Journal of Research High School

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

The Impact of Artificial Light at Night on the Flight and Stopover of Nocturnal Migratory Birds in the South San Francisco Bay Area

Jingzhi Gu[1] *

[1]Henry M. Gunn High School, Palo Alto, CA, USA
*Corresponding Author: jeangu123580@gmail.com

Advisor: Junyan Wang, junyan.wang@cedarnature.org

Received August 14, 2026; Revised September 1, 2026; Accepted September 14, 2026

Abstract

Artificial light at night (ALAN) is one of the fastest-growing forms of global environmental change, and nocturnal migrants that orient by celestial cues are expected to be highly sensitive to it. Landscape-scale radar studies report that brighter skies predict higher stopover density, but whether that reflects behavior or the geography lighting accompanies depends on how separable brightness is from terrain and habitat in the region sampled. Two hypotheses were tested on the same nights with one model and one evidence standard: (H1) the stopover association would not survive control for habitat, elevation, and spatial autocorrelation; (H2) an in-flight association with flight coherence, indexed by Doppler spectrum width, would survive that control and replicate. Data came from one mountaintop weather radar over the South San Francisco Bay across three springs and 49 clear-sky nights, with coefficient sign agreement across all three seasons required to accept an effect. The results were consistent with both predictions: the stopover association was significant in one of three years and vanished when the brightest 6% of pixels were removed, while the in-flight association held in all three years and on all 49 nights, emerged only above a repeatable threshold near 19 nW cm^{-2} sr^{-1}, and persisted at 2.3 to 2.7 km, where ground clutter is geometrically impossible. Because the design is observational and cross-sectional, these are associations; the study cannot establish that artificial light caused the wider spectra.

Keywords: Artificial light at night, Nocturnal bird migration, Weather surveillance radar, Doppler spectrum width, Spatial confounding, Skyglow threshold

1. Introduction

Artificial sky brightness is expanding by several percent per year, and its ecological consequences are recognized as a distinct class of anthropogenic disturbance (Longcore and Rich, 2004; Falchi et al., 2016; Kyba et al., 2017). Billions of songbirds migrate at night across North America (Dokter et al., 2018), navigating by celestial and horizon cues that make them mechanistically vulnerable to ALAN (Burt et al., 2023), and passage areas overlap the brightest landscapes (Cabrera-Cruz et al., 2018; Horton et al., 2019). The South San Francisco Bay lies on the Pacific Flyway and is among the most densely lit regions in the western United States.

At intense point sources birds aggregate, circle, and lose airspeed (Van Doren et al., 2017), driving collision mortality at lit buildings (Van Doren et al., 2021); at landscape scale, radar and citizen-science studies report that skyglow predicts higher stopover density (La Sorte et al., 2017; Cabrera-Cruz et al., 2020; Horton et al., 2023). McLaren et al. (2018) found a scale-dependent response, and Horton et al. (2023) found skyglow only weakly collinear with habitat at continental scale. A single radar in a topographically compressed coastal basin is different: the illuminated land is the valley floor and the dark land is the mountain ridge, so brightness, elevation, and land cover may not be separable, and a model without explicit spatial structure can return a light coefficient that reflects position rather than brightness.

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

The second outcome needs a brief definition before the hypotheses are stated. Every radar sample volume contains many birds at once, and Doppler spectrum width is the standard deviation of the radial velocities returned from that volume. When the birds inside share a heading and speed, their echoes share nearly one velocity and the spectrum is narrow; when their headings scatter, the same volume returns a wider mixture of velocities. Spectrum width is used here as an indirect proxy for flight coherence, with larger values meaning the birds present agreed less about which way to go. It is a proxy rather than a measurement of orientation, and it can be widened by physical as well as behavioral causes, so Section 2.3 gives its formal definition and the screens applied to it, and Section 4.3 states its limitations.

This study tested both questions within one three-times-replicated design and one evidence standard, so the two outcomes check each other. H1: the stopover association would not survive control for habitat, elevation, geography, and spatial autocorrelation, and would not reproduce across years. H2: an in-flight association would survive the same controls and reproduce, because flight coherence is measured within a cell rather than compared across the landscape. RQ1: after full control, does ALAN predict stopover density reproducibly across three springs? RQ2: does ALAN predict increased Doppler spectrum width in the low-altitude layer through which departing migrants climb, net of density and contamination, and is the response linear or thresholded? RQ3: does altitude-scaled skyglow predict increased spectrum width across the observable airspace, including the clutter-free upper layers?

2. Materials and Methods

2.1 Radar data, night selection, and quality control

NEXRAD Level II volumes from station KMUX (WSR-88D; Mt. Umunhum, 37.155° N, 121.898° W, antenna ≈1,082 m above sea level) were accessed via NOAA NCEI and its AWS mirror (National Oceanic and Atmospheric Administration, 1991). Forty-nine clear-sky spring nights under the clear-air scan pattern VCP-35 were analyzed (2024: 16; 2025: 17; 2026: 16). Nights were chosen in two stages, both completed before any radar field was extracted and before any light layer was joined. First, the BirdCast migration dashboard (Van Doren and Horton, 2018; Cornell Lab of Ornithology, 2026), which issues nightly forecast migration traffic rates for the continental United States, was used to shortlist spring nights forecast to carry substantial passage over Santa Clara County. Second, each shortlisted night was checked against Weather Underground station records for the San Jose area and retained only if it showed zero measurable precipitation and sustained surface winds below 10 m/s through the analysis window. Selection therefore rested on forecast migration intensity and observed weather alone; no night was added or removed after its spectrum width, stopover density, or light values had been inspected. Analysis used the 5 to 35 km annulus. A terrain beam-blockage mask (Figure 1) excludes shadowed pixels from the stopover arm (3.7% of the grid); shadowed gates are removed from the in-flight arm by the velocity-texture and clutter-power screens. Texture-based gate filtering replaced a reflectivity floor that would delete weak biological echoes (Dokter et al., 2011); velocities were dealiased with Py-ART (Helmus and Collis, 2016); and dual-polarization screening removed 12.6% of in-flight cells. Light layers are annual and per-year where available: the 2024 and 2025 VIIRS composites are independent of one another, but the 2026 composite had

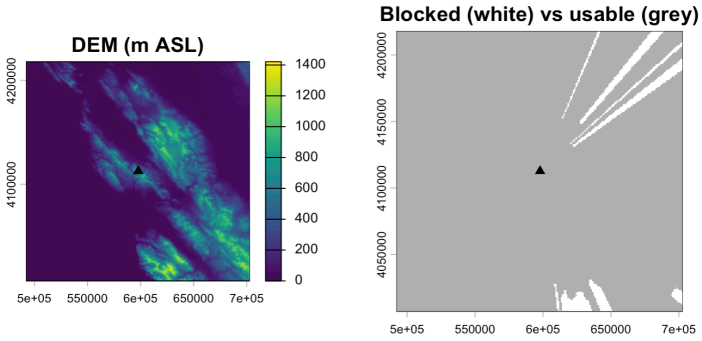


Figure 1. Terrain beam-blockage mask used to exclude azimuths the radar cannot cleanly see (EPSG:32610; black triangle = KMUX). Left: terrain elevation from the USGS 3D Elevation Program, in meters above sea level, showing the Santa Cruz Mountains ridge. Right: the resulting mask, white = azimuths blocked by terrain, gray = usable airspace. Because blockage is evaluated per bearing, comparing the steepest terrain horizon along each azimuth against the lowest beam edge, shadowing appears as narrow angular wedges radiating outward. Only gray cells enter the analysis. Created by the author in R (terra, elevatr).

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

not been released at the time of analysis, so the 2025 composite was reused for the 2026 nights. The three seasons are therefore independent replicates of the radar and bird observations, and two independent realizations of the light field; statements of three-year replication in this paper refer to independently sampled radar seasons, not to three independent light layers.

2.2 Sample structure

Extraction yielded approximately 248,000 volume-scan × sector × band samples in flight and approximately 1.17 million altitude-binned transit samples, aggregated to 3,528 night-level in-flight cells of 30° × 5 km and night-level cells in five altitude layers (≈1.3 to 3.0 km above sea level) for transit; per-night correlations use each night's 72 cells. The in-flight samples are not drawn from a restricted window after departure: they come from the lowest high-Nyquist Doppler sweep of every volume scan recorded between civil dusk and civil dawn, 3,465 volumes in total across the 49 nights, 60 to 114 per night with a median of 70. Low altitude in this paper therefore refers to the beam heights sampled, approximately 1.3 to 2.0 km above sea level, and not to a short interval of the night. The stopover grid held 30,608 pixels at 1 km, of which 3,740 fully attributed core pixels entered the per-year test; spline fits used a fixed-seed random subsample of 1,500 pixels, and alternative seeds reproduced the same non-significant coefficients. The altitude windows are anchored to the observed migration profile (Figure 2), which peaks near 1.6 to 1.7 km above sea level in all three seasons.

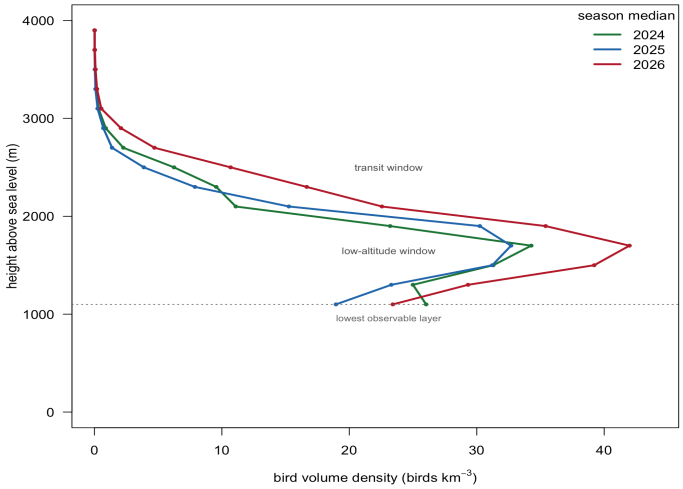


Figure 2. Season-median vertical profile of bird volume density, used to set the altitude windows. Horizontal axis: volume density in birds km^{-3} from vol2bird; vertical axis: height above sea level in meters; each line is the median across all nights of one spring within 200 m layers (2024 green, 16 nights; 2025 blue, 17; 2026 red, 16). Labels mark the low-altitude (1.3–2.0 km) and transit (2.0–3.0 km) windows. All three seasons peak near 1.6–1.7 km, fall steeply through 2.5 km, and are negligible above 3 km. The profile begins at approximately 1.1 km because the beam reaches no lower at these ranges; any denser near-ground passage is unobserved. Created by the author in R (bioRad).

2.3 Variables and light exposure metrics

Six quantities carry the analysis (Table 1). Stopover density is range-corrected reflectivity at the post-sunset emergence window, when birds that rested on the ground lift into the beam together, converted to volume density with vol2bird through bioRad (Dokter et al., 2011; Dokter et al., 2019) assuming an 11 cm² cross-section per passerine, so densities are relative. The emergence window was defined astronomically rather than by clock time. For each night the target instant was the moment the sun reached 5.5° below the horizon at the radar, approximately 30 minutes after local sunset, which is the evening-exodus timing used by McLaren et al. (2018). The two volume scans immediately bracketing that instant were processed and their plan-position-indicator surfaces interpolated linearly in time to the target, so each night's stopover surface rests on exactly two volume scans, 98 across the 49 nights. Consecutive VCP-35 volumes are separated by approximately six minutes, so the interpolated surface represents a window of a few minutes around the target instant. Doppler spectrum width, the coherence index, is the standard deviation of radial velocities within one sample volume (Doviak and Zrnic, 1993): shared headings return a narrow spectrum, scattered headings a wide one. Because gate-level values are noisy at low densities (Stepanian et al., 2016), they were averaged within 30° × 5 km cells pooling hundreds of gates; excluding windy nights removes the dominant physical broadening source (Nastrom, 1997), and a range term in every model absorbs beam broadening (Figure 3).

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

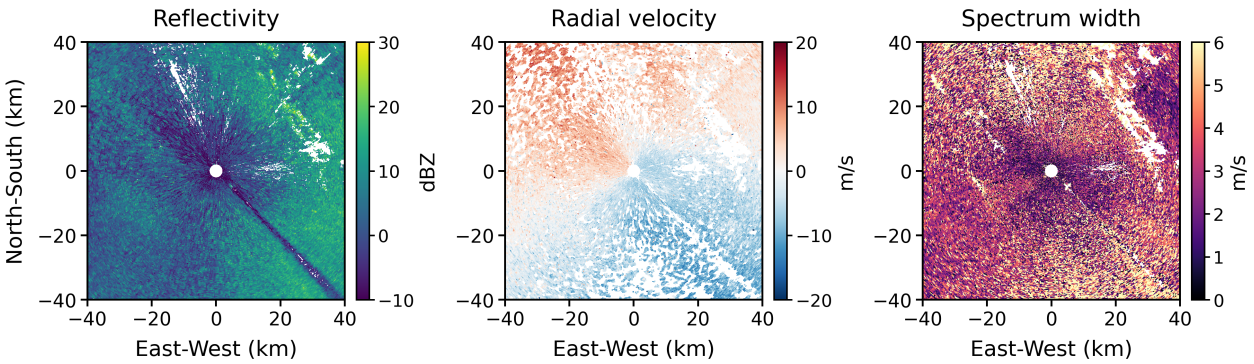


Figure 3. The three raw radar moments behind the coherence metric, for one volume scan (4 April 2026, 07:04:32 UTC). Axes are east–west and north–south distance from the radar in kilometers (white dot at origin). Left: reflectivity in dBZ; widespread moderate values confirm biological echo across the domain. Center: radial velocity in m/s; the clean two-lobed couplet is the signature of a coherent broad-front migration and confirms successful dealiasing. Right: Doppler spectrum width in m/s, the disorder metric; brighter values mark volumes in which the birds present disagreed more about heading. Data are unmasked; the low-reflectivity diagonal streak is a terrain-shadowed sector removed by the Figure 1 mask before analysis. Created by the author in Python (Py-ART, matplotlib).

Table 1. Variables used in the analysis, with the physical quantity each represents, its source, and the scale on which it entered the models.

Variable	Physical quantity it represents	Source	Scale used in models
Stopover density	Relative abundance of birds departing the ground in the first post-sunset window	KMUX Level II via bioRad / vol2bird	Birds km^{-3}; Gamma response, log link
Spectrum width (sw_mean)	Spread of flight velocities within one radar sample volume; higher means less coherent	KMUX Level II, high-Nyquist split cuts	m/s; Gamma response, log link
Ground brightness	Upward radiance from artificial lighting at the surface	VIIRS VNL V2 annual composites	ln(VIIRS+1), from raw nW cm^{-2} sr^{-1}
Skyglow	Scattered light reaching a bird at a given flight altitude	VIIRS convolved per Duriscoe et al. (2018)	Dimensionless log index
Land cover / NDVI / elevation	Habitat and terrain confounds entangled with light	NLCD; Sentinel-2; USGS 3DEP	Three-level factor; index −1 to +1; m ASL
Range, clutter power, reflectivity	Radar viewing geometry, ground-clutter residue, and echo strength	KMUX Level II	km, dB, dBZ; all z-standardized

Ground brightness is VIIRS Day/Night Band annual-composite radiance (Elvidge et al., 2021), modeled as ln(VIIRS+1) because raw radiance is right-skewed and zero over dark pixels; ln(VIIRS+1) = 3 corresponds to approximately 19 nW cm^{-2} sr^{-1}. Skyglow, used only in the transit arm, is what a bird at altitude sees rather than the brightness of the ground directly beneath it: VIIRS radiance is summed over every source pixel within the height-dependent geometric horizon of Equation 1 and distance-weighted by the kernel of Equation 2, following Duriscoe et al. (2018), after Garstang (1986).

$$d_max = \sqrt{(\ 2 \cdot (4/3) \cdot R_e \cdot h\ )} \tag{1}$$

$$skyglow = \log(\ 1 + \Sigma_j\ [\ VIIRS_j \cdot \exp(-d_j / 20{,}000) / (d_j^2 + 10^6)\ ]\ ) \tag{2}$$

In Equation 1, d_max is the geometric visual horizon in meters for a bird at height h, R_e is the radius of the Earth, and the factor 4/3 is the standard effective-radius correction for atmospheric refraction; h is terrain-corrected, computed as antenna elevation plus beam height above the antenna minus terrain elevation beneath the gate, so the horizon matches the bird's height above the ground below it rather than above the mountaintop antenna. In Equation 2, the sum runs over every VIIRS source pixel j lying within d_max, $VIIRS_j$ is the radiance of pixel j in nW cm^{-2} sr^{-1},

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

and d_j is the ground distance in meters from the bird to pixel j. The exponential term, with a length scale of 20,000 m, attenuates distant sources through the clear-sky atmosphere; the inverse-square term accounts for geometric spreading; and the added 10^6 in the denominator prevents the expression from diverging for sources directly beneath the bird. The resulting index is dimensionless, and a higher value means greater modeled exposure to artificial light at the sampled flight altitude. Because all 49 nights were screened to be clear, the clear-sky scattering assumption behind the kernel is satisfied by design.

Habitat and terrain covariates control the confound (NLCD Urban/Water/Natural, Multi-Resolution Land Characteristics Consortium, 2024; Sentinel-2 NDVI; USGS 3DEP elevation accessed with elevatr, Hollister et al., 2025; distance to water); ρHV and ZDR act as contamination screens rather than covariates, while range, clutter power, and reflectivity enter the models directly to absorb non-biological variation, and gate count serves as the weight in the sensitivity refit. All layers share a 1 km grid in EPSG:32610.

2.4 Statistical models and software

Both arms were fitted separately per year in R (R Core Team, 2025) with the spaMM package function fitme() (Rousset and Ferdy, 2014) as Gamma(log) mixed models with a Matérn spatial random field over cell or pixel centroids and a night random intercept. For in-flight cell-mean spectrum width SW in cell i on night t:

$$E[SW_it] = \exp(\beta_0 + \beta_1 \cdot light_i + \beta_2 \cdot range_i + \beta_3 \cdot clutter_it + \beta_4 \cdot refl_it + u(x_i, y_i) + b_t) \quad (3)$$

where light is ln(VIIRS+1); range, clutter power, and reflectivity are z-standardized; u(x, y) is the Matérn field with smoothness fixed at $\nu = 0.5$, the conservative rough-field choice; and b_t is the night random intercept. The stopover arm used the same family and spatial structure over 1 km pixels, with linear and quadratic light plus the habitat covariates. On the log link every coefficient is multiplicative (+0.10 ≈ +10% per unit). Generalized additive models (mgcv; Wood, 2017) were used adversarially: concurvity and residual Moran's I (Moran, 1950) on a spatial-smooth GAM are what exposed the stopover confound. Response shape was tested with an eight-level light factor and a natural spline inside the same Matérn model, with replication scored as the mean pairwise correlation among year-curves. The eight levels are the fixed-width 0.5-unit intervals of ln(VIIRS+1) that tile the observed brightness range from 0 to 4.39 across in-flight cells and 0 to approximately 4.76 across stopover pixels, with the top level left open above 3.5 because too few cells lie beyond it to support a narrower bin; the grid was therefore fixed by the range and sample size of the covariate rather than chosen from the fitted response, and the same edges were used in the stopover analysis. A width of 0.5 units was the narrowest that kept most bins populated: it divides the observed range into eight levels, whereas halving it would have produced fifteen, six or seven of which would have been empty or held two or fewer of the 72 cells. Because brightness in this domain is concentrated at the dark end, the bins are unevenly populated: of the 72 sector × band cells, 35 to 36 fall in the darkest bin, while the 1.5-2 bin holds two cells in every year and the 2.5-3 bin holds two cells in 2024 and one in 2025 and 2026. Those two bins are estimable but weakly constrained, which is why their confidence intervals are the widest in the series and why no feature of the response shape rests on them. The acceptance standard for every claim was coefficient sign agreement across all three years. Analyses were run in R 4.6.0 using spaMM 4.6.65, mgcv 1.9.4, bioRad 0.11.0, terra 1.9.27, and elevatr 0.99.1, and in Python 3.14.6 using Py-ART 2.2.2, NumPy 2.4.6, SciPy 1.17.1, and matplotlib 3.10.9.

2.5 Validation, robustness, and disclosed limitations

Four go/no-go checks addressed artifacts that could imitate a light effect. First, metric validity: a Monte Carlo simulation, run before any real data were modeled, confirmed that spectrum width rises monotonically with heading scatter, unlike the near-flat velocity-azimuth-display residual, and that realistic contrasts are detectable far below the available cell counts (Figure 4). Second, contamination: dual-polarization screening (ρHV > 0.95, ZDR > 4 dB, top-decile clutter power; Stepanian et al., 2016) removed 445 of 3,528 cells; screening strengthened the light–disorder correlation from +0.48 to +0.59 (Figure 5), the opposite of what insect-generated signal would produce; and retaining

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

every cell but adding ZDR, clutter power, reflectivity, and ρHV to the model reduced the light coefficient from +0.088 to +0.062, so 70% of the association survives simultaneous control for all four contamination signatures. Third, range: brightness is organized by compass direction rather than distance, so every 5 km band spans a wide brightness range and light is estimable separately from range (Figures 6 and 7). Fourth, effort: refitting all three in-flight models with observations weighted by gate count returned +0.031 (0.011, 0.050) in 2024, +0.032 (0.011, 0.053) in 2025, and +0.025 (0.009, 0.040) in 2026, that is 3% to 7% larger than the unweighted values reported in Table 2, so the specification used throughout is the more conservative of the two.

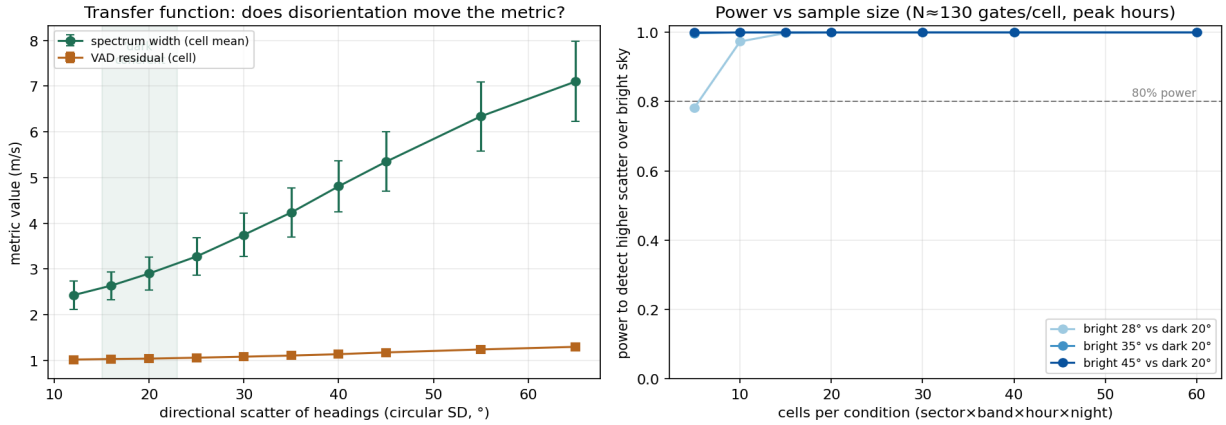


Figure 4. The coherence metric responds to heading scatter, and the study is adequately powered to detect it; both panels are Monte Carlo simulation, not observational data. Left: transfer function, imposed circular heading scatter in degrees against the resulting metric in m/s; cell-mean spectrum width (green) rises monotonically and steeply while the velocity-azimuth display (VAD) residual (orange) is nearly flat, which is why spectrum width was adopted. Error bars: interquartile range across runs; shaded band: scatter typical of real migration. Right: power, probability of detecting higher scatter over bright sky against cells per condition; dotted line marks 80% power. Against a 20° dark baseline, the 35° and 45° bright conditions are detected essentially always, and the 28° condition exceeds 80% power by approximately 10 cells, all far below the 1,152 to 1,224 cells available per year. Created by the author in R.

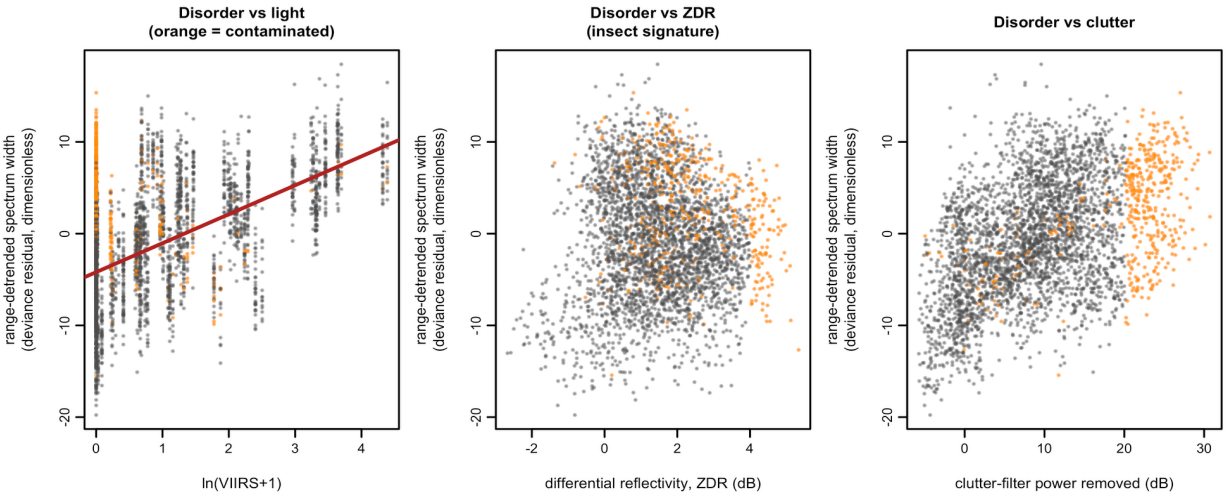


Figure 5. Disorder tracks light, and the one contamination signature it also tracks is itself unrelated to light, so the light effect is not a filtering artifact (all 49 nights; 3,528 cells; one point per cell; orange = flagged by the dual-polarization screen). Vertical axis in all panels: disorder, the weighted deviance residual of cell-mean spectrum width from a Gamma(log) model with a range spline and night factor, so zero means exactly as coherent as expected for that distance and night. Left: against ln(VIIRS+1); clearly positive (r = +0.59 on clean cells; red line = clean-cell fit). Center: against differential reflectivity, the insect signature; no relationship (r = +0.02). Right: against clutter-filter power; r = +0.46, but clutter power is almost uncorrelated with light (r = −0.01), so this association cannot generate the light effect. Created by the author in R.

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

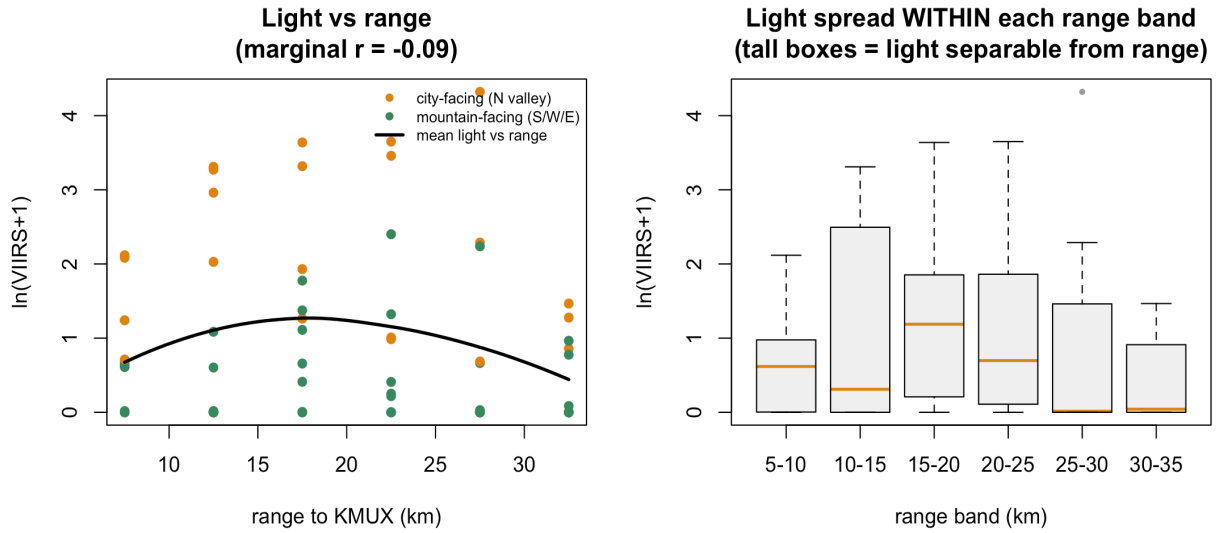


Figure 6. Ground brightness varies by compass direction rather than by distance from the radar, so it can be estimated separately from range. Left: ln(VIIRS+1) against range in kilometers for the 72 sector × band modeling cells; orange = the four city-facing sectors (300° to 60°; mean ln(VIIRS+1) 1.27 to 2.88), green = the eight mountain-facing sectors (mean 0.07 to 0.75), black line = loess fit, rising to approximately 18 km where the beam crosses the urban core and then declining. Individual cells overlap between the two groups, because mountain-facing sectors clip lit valley edges at some ranges, but the group means are well separated, and the marginal light–range correlation of −0.09 understates the directional structure visible here. Right: distribution of ln(VIIRS+1) within each 5 km range band as boxplots; interquartile ranges span 0.84 to 2.26 units, so every band contains cells across much of the brightness gradient, with the 10–15 km band the most heterogeneous. Created by the author in R.

Six limitations are disclosed. The radar sees nothing below approximately 1.1 km above terrain, so the densest near-ground passage is unobserved. Light is cross-sectional: annual composites compare brighter places, not brighter nights. Insect contamination remains a live concern for the stopover arm, whose apparent effect sits in the bright tail where contamination concentrates. Skyglow indexes urbanness broadly. The 2025 composite is reused for 2026, so 2026 is an independent replicate of the radar observations over a repeated light surface, and 2024, with its own light layer, is the decisive independent replicate of the light–response relationship. Wind screening used surface rather than flight-altitude observations.

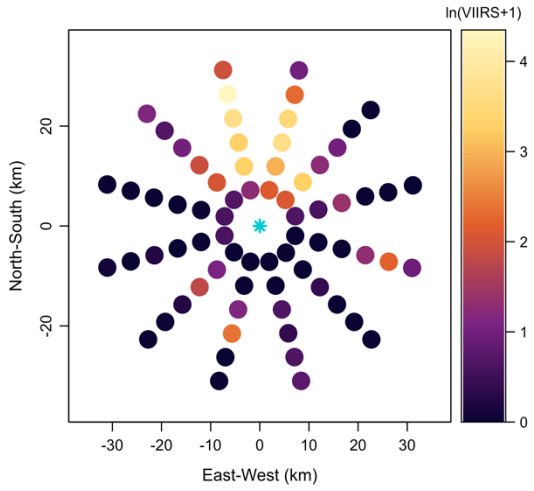


Figure 7. The ground-brightness field across the analysis domain. Each disc is one 30° × 5 km cell, plotted in kilometers east–west and north–south from the radar at the origin, colored by ln(VIIRS+1) from dark through orange to pale yellow. Brightness concentrates in a wedge extending north and northeast toward the urban core, while sectors facing the Santa Cruz Mountains remain dark at all ranges; this directional structure, not any radial trend, is what makes brightness separable from range (Figure 6). Cell geometry is fixed and the composites are annual; the 2025 composite is shown (used for 2025 and 2026). Created by the author in R.

3. Results

3.1 Baseline models without spatial control

Both regimes begin identically. A spatial GAM without explicit spatial-covariance structure returned a strong positive light effect for stopover density (light smooth $p < 2 \times 10^{-16}$; 93% deviance explained; 2026 model, similar in 2024 and 2025) and, separately, a strong positive light–disorder relationship. The remaining analyses apply the same controls to each regime.

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

3.2 Stopover density and ALAN (RQ1)

Diagnostics on the stopover GAM returned concurvity of approximately 0.96 for the light term and residual Moran's I of 0.36 even at near-maximum smooth flexibility (k = 100), so light, NDVI, and position were nearly interchangeable. The Matérn GLMM gave a pooled effect (likelihood-ratio $\chi^2 = 9.92$, df = 2, p = 0.007) rising only in the extreme bright tail (Figure 8) but failed every robustness check: significant in 2026 only, and removing the brightest 5.9% of pixels (ln(VIIRS+1) > 3.5) collapsed it to non-significance. Elevation was the only consistently significant predictor.

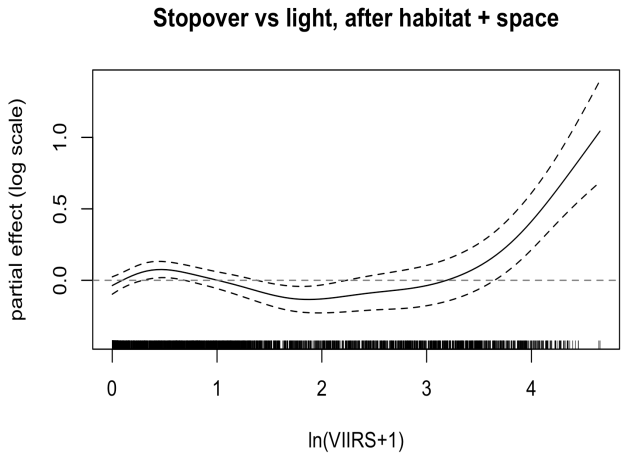


Figure 8. After habitat and space are controlled, stopover density is flat against light across almost the entire brightness gradient. Horizontal axis: ln(VIIRS+1); vertical axis: partial effect of light on the log scale of the Gamma Matérn GLMM, so zero (dashed line) means no difference from the average pixel once habitat, elevation, distance to water, and position are accounted for; dashed curves are the 95% confidence band. Bottom ticks show where pixels lie: the great majority sit below 3. The curve is statistically indistinguishable from flat across that bulk and rises only above approximately 3.5, where few pixels remain and the band widens sharply. Created by the author in R (spaMM).

3.3 In-flight coherence and ALAN (RQ2)

On the same nights, the disorder result behaved oppositely. The light–disorder correlation was positive on all 49 nights, with mean per-night Pearson r of +0.42 to +0.51 and a Wilcoxon signed-rank p below 10^{-4} in each year (Figure 9); no night in any year reversed sign. The contamination-controlled Matérn models gave a positive light coefficient with a 95% confidence interval clear of zero in every year at near-identical magnitude (Table 2); across the observed brightness range this corresponds to approximately 14% wider spectra.

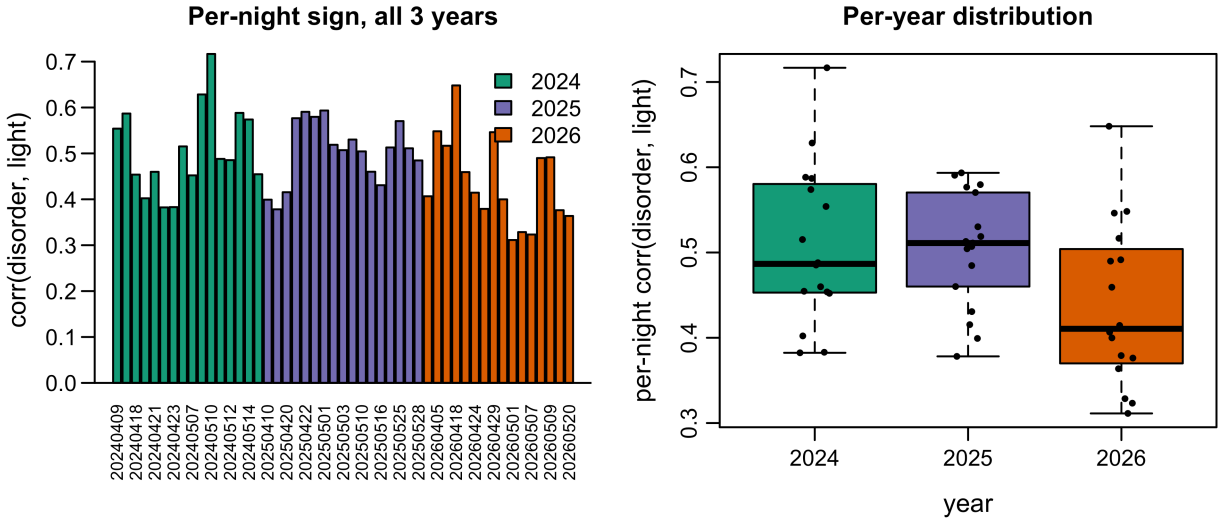


Figure 9. The light–disorder association is positive on every one of the 49 clear-sky nights, in all three years. Left: one bar per night, ordered chronologically within year (2024 green, 2025 purple, 2026 orange); the vertical axis is the within-night Pearson correlation between cell-mean disorder and cell ground brightness across that night's 72 sector × band cells. Every bar is above zero. Right: the same 49 values summarized by year as boxplots with individual nights overplotted; medians sit near +0.5 in 2024 and 2025 and near +0.43 in 2026, so the association is of similar strength in three independently sampled radar seasons. Created by the author in R.

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

3.4 A replicated brightness threshold

Fitting light as an eight-level factor, with the darkest bin (0 to 0.5) as reference, draws the response shape with no linearity assumption; the bin edges were fixed in advance from the range of the covariate, as described in Section 2.4. The response is a threshold, not a ramp: in every year all bins below ln(VIIRS+1) = 3 have intervals spanning zero, and only the two brightest bins, 3 to 3.5 and above 3.5, are elevated with intervals excluding zero (Figure 10, Table 3). The brightest bin is the highest-disorder bin in every year. The threshold sits near ln(VIIRS+1) = 3, approximately 19 nW cm^{-2} sr^{-1}, and its position, not merely its existence, is the same in three independently sampled radar seasons. The location of that threshold was read off the fitted bin coefficients rather than specified beforehand, so it is an exploratory, estimated value; what was specified in advance was the requirement that any feature of the response shape agree across all three seasons, which this one does.

Table 2. In-flight light coefficient by year, from the contamination-controlled Gamma(log) Matérn model: multiplicative log-link slopes of cell-mean spectrum width on ln(VIIRS+1), net of range, clutter power, reflectivity, sampling effort, spatial position, and night. A coefficient of +0.029 means spectrum width rises by approximately 2.9% per one-unit increase in ln(VIIRS+1), roughly a 2.7-fold increase in radiance. All intervals exclude zero.

Year	Light coefficient	95% confidence interval
2024	+0.029	(0.009, 0.049)
2025	+0.031	(0.009, 0.053)
2026	+0.024	(0.008, 0.040)

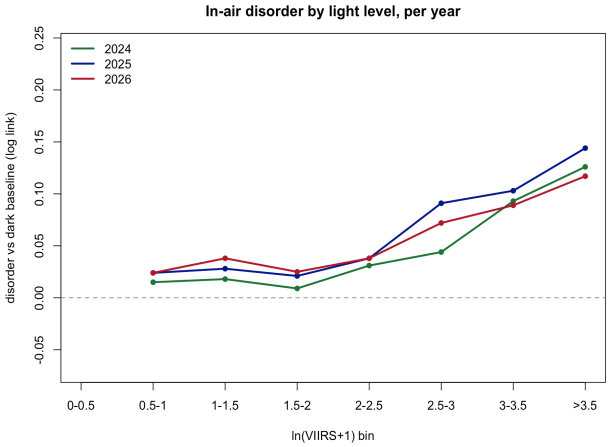


Figure 10. In-flight disorder is flat across the dark-to-moderate range and rises only above an urban-core brightness threshold, in all three years. Light was fitted as an eight-level factor, so the shape is drawn by the data. Horizontal axis: ln(VIIRS+1) bins, where 3 corresponds to approximately 19 nW cm^{-2} sr^{-1}; vertical axis: modeled change in cell-mean spectrum width relative to the darkest bin, on the log scale (+0.10 ≈ 10% wider). Each line is one spring (2024 green, 2025 blue, 2026 red). The three curves lie almost on top of one another, stay near zero out to the 1.5–2 bin, and rise sharply in the two brightest bins; estimates and intervals are in Table 3. The small irregularities in the middle of each curve fall in the most sparsely occupied bins and are not interpreted; see the occupancy note to Table 3. Created by the author in R (spaMM).

3.5 Comparison of the two response shapes

Continuous natural-spline fits compare the two regimes on the same brightness axis. The stopover curves diverge: within every land class the three year-curves cross and change shape (Figure 11). The disorder curves are nearly identical (mean pairwise year-curve correlation r = 0.986), with predicted disorder rising 16%, 17%, and 12% in 2024, 2025, and 2026 from darkest to brightest cells (Figure 12). The two analyses differ only in the response variable (Table 4).

Table 3. In-flight disorder by light bin relative to the darkest bin, from the binned-factor specification, for all three springs. Values are log-link coefficients with 95% confidence intervals; the reference category is ln(VIIRS+1) of 0 to 0.5, whose coefficient is zero by construction, so every entry is a proportional change in spectrum width relative to a dark cell. An asterisk marks bins whose interval excludes zero in all three years. No bin below ln(VIIRS+1) = 3 is elevated with an interval excluding zero in any year. Bin occupancy is uneven: the 1.5–2 and 2.5–3 bins rest on only one or two of the 72 sector × band cells, so their estimates are weakly constrained and their wide intervals should be read accordingly; the two brightest bins rest on three and four cells respectively.

Light bin, ln(VIIRS+1)	2024 coefficient (95% CI)	2025 coefficient (95% CI)	2026 coefficient (95% CI)
0 – 0.5 (reference)	–	–	–
0.5 – 1	+0.015 (−0.028, 0.059)	+0.024 (−0.020, 0.067)	+0.024 (−0.010, 0.058)
1 – 1.5	+0.018 (−0.029, 0.064)	+0.028 (−0.021, 0.076)	+0.038 (0.000, 0.076)
1.5 – 2	+0.009 (−0.069, 0.088)	+0.021 (−0.058, 0.099)	+0.025 (−0.037, 0.086)
2 – 2.5	+0.031 (−0.029, 0.092)	+0.038 (−0.020, 0.095)	+0.038 (−0.006, 0.082)
2.5 – 3	+0.044 (−0.040, 0.127)	+0.091 (−0.031, 0.213)	+0.072 (−0.023, 0.166)
3 – 3.5 *	+0.093 (0.017, 0.169)	+0.103 (0.021, 0.184)	+0.089 (0.026, 0.152)
Above 3.5 *	+0.126 (0.035, 0.217)	+0.144 (0.045, 0.242)	+0.117 (0.041, 0.192)

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

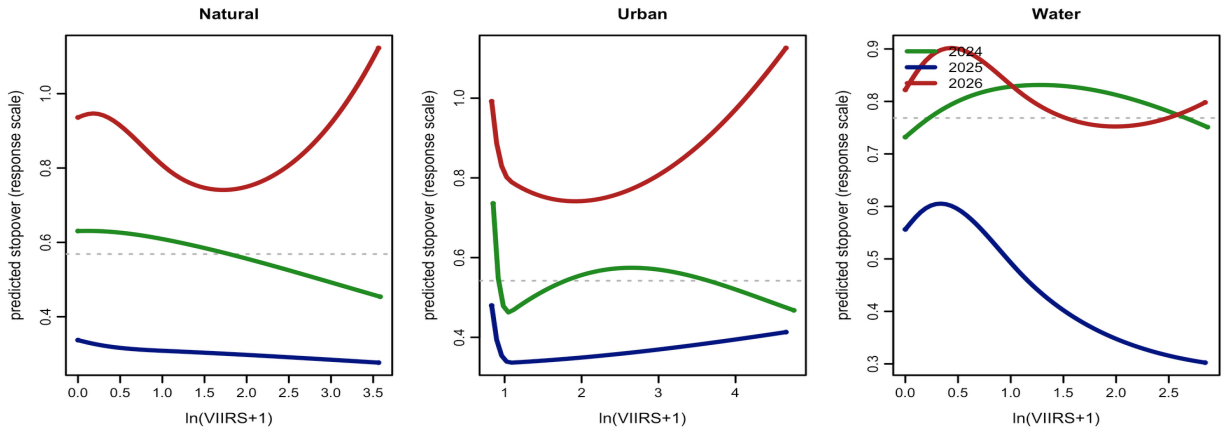


Figure 11. Stopover density against light does not reproduce across years in any land-cover class. Panels: Natural, Urban, and Water; curves: 2024 green, 2025 blue, 2026 red. Horizontal axis: ln(VIIRS+1); axis ranges differ between panels because the three land classes occupy nearly non-overlapping brightness ranges, which is itself the spatial confound described in Section 4.1. Vertical axis: predicted stopover density on the response scale in relative units; compare heights within a panel only. Within every panel the year-curves cross and disagree in shape as well as level. Compare with Figure 12. Created by the author in R (mgcv).

Table 4. The same analysis pipeline applied to both questions on the same 49 nights; the two columns differ only in the response variable being modeled.

Property tested	**Stopover (where they land)**	**In-flight disorder (how they fly)**
Radar metric	Range-corrected density	Doppler spectrum width
Baseline model, no spatial covariance	Strong positive light effect	Strong positive light effect
After Matérn GP and three-year test	Collapses (significant in 1 of 3 years)	Holds (3 of 3 years; 49 of 49 nights)
Bright-tail and contamination tests	Effect disappears	Effect strengthens
Spline shape across years	Diverges (no replication)	Nearly identical (r = 0.986)
Verdict	No robust ALAN effect	Robust, thresholded ALAN effect

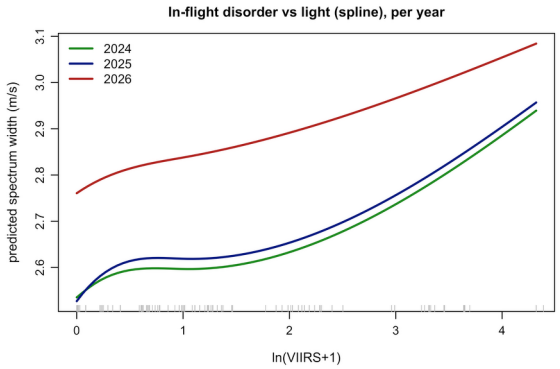


Figure 12. In-flight disorder against light reproduces almost exactly across three independently sampled radar seasons. Horizontal axis: ln(VIIRS+1); vertical axis: predicted cell-mean Doppler spectrum width on the response scale in m/s; bottom ticks show where the observed cells lie. All three curves (2024 green, 2025 blue, 2026 red) are flat from 0 to approximately 1.5 and rise steadily above 2, giving a mean pairwise year-curve correlation of r = 0.986. The 2026 curve sits approximately 0.15 m/s above the other two at every brightness, a baseline offset most plausibly reflecting seasonal wind and weather that does not alter the shape the replication standard tests. Compare with Figure 11. Created by the author in R (mgcv).

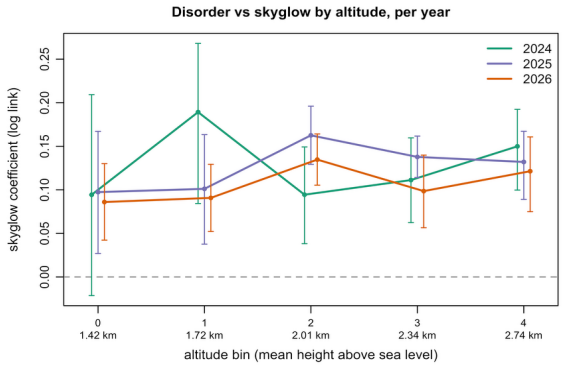


Figure 13. The skyglow–disorder association persists at every sampled altitude, including layers where ground clutter is physically impossible. Horizontal axis: altitude bin, from ≈350 m to ≈1700 m above the antenna, corresponding to mean heights of 1.42 to 2.74 km above sea level and 1.14 to 2.48 km above terrain. Vertical axis: skyglow coefficient on the log link (+0.10 ≈ 10% wider spectra per unit index); dashed line marks zero. Lines: 2024 green, 2025 purple, 2026 orange; bars are 95% confidence intervals. Every point in every year lies above zero, and the trend is flat to rising with height, the opposite of a ground-contamination artifact, whose influence necessarily weakens with altitude. Only the lowest 2024 interval touches zero. Created by the author in R (spaMM).

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

3.6 Transit altitude (RQ3)

Against altitude-scaled skyglow across five altitude bins and three years, the coefficient was positive in 15 of 15 year × altitude combinations, with 14 of 15 intervals excluding zero; the only exception was the lowest 2024 bin, +0.094 (−0.021, 0.209). In the decisive high bins at 2.34 and 2.74 km above sea level (2.07 and 2.48 km above terrain), where ground clutter is geometrically excluded, coefficients were positive with intervals clear of zero in all three years (Table 5), and the effect is flat to rising with altitude (Figure 13). Skyglow is on a different scale from ln(VIIRS+1), so only sign and consistency should be compared between Tables 2 and 5.

Table 5. Transit skyglow coefficient in the two highest, clutter-free altitude bins, with 95% confidence intervals. Bin 3 averages ≈1300 m and bin 4 ≈1700 m above the antenna, i.e. 2.34 and 2.74 km above sea level, or 2.07 and 2.48 km above terrain. Coefficients are log-link slopes of cell-mean spectrum width on the dimensionless skyglow index; all six intervals exclude zero.

Year	Bin 3 (≈2.3 km ASL)	Bin 4 (≈2.7 km ASL)
2024	+0.111 (0.062, 0.160)	+0.150 (0.100, 0.192)
2025	+0.138 (0.114, 0.162)	+0.132 (0.089, 0.167)
2026	+0.099 (0.056, 0.140)	+0.121 (0.075, 0.161)

4. Discussion

4.1 Two questions, one credible answer

Read together, the two regimes tell one story. ALAN does not reproducibly predict where migrants stop over the South Bay: the association collapses under spatial control, replication, and de-contamination, and the reason is identifiability rather than contradiction of prior work. Horton et al. (2023) found skyglow only weakly correlated with habitat continentally; in this basin the same variables are nearly interchangeable (concurvity ≈0.96, Moran's I 0.36), and the bright tail carrying the apparent effect is also where contamination is worst and land classes occupy non-overlapping brightness ranges. That describes what this design can resolve here, not evidence that the continental effect is absent; the domain also sits between the scales at which McLaren et al. (2018) found opposite-signed responses.

ALAN does, however, predict how coherently migrants fly. Stopover density compares places and requires light separable from habitat across the landscape; flight coherence compares birds within one sample volume on one night, so the spatial confound that disables the first comparison does not disable the second. Wider spectra over brighter skies are consistent with degraded celestial and horizon orientation widening heading scatter, the radar analogue of disorientation reported at intense installations (Van Doren et al., 2017; Burt et al., 2023). Consistency is not demonstration: the design is cross-sectional and compares brighter places rather than manipulating brightness, so it cannot show that light caused the wider spectra, and the metric cannot separate passive heading scatter from active maneuvering. Either would be ecologically meaningful if the association proves causal.

4.2 The threshold and its implications

The association is thresholded rather than linear. Cell-mean spectrum width is statistically indistinguishable from its dark-sky value across the whole dark-to-moderate range and is elevated only once radiance reaches urban-core levels near 19 nW cm^{-2} sr^{-1}, corresponding in the South Bay to centers such as central San Jose and the Milpitas corridor, while lower-density residential suburbs fall largely below ln(VIIRS+1) of approximately 2.5. The same approximately 6% brightest tail whose removal collapsed the stopover signal carries virtually all of the measured disorder, so the two arms point to the same small part of the landscape while attributing different consequences to it: not a change in where birds come down, but a change in how consistently they hold course above it. Because that tail carries nearly all of the measured association, dimming it would, if the relationship is causal, target most of the effect, which is consistent with the disproportionate reductions in collision mortality reported when a small number of intensely lit structures are extinguished (Van Doren et al., 2021); confirming that reducing radiance at a fixed location restores coherence would require a longitudinal intervention study. An alternative reading is that ln(VIIRS+1) = 3 marks where the biological signal clears the spectrum-width noise floor rather than a behavioral tipping point, but the threshold's stability across three seasons with independent atmospheric conditions argues against a purely instrumental origin, which would be expected to produce a variable crossing point.

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

4.3 Credibility, limitations, and future work

Inference was insulated against artifacts: the Gaussian process was deployed because diagnostics showed a simpler model was confounded; light is separable from range; same-sign results were required in three years; the metric was validated by simulation before real data; estimates were insensitive to effort weighting; and the decisive test sits at approximately 2.6 km, where ground clutter is physically impossible and high-flying insects (Chapman et al., 2010) are removed dual-polarimetrically. One feature of the contamination screen deserves explicit mention: the copolar correlation coefficient used in screening is itself associated with disorder ($r = -0.55$), because uniform targets return both high correlation and narrow spectra, so the screen systematically removes low-disorder cells. Those cells are disproportionately dark as well as low-disorder, so their deletion would be expected to weaken a positive light–disorder relationship rather than create one; the observed strengthening from +0.48 to +0.59 runs the other way. The stopover null shows the standard can return no effect, so the disorder positive is not manufactured. Two limits bound what any of this can support. Spectrum width is a proxy for flight coherence and not a measurement of orientation, and the design is observational, so the results identify an association and cannot establish that artificial light causes it. An overcast-versus-clear positive control is the key next step toward causation, since cloud cover independently degrades celestial cues and would provide a quasi-experimental contrast without a second radar. Future work should add a circular-statistics coherence endpoint, multi-radar mosaics, autumn seasons, and acoustic ground truth.

5. Conclusion

Using one radar, one framework, and a three-year sign-agreement standard, this study let two usually separate questions check each other. ALAN did not robustly predict where migrants stopped over, a well-diagnosed null, but did robustly predict how coherently they flew: positive in all three years and on all 49 nights, emerging only above a replicated urban-core threshold, and persisting clutter-free at approximately 2.6 km above sea level. Within the limits of an observational design, the measurable association with ALAN in this flyway is with flight coherence rather than site selection, and the pipeline is reusable at any WSR-88D station.

Acknowledgment

The author thanks Junyan Wang for methodological review of successive drafts. All radar, light, land-cover, and elevation data used in this study are public and are cited in the References. Analyses were performed with the open-source R packages bioRad, spaMM, mgcv, terra, and elevatr, and with the Python package Py-ART; figures were produced in R and in Python (matplotlib, geopandas), and mapping used QGIS. A large language model (Anthropic Claude) was used for code scaffolding and literature triage; all statistical results reported here were produced by the named open-source packages and were independently inspected by the author.

References

Burt, C. S., et al. (2023). The effects of light pollution on migratory animal behavior. *Trends in Ecology & Evolution, 38*(4), 355-368. https://doi.org/10.1016/j.tree.2022.12.006

Cabrera-Cruz, S. A., Smolinsky, J. A., & Buler, J. J. (2018). Light pollution is greatest within migration passage areas for nocturnally-migrating birds around the world. *Scientific Reports, 8*, 3261. https://doi.org/10.1038/s41598-018-21577-6

Cabrera-Cruz, S. A., et al. (2020). Artificial light at night is related to broad-scale stopover distributions of nocturnally migrating landbirds along the Yucatan Peninsula, Mexico. *Remote Sensing, 12*(3), 395. https://doi.org/10.3390/rs12030395

Chapman, J. W., et al. (2010). Flight orientation behaviors promote optimal migration trajectories in high-flying insects. *Science, 327*(5966), 682-685. https://doi.org/10.1126/science.1182990

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

Cornell Lab of Ornithology. (2026). *BirdCast migration dashboard: Santa Clara County, 2024-2026*. https://birdcast.org/migration-tools/migration-dashboard

Dokter, A. M., et al. (2011). Bird migration flight altitudes studied by a network of operational weather radars. *Journal of the Royal Society Interface, 8*(54), 30-43. https://doi.org/10.1098/rsif.2010.0116

Dokter, A. M., et al. (2018). Seasonal abundance and survival of North America's migratory avifauna determined by weather radar. *Nature Ecology & Evolution, 2*(10), 1603-1609. https://doi.org/10.1038/s41559-018-0666-4

Dokter, A. M., et al. (2019). bioRad: Biological analysis and visualization of weather radar data. *Ecography, 42*(5), 852-860. https://doi.org/10.1111/ecog.04028

Doviak, R. J., & Zrnic, D. S. (1993). *Doppler radar and weather observations* (2nd ed.). Academic Press.

Duriscoe, D. M., et al. (2018). A simplified model of all-sky artificial sky glow derived from VIIRS Day/Night band data. *Journal of Quantitative Spectroscopy and Radiative Transfer, 214*, 133-145. https://doi.org/10.1016/j.jqsrt.2018.04.028

Elvidge, C. D., et al. (2021). Annual time series of global VIIRS nighttime lights derived from monthly averages: 2012 to 2019. *Remote Sensing, 13*(5), 922. https://doi.org/10.3390/rs13050922

Falchi, F., et al. (2016). The new world atlas of artificial night sky brightness. *Science Advances, 2*(6), e1600377. https://doi.org/10.1126/sciadv.1600377

Garstang, R. H. (1986). Model for artificial night-sky illumination. *Publications of the Astronomical Society of the Pacific, 98*(601), 364-375. https://doi.org/10.1086/131768

Helmus, J. J., & Collis, S. M. (2016). The Python ARM Radar Toolkit (Py-ART), a library for working with weather radar data in the Python programming language. *Journal of Open Research Software, 4*(1), e25. https://doi.org/10.5334/jors.119

Hollister, J., et al. (2025). *elevatr: Access elevation data from various APIs* (R package). https://doi.org/10.5281/zenodo.8335450

Horton, K. G., et al. (2019). Bright lights in the big cities: Migratory birds' exposure to artificial light. *Frontiers in Ecology and the Environment, 17*(4), 209-214. https://doi.org/10.1002/fee.2029

Horton, K. G., et al. (2023). Artificial light at night is a top predictor of bird migration stopover density. *Nature Communications, 14*, 7446. https://doi.org/10.1038/s41467-023-43046-z

Kyba, C. C. M., et al. (2017). Artificially lit surface of Earth at night increasing in radiance and extent. *Science Advances, 3*(11), e1701528. https://doi.org/10.1126/sciadv.1701528

La Sorte, F. A., et al. (2017). Seasonal associations with urban light pollution for nocturnally migrating bird populations. *Global Change Biology, 23*(11), 4609-4619. https://doi.org/10.1111/gcb.13792

Longcore, T., & Rich, C. (2004). Ecological light pollution. *Frontiers in Ecology and the Environment, 2*(4), 191-198. https://doi.org/10.1890/1540-9295(2004)002[0191:ELP]2.0.CO;2

McLaren, J. D., et al. (2018). Artificial light at night confounds broad-scale habitat use by migrating birds. *Ecology Letters, 21*(3), 356-364. https://doi.org/10.1111/ele.12902

Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika, 37*(1/2), 17-23. https://doi.org/10.2307/2332142

Vol. 2026 (11) 903 – 916
ISSN 2688-3651 [online]

Multi-Resolution Land Characteristics Consortium. (2024). *National Land Cover Database (NLCD)*. U.S. Geological Survey. https://doi.org/10.5066/P94UXNTS

Nastrom, G. D. (1997). Doppler radar spectral width broadening due to beamwidth and wind shear. *Annales Geophysicae, 15*(6), 786-796. https://doi.org/10.1007/s00585-997-0786-7

National Oceanic and Atmospheric Administration. (1991). *NEXRAD Level II base data, station KMUX*. NOAA National Centers for Environmental Information. https://doi.org/10.7289/V5W9574V

R Core Team. (2025). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.R-project.org/

Rousset, F., & Ferdy, J.-B. (2014). Testing environmental and genetic effects in the presence of spatial autocorrelation. *Ecography, 37*(8), 781-790. https://doi.org/10.1111/ecog.00566

Stepanian, P. M., et al. (2016). Dual-polarization radar products for biological applications. *Ecosphere, 7*(11), e01539. https://doi.org/10.1002/ecs2.1539

Van Doren, B. M., & Horton, K. G. (2018). A continental system for forecasting bird migration. *Science, 361*(6407), 1115-1118. https://doi.org/10.1126/science.aat7526

Van Doren, B. M., et al. (2017). High-intensity urban light installation dramatically alters nocturnal bird migration. *Proceedings of the National Academy of Sciences, 114*(42), 11175-11180. https://doi.org/10.1073/pnas.1708574114

Van Doren, B. M., et al. (2021). Drivers of fatal bird collisions in an urban center. *Proceedings of the National Academy of Sciences, 118*(24), e2101666118. https://doi.org/10.1073/pnas.2101666118

Wood, S. N. (2017). *Generalized additive models: An introduction with R* (2nd ed.). Chapman & Hall/CRC. https://doi.org/10.1201/9781315370279