(1-\alpha/2)$ and $p_{\text{hat}}=k/n$:

$$\text{Wilson} = \frac{\hat{p} + z^2/(2n) \pm z \sqrt{\hat{p}(1-\hat{p})/n + z^2/(4n^2)}}{1 + z^2/n}. \quad (\text{S18})$$

The Clopper–Pearson endpoints are

$$p_L = B_{k,n-k+1}^{-1}(\alpha/2), \quad p_U = B_{k+1,n-k}^{-1}(1-\alpha/2), \quad (\text{S19})$$

with $p_L=0$ at $k=0$ and $p_U=1$ at $k=n$. This is an individual binomial interval; no simultaneous family coverage is claimed for displaying many such intervals. The intervals describe repeated synthetic trials under this specified generator, not uncertainty in real cloud populations or correctness of the source model. The review recomputes both kinds and compares them with the published values.

Numerical accuracy and the browser cross-reference

The review compares SciPy tails with a positive Poisson-PMF recurrence, retaining a decreasing geometric bound for truncated terms. Small-count on/off tails are also summed directly. No normal/Gaussian approximation replaces the count probabilities.

“Exact Poisson” identifies the discrete statistical model and tail construction; ordinary floating-point code is not formal interval arithmetic. Near a mean of 10^6 and a roughly 5σ upper tail, the retained independent high-precision audit found SciPy's incomplete-gamma survival value differed by about 4.6 parts per million from direct summation. The tested integer thresholds remain the same. The notebook reports this reference-precision qualification rather than hiding it.

The copied reference/statistics/browser_math.js provides the live web-engine functions sf, critical, threshold and upper. The source map gives their line numbers and hash. If Node is available, an optional check actually executes selected browser thresholds and compares them with Python. If Node is unavailable, the result is **NOT_RUN**, not a JavaScript validation claim. Node is not required for the Python notebook.

The previous tiny-background cancellation defect was corrected by computing $P(N \geq 1) = -\exp(-b)$, and zero-count mass/CDF as $\exp(-b)$. The reference copied into this package is the corrected version. The browser has an explicitly bounded live domain; use the offline code beyond its stated limits.

Known-hypothesis oracle: what the ideal bounds mean

For independent Poisson bins with two fully specified mean vectors μ and ν , the Bhattacharyya affinity is

$$BC(P_\mu, P_\nu) = \prod_j \sum_{n=0}^{\infty} \sqrt{p_{\mu_j}(n)p_{\nu_j}(n)} = \exp\left[-\frac{1}{2} \sum_j (\sqrt{\mu_j} - \sqrt{\nu_j})^2\right]. \quad (S20)$$

For one bin, the sum is the exponential series for $\exp(\sqrt{\mu\nu})$, multiplied by $\exp[-(\mu + \nu)/2]$. Independence multiplies affinities. No Gaussian approximation is involved.

Under equal prior probabilities and zero-one classification loss, the optimal binary error satisfies

$$\frac{1 - \sqrt{1 - BC^2}}{2} \leq P_e^* = \frac{1 - TV(P, Q)}{2} \leq \frac{BC}{2}. \quad (S21)$$

The upper bound follows from $\min(P, Q) \leq \sqrt{PQ}$. For the lower bound, write $|P - Q| = |\sqrt{P} - \sqrt{Q}|(\sqrt{P} + \sqrt{Q})$ and apply Cauchy-Schwarz; then $TV \leq \sqrt{1 - BC^2}$.

The review independently enumerates three small two-bin joint PMFs and checks their affinity and exact finite-grid Bayes errors against these expressions. The finite grids contain probability within $1e-13$ of one. **The original 320-row oracle grid is retained but not independently rebuilt by this light notebook.**

These bounds concern known simple hypotheses, equal priors and a specified loss. Substituting best-fit unknown amplitudes/onsets/backgrounds does not retain them. They also do not rank a practical accepted-only error fraction without accounting for rejected events and its different decision problem.

Statistical bindings in the FOV and auroral tables

The downstream audit recomputes critical integers and source means for 95% detection for all 216 FOV rows and 192 auroral rows. It caches repeated background/false-alarm pairs, while comparing every stored row.

Auroral rows state their known background and false-alarm alpha directly. The FOV table's variable called alpha is the **hypothetical population slope**, not the false-alarm allocation. Its retained source uses

count_requirements.ALPHA_OPERATIONAL=0.1/(50,000×86,400×3) and B=(sky count rate+50 dark counts/s)×window. These constants are read through the Python AST; the retained source is not executed.

This checks only the statistical-column binding to the declared background, allocation and power. It does not independently validate FOV geometry, sky radiance, luminosity law, orbital coverage, real counter response or mission event rate. Source hashes and the maximum numerical differences are recorded in the check output.

What the verification does not establish

There is no analog post-counting low-pass in the Poisson shape generator. A deterministic mean RC illustration elsewhere does not make filtered samples independent Poisson variables. Deadtime occupancy flags do not restore missing arrivals or Poisson likelihoods.

A 10 kS/s stream has 100 μs bins; 50 kS/s gives 20 μs; 100 kS/s gives 10 μs. These are sampling/integration rates, not analog noise bandwidths. The input shapes refer to an expected received count rate after transport, not intrinsic nitrogen excitation histories.

No absolute source luminosity distribution, physical-depth labels, end-of-life hardware response, thermal alternative ratio, polar background climatology, operational false-event rate or mission useful-event rate is determined by these checks.

Source and function crosswalk

Equation IDs S01–S02 map to photon_planes and common.py constants; S03–S12 to critical_known, required_signal_known, onoff_pvalue, source_upper_known and source_upper_off; S13–S16 to primitive, q_integrals, poisson_loglik, nuisance_profile and worked_example; S17–S19 to confusion_checks and mc_intervals; S20–S21 to affinity_bound_values and oracle_affinity_controls.

checks/statistics_results.json records exact source hashes and function lines for retained originals. In the original scripts:

Original file	Function	Original line
practical_morphology.py	fit_one	See generated source map
practical_morphology.py	draw_case	138
practical_morphology.py	analyze_case	154
count_inference/count_inference.py	count_threshold_known_background	78
count_inference/count_inference.py	exact_onoff_pvalue	96
count_inference/count_inference.py	poisson_total_mean_upper/lower	133 / 140
count_inference/count_inference.py	red_source_upper_known_background	147
count_inference/count_inference.py	red_source_upper_finite_baseline	156

Line numbers are for the retained versions; the recorded hashes and generated map take precedence if formatting changes. Retaining old code is provenance, not independent theory validation.

Primary implementation documentation: [SciPy Poisson](#), [SciPy binomial](#), [SciPy chi-square](#), and [NumPy Poisson generator](#). The mathematical likelihood and coverage arguments are derived above, not inferred from software documentation.