Pittsburgh in a Grain of Sand

Lessons for Mars rovers and autonomous self-driving cars

CMU Skyline Monitor — `pandr` / `akdeniz`*

2026-07-28 (v11)

Abstract

A grain on a Porter Hall window at Carnegie Mellon — consistent with 0.25 millimetres of debris, three centimetres before a sky monitor's $f/4$ pupil — prints a 76-pixel shadow that is a single-particle optics experiment in disguise. Shadow depth sits at the projected-area value, the extinction paradox's forward-scattered half recollected rather than registered; recollected blue scatter approximately predicts the green channel; and the shadow resolves millimetre-scale sash flexure (uncalibrated). The geometry reproduces the Mars rovers' dust-shadow optics; to our knowledge no prior work continuously tracks the apparent image position of a single adherent near-lens occluder. An autonomous, offline-validated tracker erases the artifact (median dip 4.5 to 0.5 per cent, retrospective; live at proofsandreasons.io), motivating physics-first handling of self-driving cars' soiled cameras. To provide information to the wider scientific community, raw data, materials, methods, and analysis code are available at github.com/simon-dedeo/cmu-skyline-monitor.

Figure 1: The same three raw brackets fused without (left) and with (right) the correction — 2026-07-25 19:17 EDT. The blemish is the dark smudge in the red square at the top of the sky, left of centre; displayed dip 5.46 % → −0.5 % (nulled to within measurement noise). The corrected fusion is the production chain's output.

Contents

*Operated by Simon DeDeo (Carnegie Mellon University). Analysis and code developed with AI assistance (Claude, Anthropic); the human operator is responsible for all claims, citations and errors.

1 The instrument and the defect

A 4 K camera (3840×2160) images the CMU skyline through a fixed window. Every five minutes it captures a three-exposure bracket, fused by Mertens exposure fusion [1] and finished with a local-contrast pass (CLAHE [2, 3]) plus global saturation and contrast curves. Brackets are archived losslessly, so all photometry runs on raw frames; the correction is applied to the raw brackets *before* fusion, because the enhancement stage amplifies any residual dip roughly 2.9×, scene-dependently. The camera (Elgato Facecam 4K) is an $f/4.0$ fixed-focus design — Elgato Prime Lens, 21 mm equivalent on a 1/1.8" sensor (actual focal length ≈ 4.3 mm, entrance pupil ≈ 1.1 mm), focus range 30 cm to 120 cm — mounted with its lens under 1 cm from the pane. These numbers matter: the blemish is imaged as an out-of-focus pupil shadow, and every inference about its physical nature runs through them (§9). The blemish currently projects to $\sim (1300, 132)$ — only 130 px below the top edge of the sensor, which constrains the correction geometry (§5).

What is deployed and what is offline. This note describes code in three states, and labels each claim accordingly: *deployed* (running on the capture host `pandr` or the compute host `akdeniz` as of 2026-07-28), *offline* (analyses run to produce the results here, with code retained under `spot_paper/code/`), and *proposed*. The authoritative component-by-component table, with byte-copies of the deployed files, is `PRODUCTION_STATUS.md` and `production/` in the project directory. An earlier draft did not make these distinctions and correctly drew a reviewer's objection.

2 Statistical model and inference

2.1 Photometric model

Per colour channel, in linear light, the blemish maps the background radiance it occludes as an affine transformation — the standard model for a defocused contaminant layer, which both attenuates the scene and scatters stray light toward the sensor [4] —

$$L_{sp} = t\, L_{bg} + s, \qquad \text{so the fractional dip} \qquad d = 1 - L_{sp}/L_{bg} = (1-t) - s\, L_{bg}^{-1}. \tag{1}$$

The intercept of d against L_{bg}^{-1} identifies the transmittance t; the slope is exactly $-s$, so a nonzero slope is a direct test for an additive component (reported in §5). One structural fact shapes everything: *an exposure bracket cannot separate t from s.* Changing exposure scales L_{sp} and L_{bg} together, so within a bracket the three points are collinear with slope $t + s/L_{bg}$. Brackets calibrate the tone curve, nothing more; separating t from s requires natural variation of L_{bg} at fixed window irradiance (cloud behind the blemish, the twilight ramp).

2.2 Locating the blemish: spatial models (offline analysis)

Equation (1) uses scalar aperture summaries and carries no spatial information. Two spatial estimators exist, and it matters which is which. The *production* nightly estimator (deployed, `spotnight.py`) is deliberately simple: median-stacked ratio maps, position pinned across brackets, a half-depth centroid, and a recursive prior with step and spread gates. The *Bayesian pixel model below is the offline analysis* that validates it: fitted independently, the two agree along the whole eight-day track to ~ 3 px, which is the justification for running the cheap estimator nightly. In the offline model, a level-normalised image window is a smooth background minus a blob,

$$y_i \sim \text{Student-}t\Big(\nu,\ b_{0,m} + \mathbf{x}_i^\top \boldsymbol{\beta}_m - A_{k(m)}\, e^{-Q_i/2},\ \sigma_m\Big)\,, \tag{2}$$

with m indexing image stacks, $k(m)$ the exposure bracket of stack m (one amplitude per bracket class), $\mathbf{x}_i$ the quadratic basis (u, v, u^2, uv, v^2) standardised per stack, and

$$Q_i = \frac{1}{1-\rho^2}\Big[\tfrac{(u_i-c_x)^2}{\sigma_x^2} - \tfrac{2\rho(u_i-c_x)(v_i-c_y)}{\sigma_x\sigma_y} + \tfrac{(v_i-c_y)^2}{\sigma_y^2}\Big],$$

with priors $c_x, c_y \sim \mathcal{N}(0, 30)$ px about the drift-predicted centre and a non-centred random walk of scale τ px/night tying days [5, 6]; $\sigma_{x,y} \sim \mathcal{N}(33, 7)$ bounded to $[15, 60]$ px; $\rho \sim \mathcal{N}(0, 0.35)$ on $[-0.85, 0.85]$; $A_k \sim \mathcal{N}(0.02, 0.012)$ bounded to $[0.002, 0.10]$ (a positive floor, see below); Student-t errors ($\nu \sim$ Gamma, $\nu \geq 2$) keep cloud edges from dragging the centre [7]. Exact code, bounds, seeds and the multi-start command lines are in `spot_paper/code/`. Three structural choices carry the statistical weight:

- **One position, shared across brackets.** The three brackets of a tick see the same glass at the same instant. Pinning the centre across brackets (after normalising each bracket's dip by its own peak, so shape is compared rather than depth) reduced per-bracket disagreement from $\sim$60 px to a few px.
- **Positions tied across days by a drift-informed random walk.** Coordinates are expressed relative to each day's predicted centre, so the parameters are small deviations, $c^{(d)} = c^{(d-1)} + \tau\,\eta_d$, $\eta_d \sim \mathcal{N}(0, 1)$, with τ a few px: the geometry changes smoothly, it does not teleport.
- **Depth free per exposure; shape shared.** Apparent depth is tone-curve dependent; the blob's location and width are physical. A depth that scales with exposure at a fixed location is a signature a smooth background cannot mimic, and it is what identifies the blob against the background polynomial.

Identifiability requires care rather than luck [8]: the background basis must be standardised per stack, the blob width bounded well below the window half-width (a blob as wide as its window *is* a background quadratic), and the amplitude given a positive floor (at $A = 0$ the position has no likelihood gradient and reverts to its prior). With those constraints, MAP estimation (CmdStan [9] L-BFGS [10], multiple starts, best objective kept) recovers the position in seconds. The output is a *MAP state*, not a posterior — the posterior is never sampled — and the trust decision is made by an acceptance gate (§7) rather than by a credible interval. The fitted half-depth radius, 37.8 px, independently matches the 38.2 px measured by a completely separate empirical estimator.

2.3 A recursive filter with honest epochs

For unattended operation the nightly fit is a filter in the classical predict–update sense [11, 12]: tonight's prior is last night's accepted state (a point estimate with gates, not a posterior)

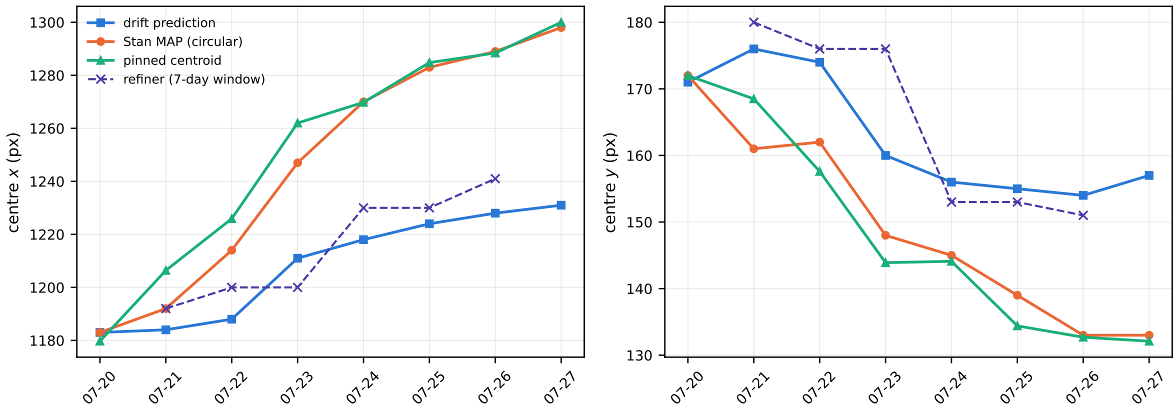

Figure 2: Estimates of the blemish centre across eight days. The scene-drift prediction systematically lags: blemish motion is *not* proportional to scene motion (§5). The Bayesian MAP track and the independent pinned-centroid track agree to ~ 3 px at both ends while sharing no machinery.

propagated by the blemish's own observed velocity (self-calibrating — no parallax constant, no scene-drift proxy), with implausible steps and low-confidence fits rejected in favour of carrying the previous state forward.

One discipline deserves emphasis because everything downstream depends on it: **an estimate carries its data's epoch, not its run date**. The nightly position is fitted from the two previous days of frames, so its effective epoch is the mean timestamp of those frames (≈ 1.5 days before the run). The estimator stores that epoch in its state, velocity is the centre step divided by the *elapsed time between epochs* (px/day — robust to skipped nights, gaps and replays), and the publisher stamps the epoch into the model bundle so all placement extrapolation is referenced to it. Mislabelling the epoch by a day shifts every placement by $v\,\Delta t$ and prints a bright/dark dipole at the blemish — which is how the requirement was learned, twice: first as a run-date-vs-data-date confusion, then as a reviewer's observation that the state carried no epoch at all.

3 The empirical transmittance map

The blemish's radial profile is not Gaussian: its wings fall faster than a Gaussian matched at half depth (Figure 3a). The parametric blob is therefore used only to *locate*; depth and shape come from an empirical map,

$$T_c(\mathbf{r}) \;=\; \mathrm{median}_j\, \frac{I_{j,c}(\mathbf{r})}{B_{j,c}(\mathbf{r})}, \tag{3}$$

the per-channel median, over selected frames, of each frame divided by its own quadratic background, in linear light — in effect a defect-local flat field [13], built by registered, quality-selected stacking in the same spirit as astronomical lucky imaging [14]. Moving cloud averages out; fixed attenuation survives. Four details matter:

- **Registration.** Each frame's window is extracted at that frame's velocity-interpolated centre, so the dip stacks aligned instead of being smeared by its motion across the stacking window.
- **Blemish-local, sigma-clipped background rings.** The background is fitted on an annulus centred on the *blemish* (never the window centre), with MAD-based outlier rejection; otherwise the dip leaks into its own background and the map records roughly half the true depth.

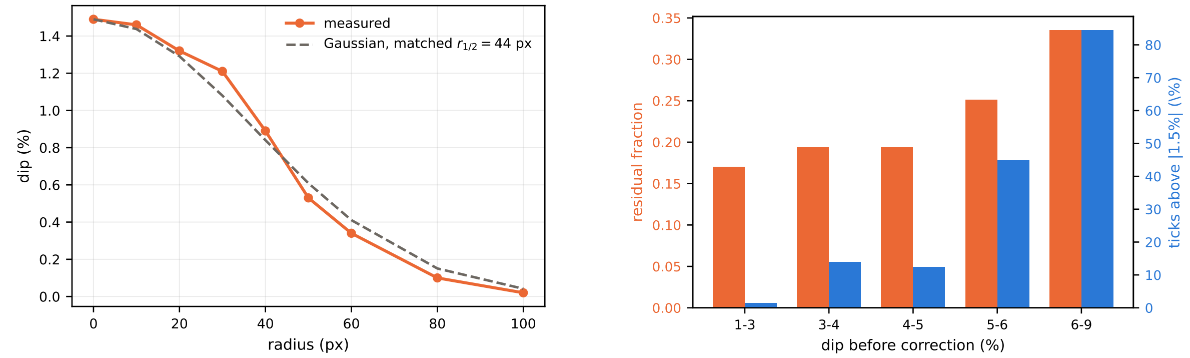

(a) Measured radial profile vs. a Gaussian matched at half depth: the wings fall faster. Parametric fits locate; the empirical map corrects.

(b) Why offline reprocessing solves its gain in display space: with a fixed linear-space gain, the residual is proportional to the input dip.

Figure 3: Shape and gain measurements that fixed the design.

- **Frame selection: bright first, then flat.** The dip is multiplicative, so signal scales with scene brightness. Selecting on flatness alone picks uniform overcast — the flattest sky there is, and exactly where the dip sits at the noise floor.
- **Edge-aware geometry.** The window is clamped to the frame — the blemish sits closer to the top edge than the window half-size — and the blemish's local coordinates inside the window, plus the measured centroid of the map's dip content, are published with the map. The corrector aligns the *content*, never the window origin.

The publisher refuses to publish rather than publish junk: minimum frame counts per channel and a sanity range on the map's peak depth. On any failure the live system keeps the previous good map.

4 Applying the correction

4.1 Position: three layers, fastest wins

The blemish migrates intraday, so placement cannot rely on the nightly fit alone. At every tick the corrector chains three estimators, each bounded by the one below it:

1. **Nightly posterior** (position, velocity, epoch) — the slow, validated anchor.
2. **Velocity extrapolation** from the fit epoch to now (fractional days), total displacement capped at 40 px.
3. **Persistent template lock** — the measurement update. The corrector cross-correlates the frame's own dip (frame divided by its local background) against the map's dip template within ±25 px [15]; a decisive lock (normalised correlation ≥ 0.45, unambiguous runner-up margin, bounded displacement) snaps the map onto the measured position *and is saved*; ring cleanliness gates the amplitude self-calibration, not the lock. A lock must also be unambiguous: the runner-up correlation peak (outside the winner's neighbourhood) must trail by a margin, so a cloud edge or contrail that ties the blemish cannot capture the tracker. The next tick starts from the last saved lock (if fresh and within 40 px of the prediction) rather than from the prediction alone — a tick-cadence recursive filter whose prior is minutes old instead of a day old. The nightly prediction remains the fallback and the sanity bound, so one corrupted lock cannot walk the correction away. The published

bundle itself is integrity-checked on load (schema version, array hash, finiteness); a torn or corrupt publish is rejected and the corrector falls back to its self-referential ring fit for that tick (the last accepted bundle is not cached in memory). Every tick appends its provenance — model timestamp, map hash, code version, enable-flag state, lock diagnostics — to `spot_provenance.csv`.

The map content is placed by warping (bilinear, border transmittance 1.0) with the window origin fixed at its published, frame-clamped location — necessary because the blemish sits too close to the frame edge for the origin itself to follow it.

4.2 Amplitude: exposure scaling plus gated self-calibration

The map records depth at the reference exposure. Two mechanisms adapt it per frame:

- **Exposure scaling.** A measured depth-vs-level curve (fitted from within-tick bracket triples) scales the map's deficit by $s(L)/s_{\mathrm{ref}}$, where L is the frame's local background level.
- **Per-tick self-calibration, gated.** The tone curve is uncalibrated, so on dark backgrounds the true dip runs deeper than the scaled map. When the surrounding ring is clean enough to trust (relative MAD < 0.05), the corrector measures the frame's own core deficit and trims the amplitude, clamped to $[0.75, 1.35]\times$ the model. When the ring is cloudy the gate closes and the model amplitude stands — chasing a cloud-inflated "dip" is how inverse bright spots are made.

The correction itself divides rather than reconstructs: $I \mapsto I/\tilde{T}$ with $\tilde{T}$ feathered to unity outside the footprint — possible precisely because a defocused occluder's artifact is low-frequency and multiplicative, so scene radiance is recoverable pointwise [4]. There is deliberately *no scene-referenced cap*: a clamp referenced to the observed image or a local sky fit eats the correction exactly where a bright cloud passes behind the blemish (the dimmed cloud becomes the reference), and an earlier pixelwise version also silently deleted the upper half of the noise distribution, leaving an unnaturally smooth disc. Safety is instead *model-referenced*: the correction's shape is the quality-gated empirical map, and its amplitude is bounded relative to the characterized deficit — the closed-loop gain is clamped to $[0.4, 1.8]$, the live amplitude trim to $[0.75, 1.35]\times$ the model with the exposure scale capped at 2.5, and the total applied deficit floored at $T \geq 0.85$. The correction can therefore never exceed a fixed multiple of what the blemish could physically account for, whatever the scene does; clouds neither evade the bound nor get eaten by it, and the grain's noise is never explicitly filtered; the multiplicative correction rescales local noise amplitude together with the signal, by design.

4.3 Offline reprocessing: closed-loop gain

For archive reprocessing, where four fusions per tick are affordable, the gain is solved in *display* space: fuse at adaptive probe gains bracketing the root, secant-step to the gain that nulls the *fused* dip, and refine once. Because the gain is chosen by evaluating the corrected residual on the frame itself, the archive numbers it produces are an oracle-style upper bound, reported separately from the live path. This removes all dependence on the scene-dependent enhancement amplification (Figure 3b); the fitted gains (median 2.83) independently recover the measured $\approx 2.85\times$ CLAHE amplification.

5 Physics: what the data established

The image's motion is two-process — but the object need not move. Step-level regression of daily blemish-image motion on scene drift (phase correlation [16] on the building band) gives $\Delta x_{\text{image}} \approx 0.93\,\Delta x_{\text{scene}} + \epsilon$ with med $|\epsilon| \approx 11$ px/day: camera rotation passes through at unit gain, plus an independent component — episodic, up to 27 px/day, continuous within days (7 px to 18 px, measured by splitting days into thirds). With the true shadow geometry (§9) the independent component requires no motion of the object at all: at ∼3 cm standoff the translational lever is $f_{\text{px}}/d \approx 63$ px/mm to 73 px/mm on axis (focal length $f_{\text{px}} \approx 2200$ px from the advertised 90° diagonal field of view; $d = 30$ mm to 35 mm), rising toward ∼75 px/mm to 87 px/mm at the blemish's off-axis position. The week's 120 px walk is then 1.4 mm to 1.9 mm of cumulative flexure and the intraday motion 0.08 mm to 0.3 mm — unremarkable for a sun-struck sash with a camera resting against it. These are uncalibrated estimates: the local plate scale at the blemish has not been measured, and a controlled sub-millimetre camera translation (§11) is the direct calibration. (An earlier draft quoted 92 px/mm from an average-pixels-per-degree shortcut that overstates the paraxial scale; a referee's recomputation prompted this correction.) Day-by-day step ratios ranging from 1.06 to 26.7 simply reflect days when the mount flexed without rotating and vice versa. No constant ratio, and hence no feed-forward scheme, can track this; a near-field fiducial (a tape dot on the glass) would separate camera flexure from any residual object motion if the distinction ever matters. Operationally it does not: the corrector tracks the apparent position either way.

A response-curve-confounded additive residual. The slope of equation (1) is significantly nonzero in all channels and stable across nights (blue $\approx +0.07$ at $p \sim 10^{-9}$; green $+0.022$; red $+0.010$), strongly chromatic (B>G>R by $\approx 6.7\times$, a Rayleigh-like ordering). Its sign implies a fixed *dark* offset — occlusion-like — rather than added scattered light. Three caveats mean this is exploratory, not established: the photometry aperture was fixed at the original seed and drifted off the blemish for part of the record (it tracks the published centre as of 2026-07-28, so future nights measure the feature itself); the regression treats autocorrelated frames as independent; the linearisation is assumed inverse-sRGB, and a tone-curve error generates exactly this signature; and one early night shows the opposite sign. Until the response curve is calibrated per channel, the term should be read as a response-curve-confounded residual rather than physics. The correction is deliberately robust to the ambiguity: per-frame self-calibration corrects whatever depth each exposure actually exhibits, without needing to know why.

Frame-edge geometry. The blemish sits ∼130 px below the frame top while its background annulus extends to 170 px: the ring is partially clipped and the map window must be clamped (nothing can be cropped larger — the frame is the full sensor readout). Because the blemish is near-field, translating the camera up a few millimetres would move it ∼100 px down in frame while leaving the distant framing essentially untouched; that remains the physical fix for edge margin, with the caveat that migration continues regardless, so the tracker stays load-bearing.

6 System architecture: where each computation runs

The system is two-tier by design (Figure 4). Everything statistical and expensive runs *once per night* on a 32-core compute host (`akdeniz`); everything that runs *per tick* is closed-form, operates on the published 300×300 window rather than the 4 K frame, and lives on the capture machine (`pandr`, an M1 MacBook Air) inside the corrector itself. Measured on real frames, the entire per-tick correction — placement, template lock, amplitude calibration, divide — costs 84 ms per bracket, about 0.25 s of a 15 s to 28 s capture cycle.

Two points about the split deserve emphasis. First, **the template lock is computed on the Air, not on `akdeniz`**: it is part of the corrector, executed per bracket on the local window — two small Gaussian blurs and one 120×120-in-170×170 normalised cross-correlation, 5 ms to 10 ms. The compute host never sees a live frame; its entire influence on the live system is

level	where	cadence	cost
nightly estimator (pinned centroid, recursive prior)	`akdeniz`	nightly	seconds
MAP position/shape fit (multi-start L-BFGS; offline validation)	`akdeniz`	on demand	seconds
registered transmittance map (∼150 frames)	`akdeniz`	nightly	∼2 min
acceptance validation (20 ticks × 4 fusions)	`akdeniz`	nightly	∼4 min
publish (map + model, ∼1.1 MB)	`akdeniz`→`pandr`	nightly	seconds
per-tick correction (place, lock, calibrate, divide)	`pandr`	every tick	84 ms/bracket
archive reprocessing (closed-loop gain)	`akdeniz`	on demand	∼7 s/tick, 26× parallel

Table 1: Levels of computation. Only the nightly artifact crosses the network; the live path has no per-tick dependency on the compute host.

the nightly `spot_model.json` + `spot_tmap.npy` pair. After the MAP fit, nothing per-tick is an optimisation — a warp, a correlation, two 6-parameter least-squares fits, an elementwise divide.

Second, the split is about *availability* as much as cost: the corrector sits in the serving path, so a per-tick network dependency would turn any compute-host or network outage into a real-time degradation of the public site. As built, `akdeniz` can disappear for days: velocity extrapolation and the persistent lock carry the placement, staleness tiers step down gracefully, and the acceptance gate catches any quality regression when the nightly service resumes.

7 Autonomous operation

A nightly job on `akdeniz` runs the whole cycle unattended: sync frames from the mirror; select frames; fit the position (recursive prior, pinned across brackets); build the registered maps; validate; publish to the capture host only if validation passes. Validation runs on clean-sky ticks — level and ring-structure thresholds measured *outside* the footprint, so the criterion cannot be circular — and the gate requires zero inverse spots, a floor on the median removal fraction, a bounded tail, and no regression against a rolling baseline of recent accepted nights. Failure paths are first-class: insufficient clear frames, low-confidence positions, and implausible jumps all carry the previous parameters forward with a log line, and repeated failures escalate. The corrector applies tiered staleness policies (beyond seven days it stands down to a gated fallback). Success is asserted from output artifacts, never process exit codes. Historical reprocessing uses the identical code path (`--replay`), which is also the honest test of autonomy: replaying the archive exercises cold start, carry-forward, gate rejections, and recovery exactly as a live outage would.

8 Performance

Paired evidence is retained in machine-readable form (`spot_paper/evidence/live_paired.csv`: per tick, the archived served frame vs. a re-fusion of its own raw brackets uncorrected, with positions, levels, ring structure and model version). A referee correctly observes that these rows are *retrospective*: the analysis script stamps each row with the model current at analysis time, not the bundle that served the tick, and archive frames may have been reprocessed — so the comparison establishes paired-reprocessing performance, not contemporaneous live performance. Per-tick provenance (model hash, code version, enable state, lock diagnostics; `spot_provenance.csv`) now accumulates append-only on the capture host, and a prospective frozen-model trial over complete clear, mixed and cloudy days is the planned test. On the same clean-tick selection as Table 2: the clear evening of 07-27 gives $n = 20$, median removal 92.9 %, median served residual 0.35 %, range $[-0.80, +1.44]$; 07-28 — a heavily clouded day during which

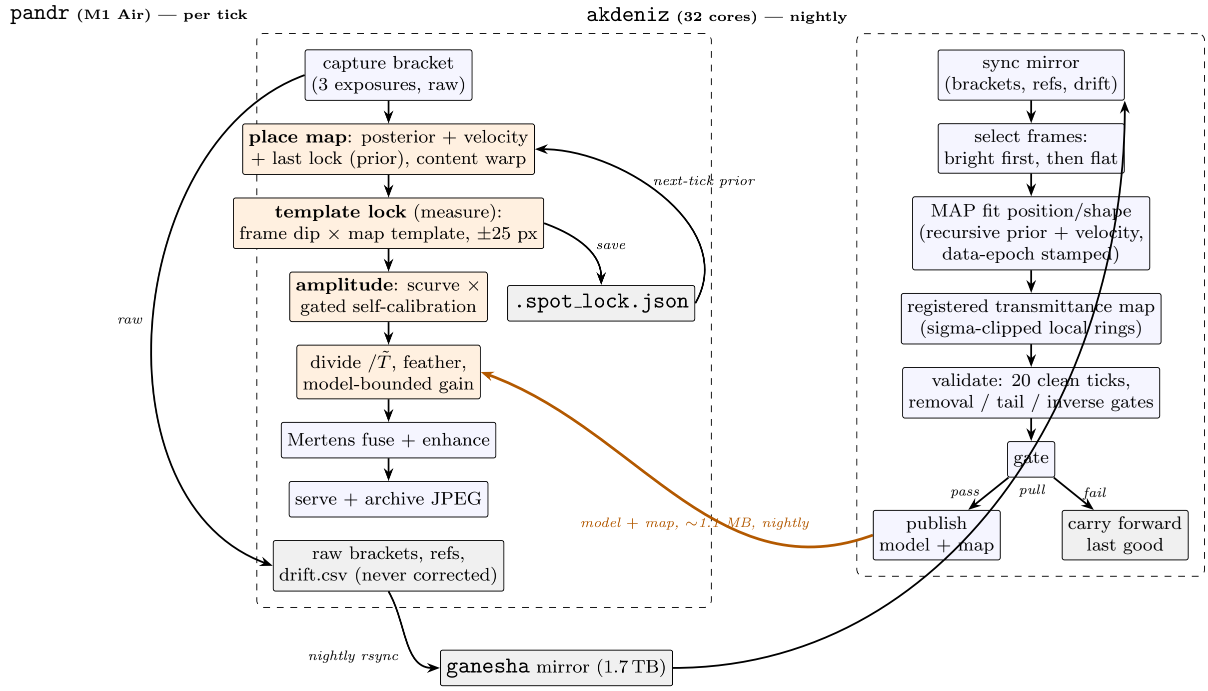

Figure 4: Data flow. Orange nodes run inside the corrector on the capture machine at every tick; the template lock and its persistence loop are local to the Air. The compute host's only influence on the live system is the nightly published model+map (orange arrow); its only input is the raw-data mirror. Raw brackets and reference frames bypass the corrector entirely, so the measurement stack is never contaminated by its own correction.

the tracking chain was deployed and revised mid-day — gives $n = 13$ post-deployment clean ticks with median removal 69.7 % and median residual 1.70 %. The honest reading is that clear-sky live performance matches the offline table and cloudy-sky "clean" ticks remain contaminated by weather at the measurement aperture; a full clear-day out-of-sample evaluation is pending and the evidence file will accumulate it. Centring is monitored separately by a dipole statistic — the signed difference between the dip in the leading and trailing half-disks along the motion direction; a mis-centred correction is *stationary* and oriented, where cloud moves. After the epoch and measurement-update corrections this statistic fell from +2.5% ($n = 17$ frames) to +0.4% ($n = 2$), with the remainder symmetric weather signal.

9 What is the blemish? A physical identification

The answer required one datum that no amount of image analysis could supply: the camera's actual aperture. Two successive wrong identifications — a migrating condensation droplet, then a reflection of a room object — were each internally consistent with an *assumed* webcam aperture of $f/2$–$f/2.4$. The published specification for this camera is $f/4.0$, fixed focus at 30 cm to 120 cm: an entrance pupil of ≈ 1.1 mm, not 2.7 mm to 3.9 mm. With that single correction, every measurement converges on the mundane truth, which is also what direct inspection of the pane suggests: **a** ~ 0.25 mm **speck of debris on the glass**.

The solve. A pixel viewing the distant scene collects a ray bundle that crosses the pane as a patch of diameter $\approx a\,(1 - d/F)$, where d is the pupil-to-pane distance and $F \approx 0.5$ m the fixed focus distance — the circle-of-confusion geometry of any defocused imaging system [17, 18]. A speck of size s removes the fraction $(s/a)^2$ of a pixel's light and images with a support width $\approx \big(s + a(1 - d/F)\big)/d$ (a support scale, not automatically the measured FWHM; the

(a) before (b) after (c) difference, 6×

Figure 5: Pixel-scale detail of Figure 1 — the same three raw brackets through the same production chain (registered map, template lock, adaptive closed-loop gain, model-referenced amplitude bounds): 5.46 % → −0.5 %. The difference (right, 6×) is a single compact blob; nothing outside the footprint is touched, and the noise texture inside it is preserved.

stratum (offline, closed-loop gain)	n	before	after	removed	$\lvert after \rvert > 1.5\%$
clean daylight frames	495	4.46 %	0.52 %	88.1 %	8.3 %
+ ring-structure gate	301	—	0.51 %	—	7.0 %

Table 2: Offline reprocessing statistics from `spot_fit_check/sweep2.csv`. Selection, disclosed: displayed frame level ≥ 140 and uncorrected measured dip in [1%, 9%] (row 2 adds ring-structure < 3). The dip window *selects on a measured quantity*: it estimates performance on frames where a blemish-scale dip is detectable, and excludes negative, weak and cloud-inflated cases — it is not a generic daylight population. Residual distribution: 66 of 495 frames (13%) end negative, six below −0.5%, one below −1% (minimum −1.28%); "inverse bright spot" in the text means a residual below −1%. The 495 frames span 8 days (11 to 111 per day); per-day median removal runs 83.2 % to 94.3 %, and a day-block bootstrap gives removal 88.1 % with 95% CI [85.6 %, 90.9 %] and median residual 0.52 % with CI [0.40, 0.66].

difference is absorbed into the fitted range of d, and the pixel-bundle factor $1 - d/F$ assumes the nominal 0.5 m focus conjugate of a spec-sheet focus *range*). The measured depth (3 % to 7 %) gives $s \approx 0.22a \approx 0.24$ mm; the measured 76 px FWHM then requires $d \approx 30$ mm to 35 mm. (That a thin occluder far from focus prints as a low-frequency attenuation pattern — the occluder convolved with the aperture — is the standard image-formation model for near-lens contaminants [4, 19–21]; multi-frame parallax approaches to removing such occluders exist when the camera moves [22], which ours, deliberately, does not.) Of that, at most ∼ 13 mm is air gap and glass, implying the entrance pupil sits ∼ 2 cm behind the front of the lens — deep, but characteristic of short-focal-length inverted-telephoto designs, and a checkable prediction. The faster-than-Gaussian wings (Figure 3a) are the compact-support signature of exactly this: a hard-edged occluder convolved with a hard-edged pupil. The neutral colour (integrated deficits equal across channels to 8 %; the red channel deeper but narrower — a point-spread difference, not pigment, traced below to wavelength-dependent forward scattering) fits nondescript grey debris. The two or three fainter compact features elsewhere on the pane are siblings: other motes.

Why the image moves. The same geometry that blurs the speck amplifies motion: image shift equals camera–pane translation times the local plate scale over d — 63 px/mm **to** 87 px/mm, uncalibrated (§5). The feature's entire week of "migration" (120 px) is 1.4 mm to 1.9 mm of accumulated flexure; its intraday drift is 0.08 mm to 0.3 mm — the thermal breathing of a sun-struck sash with a camera resting against it. No motion on the glass is required, although only a fixed fiducial can prove none occurred. The kinematic anomalies that motivated more exotic identifications (episodic steps, opposite-signed step days, onset with the hot spell) are all natural behaviours of mm-scale structural flexure seen through a 30× lever.

The colour of the blemish. The registered per-channel transmittance maps permit a measurement most defects never get: the speck's extinction spectrum, in three bands. The subtlety is separating the *object's* colour from the *optics'* colour — longitudinal chromatic aberration gives each channel a slightly different defocus blur, and a tighter point-spread function makes the same grey object look deeper and narrower in that channel. The discriminator is the *integrated* deficit, $\int(1-T)\,dA$, which is invariant to blur:

	blue (∼460 nm)	green (∼540 nm)	red (∼610 nm)
core depth	4.49 %	4.89 %	5.38 %
half-depth radius	51.4 px	50.7 px	48.8 px
integrated deficit (rel.)	1.00	1.08	1.07

Red is 20% deeper than blue but also narrower, and the integrated deficits agree to within 8% across all three channels. The depth/width pattern is *not* lens chromatism — for this geometry every axial mechanism (normal dispersion, achromat secondary spectrum, edge diffraction at Fresnel number ∼1 (radius convention), sensor carrier diffusion) predicts the red shadow *broader*, the wrong sign. What produces deeper-and-narrower red at constant integral is the speck itself, through a mechanism with a long history in particle optics: *the apparent extinction of a particle depends on how much of its near-forward scattered light the measuring aperture recollects.* A particle large compared to the wavelength ($\pi s/\lambda \sim 10^3$ here) removes *twice* its geometric blocking, half of it deposited in a narrow forward diffraction lobe — the "extinction paradox" [23–25] — and any detector of finite acceptance recovers part of that lobe and registers correspondingly less extinction, a bias quantified for transmissometry by Deepak and Box [26]. Our geometry sets a generous acceptance: light scattered at the pane by less than $\theta_{\rm acc} \approx a/(2d) \approx 18\,$mrad half-angle (34 mrad full cone) re-enters the pupil. The speck's whole-body diffraction lobe ($\theta \sim \lambda/s \approx 2\,$mrad) lies entirely inside that cone — which is why the core reads simple geometric blocking, $(s/a)^2$, in the first place. Scattering from micron-scale surface texture, however, spreads into lobes of width $\theta \sim \lambda/d_{\rm f}$ for feature size $d_{\rm f}$; only the fraction falling inside the acceptance cone, $\propto (\theta_{\rm acc}\, d_{\rm f}/\lambda)^2$, is recollected, and that fraction is $(610/460)^2 \approx 1.8\times$ larger for blue than for red. The recollected blue re-enters the ray bundles — refilling the shadow core (blue shallower), feathering its skirt (blue wider, by $\sim 30\,\mu$m at the pane), and slightly lowering the blue integral, exactly as measured. Red, its scattering spread over the widest lobes, is recollected least and casts the cleanest geometric shadow. The blur-invariant total extinction remains **neutral to** ±8%. These are consumer-camera bands after demosaicing, white balance and colour processing, linearised by an assumed inverse-sRGB curve; the integrals are computed within $r < 90\,$px with negative map values clipped, and carry no formal uncertainty — the channel comparison is indicative, not calibrated. Subject to that caveat, the bound disfavours strongly pigmented candidates: rust (× several in the blue), organic tar or sap, pollen, chlorophyll debris. What remains is the unglamorous family that grey, wavelength-flat extinction describes: mineral dust, paint or plaster fleck, paper fibre, insect frass. In transmission at $f/4$ the speck is, to the limit of what three colour channels can say, *grey* — consistent with the direct inspection finding an ordinary speck on the pane. (The blue-heavy *additive* residual of §5 is a separate, response-curve-confounded term and should not be read as the object's colour.)

Morals. Three. First, *an assumed instrument parameter is a prior like any other, and a strong one*: two plausible physical stories were built downstream of $f/2$, and both collapsed the moment the real $f/4$ was looked up. Second, the apparent motion of a near-field feature is a strain gauge, not a trajectory: at these standoffs the image measures *mechanics* (of mount and sash) more than it measures the object. Third, the correction architecture was right to track the *apparent* position agnostically — every one of the candidate physical stories moved the image the same way, and the corrector never needed to know which was true. The remaining physical action item is pleasingly small: wipe the pane at the predicted spot (inner surface first, then outer — the 3 mm path difference changes the blur width measurably, so before-and-after even identifies which surface it was on), and keep the tracker running for the next mote.

10 Validating and extending the literature

The system was built to clean pictures. Measured carefully enough, it turns out to double as a small test bench for three literatures — particle optics, contaminant-artifact removal, and robust hierarchical estimation. This section separates what our data *confirm*, where they fill (or open) *gaps*, and what they can and cannot say about live *disagreements*.

10.1 Confirmations of theory

The pupil-shadow model, at its own validation point. The geometric model of §9 — artifact = particle silhouette convolved with the entrance pupil's footprint on the glass, depth = area ratio — is exactly the optical model Willson et al. [20] derived for Mars-rover camera dust, and their validation regime (quarter-millimetre particles against a 1.1 mm entrance pupil) matches ours closely — but their quantitative validation compared artifact *diameters*; reference-particle variability prevented a predicted-vs-measured optical-density test. Our depth photometry is therefore complementary evidence in the same geometry rather than a confirmation of their quantitative predictions, with two extensions their achromatic, deposition-only setting did not need: per-channel photometry and continuous position tracking of the apparent artifact position.

Extinction energetics in the geometric limit. The particle-optics prediction is quantitative. From the speck, the pupil subtends a half-angle of ≈ 18 mrad (1.05°) — roughly 7–10$\times$ *wider* than the speck's whole-body forward-diffraction lobe ($\lambda/2s \approx 0.10^\circ$ blue to 0.15° red), yet $\sim 57\times$ *smaller* than the $\sim 60^\circ$ a circular on-axis aperture must subtend before a transmission measurement converges to the true extinction cross section [27, 28]. The apparatus therefore counts essentially all forward-diffracted light as *not removed*, and theory predicts it should register an effective attenuation at geometric blockage, not the asymptotic $Q_{\text{ext}} = 2$. It does: the inspected ~ 0.25 mm speck predicts $(s/a)^2 = 5.2\%$ obstruction; the measured per-channel core depths run 4.5 % to 5.4 % (Table, §9), where doubling would give 10.3%. We do not claim this *excludes* $Q_{\text{ext}} = 2$ — Q_{ext} is a far-field plane-wave cross-section, and the near-field, wide-acceptance conditions here are precisely those in which the paradoxical factor is recollected rather than registered; the measurement instantiates that regime, it does not adjudicate the cross-section. (Markel's $\sim 60^\circ$ convergence criterion likewise concerns a plane-wave integrating-detector configuration; we use it as an order-of-magnitude guide, not an exact correspondence.) Achromaticity corroborates independently: geometric blocking by an opaque silhouette is wavelength-flat, and the integrated deficits agree across channels to 8%, whereas any partially developed factor-of-two would print a strong monotonic trend through the channel-dependent Fresnel number ($N_F = (s/2)^2/\lambda d \approx 0.85$–$1.13$ across RGB — right at the shadow/Fraunhofer boundary, where the paradox asymptotics have not switched on).

The λ^{-2} recollection law, with an out-of-sample check. Inverting the blue–red integrated-deficit split of §9 under the acceptance-cone model gives recollected texture-scatter fractions $\rho_R \approx 8\%$ and $\rho_B \approx 14\%$ of the blocked light — physically plausible for micron-textured debris. The green channel, unused in that inversion, is then *predicted* at a relative integral of 1.045 (no formal uncertainty is propagated through the maps and colour pipeline — a day-level bootstrap is future work); it is measured at 1.08, within the few-percent scatter of the maps. Small, but it is the difference between a story that accommodates the data and one that predicts a number.

The affine contaminant model, per channel. The intercept/slope decomposition of equation (1) — attenuation plus additive scatter, the Gu et al. [4] image-formation model — holds in linear light in all three channels with night-to-night stable t, at single-defect rather than whole-lens-aggregate resolution.

10.2 Gaps: two filled, one opened

The static-contaminant assumption. The artifact-removal literature partitions cleanly: static occluders with physics [4, 19, 20], where the contaminant map is calibrated once and time-invariance is baked into the procedure; dynamic occluders with the wrong physics for our case — adherent raindrops modelled refractively, as little wide-angle lenses rather than attenuators, and without positional tracking [29]; dynamic contamination with no physics at all — automotive soiling handled as segmentation-then-inpainting [30]; and population dynamics without individual tracking — the InSight lander's lens-dust census, whose temporality is deposition and wind removal of blobs, not migration of one [31]. To our knowledge no prior work continuously tracks the *apparent image position of a single adherent occluder* — the shadow migrates across the frame whether or not the grain moves on the glass — inferred through its defocused pupil shadow; the hierarchical state-space estimator of §2 has no antecedent in this literature. Our data show the gap is not academic: intraday migration reaches 24% of the shadow's FWHM (7 px to 18 px of 76 px), and the calibrate-once assumption fails measurably — a placement built from ∼1.5-day-old data printed a ±2.5% bright/dark dipole that per-tick tracking reduced to ±0.4% (§8).

Chromatic signatures of contaminant shadows. Willson et al. [20] declare their model achromatic by construction; the removal and soiling lines contain no spectral analysis; we found no dedicated study of the RGB signature of a near-lens contaminant shadow. Astronomical flat-fielding practice has long known *anecdotally* that dust features are filter-dependent and drift between nights, so we claim the first *systematic per-channel measurement*, not the first awareness — and the blur-invariant integrated deficit of §9 is the methodological piece that makes the measurement meaningful, separating the object's colour from the optics' colour.

A gap opened: the shadow as metrology. The 63 px/mm to 87 px/mm (uncalibrated) lever that made the artifact hard to track also makes it a free strain gauge: sub-millimetre mount and sash flexure, invisible in the distant scene, is resolved at 0.08 mm to 0.2 mm intraday through an accidental fiducial on the glass. We know of no literature that uses contaminant shadows for camera-mount metrology; a deliberate near-field fiducial (§5) would make the technique intentional.

10.3 Controversies: evidence either way, and honest silence

What makes $Q \to 2$: a one-sided dispute. Berg et al. [25] reject *both* canonical explanations of the extinction paradox — van de Hulst's edge diffraction and Brillouin's shadow-forming interference — relocating the cause in the Ewald–Oseen cancellation of the incident wave inside the particle, and calling the result a reworking of sixty years of accepted understanding. We searched for published rebuttals and found none; the visible fault line is narrower — Markel [27] declines to condemn Brillouin while agreeing that edge diffraction is the wrong picture. Our measurement is deliberately silent here: every account predicts the same recollected forward flux, so no fixed-acceptance energetic measurement — ours included — can discriminate the mechanism. We confirm the operational physics and abstain on its interpretation.

Is "extinction" detector-independent? No — and our channels show it. Markel [27] argues extinction is not a detector-independent observable: the extinguished power lives in the incident–scattered interference cross-term, and what an instrument registers depends on its aperture's size, shape and position, degenerating to *absorption* for large collection angles. Our system is a clean datum at that asymptote: at 57× below the aperture-convergence criterion we measure geometric blocking, and the paradoxical factor of two is nowhere in the data. The three colour channels then act as a second knob: each samples a different lobe-width-to-acceptance ratio, and the deficit slides by 7% from blue to red exactly as the aperture-dependent view requires — a three-band, single-particle demonstration that the measured quantity is a property of the *instrument–particle pair*, not the particle alone. On the related beam-geometry continuum of Lai et al. [32] — actual removed power running from 1× geometric (narrow beam, near field) to 2× (broad beam, far field) — our point sits at the 1× end ($N_F \approx 1$, standoff well inside the shadow's Rayleigh range), as measured.

Symmetric Student-t under one-sided contamination. The robustness literature's warning lands squarely on us: symmetric-t likelihoods are robust to symmetric heavy tails, but real contamination is typically asymmetric, and under it the symmetric-t location estimand is only implicitly defined and drags toward the contamination [33, 34]. Cloud structure through a window is exactly such one-sided contamination. Our defence is empirical and structural rather than theoretical: the MAP track agrees with an independent pinned-centroid estimator to $\sim$3 px along eight days that include heavy cloud, and the pooling across stacks, the positive amplitude floor, and the shared shape limit how far any one contaminated stack can drag the centre. The literature-sanctioned sensitivity check — refitting with a skew-t or an explicit one-sided contamination component — is acknowledged future work.

To centre or not to centre. Our day-linking random walk is non-centred; the honest citation status is that Papaspiliopoulos et al. [5] established the choice as data-dependent, Yu and Meng [35] proved that neither parameterization dominates and that interweaving both can beat either — and that those are theorems about *MCMC mixing*, not MAP optimization, so our usage is an analogy (conditioning of the optimization surface when per-day information is weak), not a covered claim. Multi-start convergence to a common optimum is consistent with, but does not prove, the analogy.

Selection bias in stacking. Lucky-imaging practice knows that statistics evaluated at the point used to select or align frames are artificially enhanced — regression to the mean in image form [36]. We confirmed the caveat in the wild, from the other side: selecting map frames on *flatness* (the statistic nearest the measurement) chose uniform overcast and drove the dip to the noise floor, while selecting on *brightness* — exogenous to the dip — recovers it (§3). The general rule our data support: select on a statistic independent of the one you intend to measure.

Code, data and reproducibility

Deployed code is byte-copied under `production/` with a component-status table in `PRODUCTION_STATUS.md`. Offline analysis code (Stan models, drivers, map builders, the intraday and colour measurements) is under `spot_paper/code/`; evidence files are `spot_fit_check/sweep2.csv` (offline table, filters as disclosed) and `spot_paper/evidence/live_paired.csv` (live paired measurements). The synthetic end-to-end test (a frame carrying exactly the published defect must null through the deployed corrector to $< 0.5\%$) is `tests/test_synthetic_end_to_end.py`; the current chain passes at 0.000 %. A fuller test matrix (edge-clamped crops, injected motion, corrupted bundles, template-lock distractors) is specified in the 2026-07-28 review and remains future work; the synthetic test now exits nonzero on failure. The full repository, including a raw-bracket sample sufficient to reproduce Figure 1 and the deployed code as byte-copies, is public at github.com/simon-dedeo/cmu-skyline-monitor; the complete bracket archive (∼hundreds of GB) is available on request. A point-by-point response to the second review is `REFEREE2_ACTIONS.md`.

11 Limitations

- Live post-deployment evidence is one mixed-conditions day; the offline table is the strong result, and live parity is demonstrated only for one clear evening.
- Cloudy-sky residuals at the blemish position are dominated by real cloud; the meaningful acceptance metric is clean-sky removal, and eye tests should distinguish stationary structure (defect) from structure that moves with the weather.
- The acceptance metric shares its centre with the estimator; a common-mode position error would pass unnoticed. The population median of the uncorrected dip is the cross-check — a mis-centred night reads visibly low.
- The nightly state model is position-plus-velocity; intraday migration is handled by the template lock, not modelled. If migration accelerates beyond the ± 25 px lock search radius within a day, tracking degrades gracefully to the prediction, and the gate catches the quality regression the next night.
- The frame-edge geometry worsens at ~ 2 px/day; the camera lift is physical maintenance that software cannot substitute for.
- One-minute golden-hour frames keep no raw brackets, so their era corrections cannot be reprocessed; they are labelled rather than silently mixed with reprocessed frames.
- The plate scale at the blemish is uncalibrated; all flexure figures inherit a $\sim 30\%$ scale uncertainty until a controlled camera-translation calibration is done. The colour-channel comparison is likewise uncalibrated (consumer pipeline, assumed inverse-sRGB); narrow-band or RAW measurements would make it quantitative.
- The physical identification is a consistent solve, not a controlled measurement: known-size calibration dots, a measured entrance pupil, a fixed pane fiducial and the wipe test (§9) are the decisive experiments, none yet performed.
- Paired performance evidence is retrospective (see §8); no false-correction study under contrails, birds, glare or saturation has been run.

References

[1] Tom Mertens, Jan Kautz, and Frank Van Reeth. Exposure fusion: A simple and practical alternative to high dynamic range photography. *Computer Graphics Forum*, 28(1):161–171, 2009. doi: 10.1111/j.1467-8659.2008.01171.x.

[2] Stephen M. Pizer, E. Philip Amburn, John D. Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart M. ter Haar Romeny, John B. Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3):355–368, 1987. doi: 10.1016/S0734-189X(87)80186-X.

[3] Karel Zuiderveld. Contrast limited adaptive histogram equalization. In *Graphics Gems IV*, pages 474–485. Academic Press, 1994. doi: 10.1016/B978-0-12-336156-1.50061-6.

[4] Jinwei Gu, Ravi Ramamoorthi, Peter N. Belhumeur, and Shree K. Nayar. Removing image artifacts due to dirty camera lenses and thin occluders. *ACM Transactions on Graphics*, 28 (5), December 2009. doi: 10.1145/1618452.1618490. Proc. SIGGRAPH Asia 2009.

[5] Omiros Papaspiliopoulos, Gareth O. Roberts, and Martin Sköld. A general framework for the parametrization of hierarchical models. *Statistical Science*, 22(1):59–73, 2007. doi: 10.1214/088342307000000014.

[6] Michael Betancourt and Mark Girolami. Hamiltonian Monte Carlo for hierarchical models. In Satyanshu K. Upadhyay, Umesh Singh, Dipak K. Dey, and Appaia Loganathan, editors, *Current Trends in Bayesian Methodology with Applications*, pages 79–101. Chapman and Hall/CRC, Boca Raton, FL, 2015. doi: 10.1201/b18502-5. Also available as arXiv:1312.0906.

[7] Kenneth L. Lange, Roderick J. A. Little, and Jeremy M. G. Taylor. Robust statistical modeling using the t distribution. *Journal of the American Statistical Association*, 84(408): 881–896, 1989. doi: 10.1080/01621459.1989.10478852.

[8] Andrew Gelman, Aki Vehtari, Daniel Simpson, Charles C. Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Bürkner, and Martin Modrák. Bayesian workflow, 2020. arXiv:2011.01808.

[9] Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1):1–32, 2017. doi: 10.18637/jss.v076.i01.

[10] Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989. doi: 10.1007/BF01589116.

[11] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering (Transactions of the ASME, Series D)*, 82(1):35–45, 1960. doi: 10.1115/1.3662552.

[12] Simo Särkkä and Lennart Svensson. *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, Cambridge, second edition, 2023. ISBN 978-1-108-92664-5. doi: 10.1017/9781108917407.

[13] Steve B. Howell. *Handbook of CCD Astronomy*. Cambridge Observing Handbooks for Research Astronomers. Cambridge University Press, Cambridge, 2 edition, 2006. ISBN 978-0-521-85215-9.

[14] N. M. Law, C. D. Mackay, and J. E. Baldwin. Lucky imaging: high angular resolution imaging in the visible from the ground. *Astronomy & Astrophysics*, 446(2):739–745, 2006. doi: 10.1051/0004-6361:20053695.

[15] J. P. Lewis. Fast template matching. In *Vision Interface 95*, pages 120–123, Quebec City, Canada, 1995. Canadian Image Processing and Pattern Recognition Society. Expanded and corrected version: *Fast Normalized Cross-Correlation*, Industrial Light & Magic, http://scribblethink.org/Work/nvisionInterface/nip.pdf.

[16] C. D. Kuglin and D. C. Hines. The phase correlation image alignment method. In *Proceedings of the IEEE International Conference on Cybernetics and Society*, pages 163–165, 1975.

[17] Joseph W. Goodman. *Introduction to Fourier Optics*. W. H. Freeman, New York, 4 edition, 2017. ISBN 978-1-319-11916-4.

[18] Michael Potmesil and Indranil Chakravarty. A lens and aperture camera model for synthetic image generation. In *Proceedings of the 8th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '81)*, pages 297–305. ACM, 1981. doi: 10.1145/800224.806818.

[19] Changyin Zhou and Stephen Lin. Removal of image artifacts due to sensor dust. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Minneapolis, MN, June 2007.

[20] Reg G. Willson, Mark W. Maimone, Andrew E. Johnson, and Larry M. Scherr. An optical model for image artifacts produced by dust particles on lenses. In *Proc. 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Munich, Germany, September 2005. ESA. ESA SP-603.

[21] Jinwei Gu, Ravi Ramamoorthi, Peter Belhumeur, and Shree Nayar. Dirty glass: Rendering contamination on transparent surfaces. In *Rendering Techniques 2007 (Proc. Eurographics Symposium on Rendering)*, pages 159–170, 2007. doi: 10.2312/EGWR/EGSR07/159-170.

[22] Tianfan Xue, Michael Rubinstein, Ce Liu, and William T. Freeman. A computational approach for obstruction-free photography. *ACM Transactions on Graphics*, 34(4), 2015. doi: 10.1145/2766940. Proc. SIGGRAPH 2015.

[23] Hendrik C. van de Hulst. *Light Scattering by Small Particles*. John Wiley & Sons, New York, 1957. Reprinted: Dover Publications, New York, 1981; ISBN 978-0-486-64228-4.

[24] Craig F. Bohren and Donald R. Huffman. *Absorption and Scattering of Light by Small Particles*. John Wiley & Sons, New York, 1983. Reprinted: Wiley Professional Paperback Edition, Wiley-VCH, Weinheim, 1998; ISBN 978-0-471-29340-8.

[25] Matthew J. Berg, C. M. Sorensen, and A. Chakrabarti. A new explanation of the extinction paradox. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 112(7):1170–1181, 2011. doi: 10.1016/j.jqsrt.2010.08.024.

[26] Adarsh Deepak and Michael A. Box. Forwardscattering corrections for optical extinction measurements in aerosol media. 1: Monodispersions. *Applied Optics*, 17(18):2900–2908, 1978. doi: 10.1364/AO.17.002900.

[27] Vadim A. Markel. What is extinction? operational definition of the extinguished power for plane waves and collimated beams. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 246:106933, 2020. doi: 10.1016/j.jqsrt.2020.106933.

[28] Michael I. Mishchenko, Matthew J. Berg, Christopher M. Sorensen, and Cornelis V. M. van der Mee. On definition and measurement of extinction cross section. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 110(4-5):323–327, 2009. doi: 10.1016/j.jqsrt.2008.11.010.

[29] Shaodi You, Robby T. Tan, Rei Kawakami, Yasuhiro Mukaigawa, and Katsushi Ikeuchi. Adherent raindrop modeling, detection and removal in video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9):1721–1733, 2016. doi: 10.1109/TPAMI.2015.2491937.

[30] Michal Uricar, Pavel Krizek, Ganesh Sistu, and Senthil Yogamani. SoilingNet: Soiling detection on automotive surround-view cameras. In *Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 67–72, Auckland, New Zealand, 2019. doi: 10.1109/ITSC.2019.8917178.

[31] H. Chen-Chen, S. Pérez-Hoyos, A. Sánchez-Lavega, and J. Peralta. Characterisation of deposited dust particles on Mars InSight lander Instrument Context Camera (ICC) lens. *Icarus*, 392:115393, 2023. doi: 10.1016/j.icarus.2022.115393.

[32] H. M. Lai, W. Y. Wong, and W. H. Wong. Extinction paradox and actual power scattered in light beam scattering: a two-dimensional study. *Journal of the Optical Society of America A*, 21(12):2324–2333, 2004. doi: 10.1364/JOSAA.21.002324.

[33] Adelchi Azzalini and Marc G. Genton. Robust likelihood methods based on the skew-t and related distributions. *International Statistical Review*, 76(1):106–129, 2008. doi: 10.1111/j.1751-5823.2007.00016.x.

[34] Carmen Fernández and Mark F. J. Steel. On Bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441):359–371, 1998. doi: 10.1080/01621459.1998.10474117.

[35] Yaming Yu and Xiao-Li Meng. To center or not to center: That is not the question—an ancillarity–sufficiency interweaving strategy (ASIS) for boosting MCMC efficiency. *Journal of Computational and Graphical Statistics*, 20(3):531–570, 2011. doi: 10.1198/jcgs.2011.203main. With discussion and rejoinder.

[36] Andrew Smith, Jeremy Bailey, J. H. Hough, and Steven Lee. An investigation of lucky imaging techniques. *Monthly Notices of the Royal Astronomical Society*, 398(4):2069–2073, 2009. doi: 10.1111/j.1365-2966.2009.15249.x.