

Supplementary Material for *SpectralBridge: An Auditable Software Workflow for Reflectance Comparisons Across NEON, Drones, and Landsat*

Ty Tuff^{1,*}, Aashish Mukund¹, Nayani T. Ilangakoon², Erick Verleye²
Megan E. Cattau⁴, Janushi Shastri^{2,3}, Pavan Appanna^{2,3}, Tyler L. McIntosh²
Matthew E. Bitters², Carol A. Wessman⁵, Jennifer K. Balch^{1,6}, and Cibebe Amaral¹

Revised 12 October 2026; source snapshots identified in text

¹Earth Lab and the Environmental Data Science Innovation & Impact Lab (ESIIL), Cooperative Institute for Research in Environmental Sciences (CIRES), University of Colorado Boulder, Boulder, Colorado, USA

²Earth Lab, Cooperative Institute for Research in Environmental Sciences (CIRES), University of Colorado Boulder, Boulder, Colorado, USA

³Department of Computer Science, University of Colorado Boulder, Boulder, Colorado, USA

⁴Human-Environment Systems, School of the Environment, Boise State University, Boise, Idaho, USA

⁵Department of Ecology and Evolutionary Biology, University of Colorado Boulder, Boulder, Colorado, USA

⁶Department of Geography, University of Colorado Boulder, Boulder, Colorado, USA

*Corresponding author: ty.tuff@colorado.edu

ORCID iDs: Ty Tuff, [0000-0001-5249-5197](#); Nayani T. Ilangakoon, [0000-0002-6607-0605](#); Erick Verleye, [0009-0009-8854-0127](#); Megan E. Cattau, [0000-0003-2164-3809](#); Tyler L. McIntosh, [0000-0002-6383-1755](#); Carol A. Wessman, [0000-0002-7325-0153](#); Jennifer K. Balch, [0000-0002-3983-7970](#); Cibebe Amaral, [0000-0001-7597-2427](#)

1 Purpose and relationship to the main text

The main paper argues that comparable reflectance has to be constructed and that researchers should be able to examine how that construction was made. This supplement provides the detail needed to assess that claim: processing contracts and transformations, the intended drone artifact sequence, the first 43-flight production attempt, the streaming bulk workflow inspected in September 2026, historical QA records, and the supplied 122-flightline analysis. Development history, neighboring-software comparisons, and the prospective QA and resource protocol are included here so that the main paper can stay focused on the scientific problem.

The organization follows the questions a reader is likely to ask. What is the software intended to produce? Which numerical transformations does it apply? How are completion and QA judged? What happened in controlled tests and real-data runs? Which conclusions can be drawn from the collection analysis, and what remains to be tested? Detailed provenance and future protocols follow the results they qualify.

The historical NEON and intended drone descriptions retain the 4 September audit at [f7e7ddd](#); the bulk and release-preparation update refers to [c86da6a](#), inspected 9 September (Earth Lab, [2026ak](#);

Earth Lab, 2026l). R10C was generated at d88e8d0 and archived at f7e7ddd. The supplied bulk results have a separate analysis run identifier, not yet linked to an executed source revision by a matching manifest. The exact revision and resolved configuration used for the 43-flight drone production attempt were not included in its supplied inventory, and no corrective revision is assigned here (SpectralBridge project team, 2026). Source inspection, historical execution, production inventory, and this revision’s reanalysis of exported bulk tables are kept distinct.

The software-lineage section uses the same revision boundaries. It draws on explicit notices in pinned source files, cited repository history, and the later repository audit summarized for this revision. No standalone HYTOOLS_PROVENANCE.md or provenance manifest was supplied with the manuscript workspace. We therefore do not describe the later attribution work as documentation that existed in historical releases. Where file-level evidence is incomplete, the tables say so rather than declaring independence.

Four terms keep the claims consistent. *Implemented* means that behavior is present in source code. *Tested* means that a controlled fixture exercises a defined contract. *Demonstrated* refers to artifacts and behavior recorded in a real-data run. *Scientifically validated* requires independent observations or a justified external reference. Thus, the drone campaign demonstrates discovery, staging, and persistence of intermediate products, but not completion of the full drone route.

2 Software scope and supported paths

SpectralBridge has two observation-level paths and one collection-level path. NEON processing produces source-product and geometry-normalized hyperspectral representations, configured target-band products, analysis tables, and QA. Drone processing ingests supported local observations while preserving their identity, correction availability, and provenance (Earth Lab, 2026am; Earth Lab, 2026m). The bulk path then catalogs eligible target products and reads aligned ENVI windows; merged Parquet remains available for compatibility. Bulk analysis starts from completed products and does not rerun geometry normalization or target-band weighting (Earth Lab, 2026a).

The project-level relationship Drone → NEON → Landsat is conceptual. NEON supplies a dense reference spectrum from which configured MicaSense- and Landsat-like band representations can be constructed. The package does not implement a direct, universal drone-to-Landsat conversion, and target-band weighting does not alter spatial support (National Ecological Observatory Network, 2026; Earth Lab, 2026am; Earth Lab, 2026s).

Each observing system brings a different strength. Imaging spectroscopy preserves contiguous spectral information (Goetz et al., 1985); the NEON Airborne Observation Platform applies it to repeatable ecological observation (Kampe et al., 2010); Landsat provides a long calibrated record at moderate resolution (Wulder et al., 2022); and compact UAV instruments provide flexible, very-high-resolution sampling while demanding careful measurement design and calibration (Aasen et al., 2018; Mamaghani and Salvaggio, 2019). SpectralBridge addresses the spectral and workflow parts of relating these systems, not the entire comparison problem.

3 Flightline processing contracts

3.1 Canonical NEON stage sequence

Table 1: Canonical NEON processing stages and their scientific interpretation (Earth Lab, 2026i; Earth Lab, 2026e; Earth Lab, 2026s).

Stage	Principal input	Principal output	Scientific purpose
Source acquisition	Flightline identifier and NEON product request	Source HDF5	Preserve the acquired observation and identity.
Source-product export	NEON HDF5	ENVI image and header	Expose reflectance and map information in a standard raster representation without implying unprocessed sensor data.
Parameterization	Reflectance and ancillary geometry	JSON model reference and geometry summaries	Record the retained BRDF and geometry information; scene SCS+C terms are not persisted.
Correction	Source-product ENVI and geometry/model information	Combined topographic-and-BRDF-corrected ENVI	Normalize modeled terrain and directional effects.
Configured target-band weighting	Geometry-normalized hyperspectral cube and configured weight arrays	Target-band representations of NEON pixels	Apply target-band weights and any separately recorded brightness adjustment without changing spatial support.
Tabulation	ENVI products	Per-product Parquet files and merged flightline table	Make pixel-level values and metadata available to statistical workflows.
Quality assurance	Stage artifacts and configured checks	JSON, metrics, HTML, visual panels, and optional combined PDF	Record evaluated rules, warnings, failures, and unevaluated questions.

Canonical paths and per-flightline directories isolate products during parallel runs and provide stable locations for restart checks (Earth Lab, 2026i; Earth Lab, 2026e; Earth Lab, 2026s). The command-line interface and worked invocation are maintained in the versioned project documentation; worker count and storage configuration are execution choices rather than scientific defaults.

3.2 Optional across-track half-flight processing

Full-flightline processing remains the default. When `split_across_track=True`, SpectralBridge downloads or reuses one source HDF5 file, divides the across-track columns into left and right windows, and writes each window to a separately suffixed product tree. If the sample count is odd, the right window receives the additional column. The two windows share source identity but are independent pipeline units for output paths, QA, and restart (Earth Lab, 2026o; Earth Lab, 2026i).

This mode changes more than memory layout. Topographic and BRDF estimation for each run sees only the pixels in that window, so half-flight products are not assumed to be numerically identical to the corresponding subset of a full-flightline result. Bulk-run records must therefore distinguish

full, left-half, and right-half processing. The mode provides an operational option for large inputs, but its effect on peak memory, run time, correction estimates, and seam behavior must be measured rather than inferred (Earth Lab, 2026o; Earth Lab, 2026ao).

4 Scientific transformations

4.1 Configured target-band weighting

Nominally similar bands can have different relative spectral responses and therefore produce different reflectance values over the same target (Chander et al., 2013; Roy et al., 2016; Claverie, 2023). SpectralBridge addresses this spectral part of the comparison by applying sampled target-band weights to the NEON spectrum. The result is a configured target-band representation of the NEON observation, not an independently observed target-sensor measurement. Some source code and historical QA records call this stage *spectral convolution*; the manuscript uses *target-band weighting* when discussing its scientific meaning.

For reflectance samples ρ_i and response weights $s_{b,i}$ already sampled on the same source-band grid, the implemented target-band quantity is

$$\hat{\rho}_b = \frac{\sum_i \rho_i s_{b,i}}{\sum_i s_{b,i}}. \quad (1)$$

The implementation normalizes each response by its discrete sum and applies a floating-point matrix product (Earth Lab, 2026aa). The wavelength vector is checked for dimensional alignment but is not used for wavelength-interval quadrature. The routine therefore assumes that the supplied sampled weights already encode the intended wavelength sampling and does not independently establish their units, ordering, spacing, provenance, or spectral overlap.

Configured outputs include representations labeled for Landsat 5 TM, Landsat 7 ETM+, Landsat 8 OLI, Landsat 9 OLI-2, and MicaSense, together with downstream affine coefficient sets for selected MicaSense-like and Landsat-like variables. Packaged center and full-width-at-half-maximum values are stored in the repository’s sensor-parameter JSON (Earth Lab, 2026q). These parameters can define Gaussian approximations but are not themselves measured relative spectral response curves; the Landsat 8 Operational Land Imager, for example, has measured bandwise relative spectral responses that determine its actual wavelength weighting (Barsi et al., 2014). More generally, quantitative UAV multispectral use depends on documented geometric and radiometric calibration, and camera-specific characterization can reveal departures from factory assumptions (Aasen et al., 2018; Mamaghani and Salvaggio, 2019; Cao et al., 2019).

Some configured Landsat-like products receive an empirical percentage adjustment after Equation 1. The inspected coefficient files describe adjustments derived from Table Mountain HLS median differences and warn that their band indices are wavelength-aligned rather than interchangeable with native sensor numbering (Earth Lab, 2026r). Product headers and QA artifacts record whether the adjustment was applied, but the purely weighted representation is not always retained as a separate canonical raster. Target-band weighting and brightness adjustment are therefore distinct measurement states even where the artifact chain combines them, and transferability of the adjustment beyond its derivation context has not been established. Evaluations of bandpass-adjustment techniques show that simple linear, spectral-library, adaptive, and nonlinear methods

have different data requirements and performance; no one fitted form is a complete definition of sensor harmonization (Claverie, 2023). HLS version 2.0 further demonstrates that operational harmonization combines atmospheric correction, cloud and shadow handling, directional normalization, bandpass adjustment, and common gridding (Ju et al., 2025).

Five controlled source/target band-count variations agreed with an independently implemented normalized-weight calculation to an absolute tolerance of 2×10^{-7} (Earth Lab, 2026ar). This tests the arithmetic contract. It does not validate weight provenance, usable spectral support, or agreement with measured satellite observations.

Declared poor-quality source bands are retained visibly in the R10C evidence, but the archived reports do not establish whether they received nonzero weight in each target band (Earth Lab, 2026w; Earth Lab, 2026y). Before scientific interpretation, the planned study must report usable source-wavelength coverage for every target-band representation, identify overlap with declared poor-quality bands, and define the status assigned when usable support falls below the recorded thresholds. Until then, an arithmetic PASS for spectral weighting should not be read as evidence that every target band has adequate scientific support.

4.2 Topographic and BRDF normalization

SpectralBridge applies sun-canopy-sensor plus C (SCS+C) topographic normalization followed by kernel-driven bidirectional reflectance distribution function (BRDF) normalization. SCS+C is attributed to its forest topographic-correction source (Soenen et al., 2005); RossThick and LiDenseR belong to volume- and geometric-scattering model lineages developed and operationalized across several studies (Ross, 1981; Roujean et al., 1992; Li and Strahler, 1992; Wanner et al., 1995; Schaaf et al., 2002). The source implementation has a separate provenance: `calc_cosine_i`, `calc_volume_kernel`, and `calc_geom_kernel`, together with portions of the surrounding correction module, explicitly identify adaptation from the GPLv3-licensed HyTools project (Chlus et al., 2025; Earth Lab, 2026z). In the default mode, one SCS+C term per band is estimated from pixels passing the source no-data and finite-value masks, using streamed sufficient statistics across the flightline. Correction is applied in full-width row strips. Scene-level BRDF coefficients are then applied using those kernels, scene-mean solar zenith, local view geometry, and a conventional normalization target of 45-degree solar zenith, nadir view, and zero relative azimuth; this is not the actual acquisition geometry of Landsat or NEON (Queally et al., 2022; Earth Lab, 2026ao). Non-finite or unsuitable correction ratios are replaced by neutral factors rather than allowed to create extreme values (Earth Lab, 2026z). The fraction of pixels and bands receiving neutral fallbacks should be retained as a QA diagnostic because numerical stability does not imply that a correction was applied.

An optional tiled mode fits independent topographic terms in non-overlapping 100×100 -pixel tiles. This reduces fitting scope but creates a specific QA question: whether grid-aligned seams appear in the normalized product. The canonical normalized cube combines topographic and BRDF operations. Separate topographic-only and BRDF-only cubes are not retained, so attribution of their individual effects requires additional runs or outputs (Earth Lab, 2026ao; Earth Lab, 2026s).

These empirical and semi-empirical models alter observed relationships with terrain and acquisition geometry. Independent reference observations are needed to determine whether the resulting reflectance is more accurate for a particular scientific use, especially when vegetation composition covaries with slope or aspect (Soenen et al., 2005; Queally et al., 2022).

4.3 Spatial support

Spectral weighting changes band space, not pixel footprint. A target-band representation of a NEON pixel and an actual Landsat footprint remain different mixtures even when they use similar band labels; a sensor’s point-spread function can also extend the contributing area beyond a nominal pixel boundary (Teillet et al., 1997; Woodcock and Strahler, 1987; Huang et al., 2002). Spatial comparison must distinguish grain, support, extent, and the scale of the ecological process (Dungan et al., 2002). SpectralBridge retains geospatial rasters and polygon-extraction products so users can define co-registration and aggregation explicitly (Earth Lab, 2026s; Earth Lab, 2026m). The aggregation rule determines which fine-resolution observations contribute, how incomplete overlap is handled, and whether the resulting quantity represents area, cover, or another ecological target.

5 Drone processing pipeline

The drone pipeline is intended to prepare the independently acquired fine-resolution side of the comparison. It recursively discovers local HDF5 inputs and supports configured local TIFF inputs, preserves provenance from original filenames, and stages a working HDF5 representation. The intended downstream sequence is working HDF5, native-drone ENVI representation, geometry normalization where supported, geometry-normalized native-drone representation, optional application of reviewed affine relationships to construct configured Landsat-like variables, full-image or polygon extraction, Parquet, and QA (Earth Lab, 2026m; Earth Lab, 2026aj). This use case follows a broader ecological need to connect reproducible UAS observations with NEON’s continental sampling framework (Koontz et al., 2022).

Topographic and BRDF normalization require the corresponding terrain, view, and solar geometry. Radiometric panel methods, empirical-line calibration, camera characterization, and downwelling-sensor geometry can each influence retrieved UAV reflectance and must be recorded as part of the measurement definition (Wang and Myint, 2015; Cao et al., 2019; Schneider-Zapp et al., 2019; Origo et al., 2025). The intended QA record distinguishes a correction that ran from one that was skipped because geometry was unavailable or that used a fallback after a quality check. Users can request all available pixels or restrict extraction to intersecting polygons. Each discovered flight receives a separate output directory with drone-native identity. Which artifacts are required depends on the resolved configuration; completion must therefore be evaluated against that configuration rather than against a universal list or process return code.

The optional affine relationships require especially careful language. The documented post-normalization equation is

$$L = a + bM, \tag{2}$$

where M is a corrected native-drone variable and L is a configured Landsat-like variable. The release-candidate documentation specifies wavelength-aware source-to-target mapping, a versioned reviewed coefficient registry, and a fixed production policy using the site-balanced relationship. The drone route neither estimates these coefficients at runtime nor performs spectral convolution; convolution remains part of the NEON hyperspectral path. If an approved production registry is unavailable, the documented default is to stop with an actionable error rather than derive a registry from summary statistics (Earth Lab, 2026aj).

The registered relationships were derived from same-source synthetic NEON representations. Ap-

plying Equation 2 is not empirical calibration against independent Landsat observations and does not create Landsat spatial support. No configured Landsat-like variables are inferred from the 43-flight production attempt because the expected downstream raster and tabular products were absent (Earth Lab, 2026d; SpectralBridge project team, 2026).

Optional external QA can search for an overlapping Landsat Collection 2 Level 2 scene, use a supplied analysis-ready raster, or reuse a cached observation manifest. The documented comparison places products on a common Landsat grid and can include a pre-existing NEON-convolved product. Reproducible use therefore requires the Landsat product identifier, acquisition time, masks, scale and offset, grid and resampling rules, overlap definition, and local cache or query record (U.S. Geological Survey, 2026). This capability is not evidence that such a comparison ran for the 43-flight campaign or that agreement was scientifically validated. Surface-reflectance validation requires independent, fully characterized reference measurements with uncertainty and traceability, together with spatial and temporal matching rules (Origo et al., 2020; Origo et al., 2025).

Table 2 states the configuration-dependent completion contract used to interpret a future campaign replay. It is the manuscript’s scientific acceptance protocol, not a claim that the production orchestrator already enforces every gate. A stage marked as not requested or unsupported must carry an explicit recorded disposition; silently omitting it is not equivalent to satisfying it.

The drone and NEON pipelines therefore meet scientifically rather than through a hidden conversion. The intended drone path preserves and prepares the observed high-resolution measurement; the NEON path constructs MicaSense- and Landsat-like band representations; and the researcher must still define temporal matching, geometric overlap, spatial aggregation, and the ecological sampling unit. The first real campaign does not yet provide completed products for that comparison.

Table 2: Configuration-dependent artifact contract for the drone route. A flight is complete only when every artifact required by its resolved configuration satisfies the stated evidence check.

State	Required when	Minimum completion evidence	Non-completing disposition
Working HDF5	A discovered observation is staged	File opens; required observation arrays and source identity are readable; provenance points to the original local input.	Missing, unreadable, or identity mismatch.
Native ENVI representation	Raster representation is requested or is a prerequisite of a requested downstream stage	Header/image pair opens together; dimensions, bands, coordinate reference information, and no-data metadata pass structural checks.	Missing member of the pair, unreadable data, or inconsistent dimensions.
Geometry-normalized native-drone representation	Geometry normalization is requested and supported by available ancillary data	Output raster opens; the correction record identifies whether normalization ran, its inputs, and any fallback fractions.	Unrecorded omission. An explicit unsupported or skipped status may be valid for a configuration that does not require normalized output, but is not evidence of normalization.
Configured Landsat-like representation	Application of a reviewed affine relationship is requested	Output opens; registry version, coefficient and estimator identifiers, wavelength mapping, source and target variables, a and b , order of application, and units or stored-value scale are recorded.	Missing output or registry evidence; a native-drone product does not substitute for this state.
Extracted Parquet	Full-image or polygon extraction is requested	File and schema are readable; flight identity, variables, row count, extraction mode, and polygon identity where applicable are recorded. A zero-row polygon result requires an explicit overlap disposition.	Missing or unreadable table, or an unexplained zero-row result.
QA evidence	A stage or complete flight is evaluated	Stage JSON parses; evaluated rules identify their artifact; a combined report preserves pass, warning, failure, and not-evaluated states.	A QA rendering alone cannot substitute for a missing scientific product.

6 File contracts, provenance, and restart behavior

The phrase “outputs are the API” means that stages communicate through named files, not only through values held in memory (Earth Lab, 2026p). Those files preserve the scientific sequence: source observation, working or source-product representation, geometry/model summary, geometry-normalized representation, configured target-band representation, extracted table, and QA record. The resolved configuration determines which of them are required. This design is consistent with FAIR workflow and research-object practices that package data, software, parameters, provenance, and results as one inspectable record (Goble et al., 2020; Soiland-Reyes et al., 2022).

This design helps localize a problem, although it does not diagnose the cause automatically. A source-product cube can show whether an unexpected pattern was already present; the geometry-normalized cube can show whether it appeared during normalization; the target-band product can reveal a change during band construction; and the Parquet output can reveal a change during extraction or merging. Some boundaries remain coarse: topographic and BRDF effects share one cube, and weighting and brightness adjustment can share one target-product stage. The files can therefore identify the stage at which a pattern appeared without always isolating the operation within that stage (Earth Lab, 2026p; Earth Lab, 2026i; Earth Lab, 2026s).

On rerun, a stage can reuse artifacts that pass its format checks and reconstruct outputs that are absent or rejected (Earth Lab, 2026f; Earth Lab, 2026aq). JSON must parse, ENVI image/header pairs must pass internal readability checks, and Parquet outputs receive schema-read checks. QA reuse is stricter: the public contract requires matching schema, parameters, thresholds, software version, and input/output fingerprints (Earth Lab, 2026af). Format checks, QA reuse, and completion are therefore separate decisions. Reproducible reuse requires retaining or recording:

- the full software commit and installed package versions;
- the command and resolved configuration;
- source identifiers and, where possible, checksums;
- packaged response-definition versions or exact response arrays;
- the expected artifacts implied by the resolved configuration and their stage dependencies;
- produced artifacts and their checksums; and
- the execution environment and resource configuration.

Two important gaps remain: the scene-wide SCS+C coefficients applied to the geometry-normalized cube and the exact per-output-band weights are not fully serialized in the canonical record (Earth Lab, 2026ao; Earth Lab, 2026aa; Earth Lab, 2026s). Both are needed to reconstruct every numerical transformation exactly.

7 Multi-flightline bulk analysis

The bulk implementation inspected on 9 September 2026 replaces the earlier Parquet-centered description with direct raster streaming. Analysis eligibility requires the requested matched products,

not a complete upstream archive or an upstream QA pass. Catalogs retain duplicate identities and ineligible inputs. Merged Parquet remains supported, but the direct route does not populate a pixel-scale observation cache (Earth Lab, 2026a; Earth Lab, 2026c).

The pipeline provides four related evidence products:

Table 3: Updated bulk-analysis products and their interpretation (Earth Lab, 2026a; Earth Lab, 2026c).

Product	Function	Scientific interpretation
Catalog and census	Inventory identities, schemas, row counts, sites, dates, and duplicate or rejected inputs.	Defines the analyzed collection and exclusions before summaries are interpreted.
Compact statistics	Reduce aligned raster windows to mergeable per-flightline moments.	Supports grouped fitting without storing another pixel-scale table.
Same-source relationship summaries	Fit pooled, per-flightline, per-site, flightline-balanced, and site-balanced relationships between wavelength-matched target products.	Describes relationships among products derived from the same geometry-normalized NEON observations; it is not empirical calibration between independently observed sensors.
Leave-one-site-out evaluation	Refit without one site and report held-out-site diagnostics.	Describes out-of-site stability of a same-source synthetic relationship; it is not independent cross-sensor validation.

Pixels are nested within flightlines, and flightlines are nested within sites. Pooled pixel regressions are therefore descriptive and must not present pixel count as the number of independent landscapes (Hurlbert, 1984). The three reported estimands differ only in group contribution: pixel-pooled fitting gives each selected pixel equal weight, flightline-balanced fitting gives each flightline equal total weight, and site-balanced fitting gives each site equal total weight. The exported **weighting** descriptions and replicate counts were checked against these definitions. The balanced summaries are sensitivity analyses, not independent replications or a hierarchical uncertainty model. Validation of spatially and hierarchically structured data should preserve that dependence in its training and test partitions (Roberts et al., 2017). Leave-one-site-out evaluation fits on two observed sites and evaluates the third; with only three sites, its three folds describe those particular site/acquisition combinations rather than a population of possible sites or a general model area of applicability (Meyer and Pebesma, 2021). All summaries inherit the limitations of the underlying products: paired synthetic bands share one corrected NEON source and do not constitute independently acquired drone or Landsat observations (Earth Lab, 2026b; Earth Lab, 2026d).

For each chunk, counts, means, centered sums of squares, and cross-products are accumulated in numerically stable bivariate moments. Merging produces flightline and collection records; inverse merging removes a held-out site’s contribution to training. A second bounded pass calculates errors that are not all recoverable from moments alone (Earth Lab, 2026a; Earth Lab, 2026c). The bulk manifest and per-flightline checkpoint records should accompany source hashes and the resolved configuration; restartability is not a substitute for exact-run provenance.

The direct reader checks dimensions, transform, and coordinate reference system and applies sensor-wide finite/no-data rules together with the configured minimum-value threshold. It reads stored raster values without itself establishing their physical scale (Earth Lab, 2026a). Concurrency, window size, and storage throughput therefore matter both to reproducibility and to the cost

of repeated reads. Explicit dataset materialization and the compatibility input path should be benchmarked separately, not described as mandatory costs of streaming.

The results interpretation layer operates on the manifest and compact model/group tables without reopening the source rasters (Earth Lab, 2026g). This manuscript’s focused examination similarly uses the supplied CSV exports. Neither workflow supplies a complete upstream QA census merely by accepting inputs into a relationship analysis.

7.1 Additional spectral-library reporting route

The release-candidate interface also exposes spectral-library preflight and reporting as first-class routes. A preflight inspects a supplied Parquet library’s schema, group counts, and anticipated rendering work before analysis. The reporting route reads that library in place and can produce species summaries, quantiles, spectral ensembles, hierarchical variability summaries, traceable extreme-spectrum diagnostics, and optional multipage reports without making a second row-level copy of the library (Earth Lab, 2026aj). This route is relevant to the package’s broader analysis design and is included in the documented exact-artifact smoke. No executed spectral-library output was supplied with the manuscript evidence bundle, however, so this manuscript treats it as implemented and bounded-test functionality rather than a demonstrated scientific result.

8 Quality-assurance model

Every run raises four questions. Did the code execute? Were the products required by the resolved configuration written and readable? What did QA report about those products? Do independent observations support their scientific interpretation? We call these process, artifact, QA, and validation status. Remote-sensing frameworks likewise treat validation as evidence assembled from representative comparisons and characterized reference measurements, not as a single binary label (Morissette et al., 2006; Origo et al., 2020). The 43-flight campaign is a concrete example: it reached intermediate and QA states but did not satisfy its downstream artifact contract (SpectralBridge project team, 2026).

SpectralBridge separately uses PASS, WARN, and FAIL within evaluated QA records. Table 4 defines that narrower vocabulary (Earth Lab, 2026s; Earth Lab, 2026w).

Status	Meaning
PASS	The evaluated artifact or numerical quantity met its configured rule.
WARN	The evaluated quantity or known condition requires interpretation; the status does not independently establish full pipeline completion.
FAIL	A required artifact or evaluated rule prevents the intended interpretation.
Not evaluated	Available artifacts do not answer the stated question; absence of evidence is not converted into a pass.

Table 4: QA status semantics. A status applies to its named diagnostic, not to process execution, complete fulfillment of a resolved artifact contract, or universal scientific validity (Earth Lab, 2026s; Earth Lab, 2026w).

QA artifacts include visual panels, machine-readable JSON and metric records, stage-level HTML and JSON reports, and an optional combined PDF (Earth Lab, 2026s). These products record observed values, thresholds, provenance fields, and reasons that a diagnostic could not be evaluated. Several recorded thresholds are provisional workflow diagnostics rather than externally validated scientific decision limits.

For a collection QA census, the emitted combined status is treated as authoritative rather than reconstructed from selected metrics. PASS, WARN, and FAIL counts are mutually exclusive. A missing or unreadable report remains a separate category in the attempted-flightline denominator. Stage-specific warning counts can overlap. Threshold sets and their provisional flags must be archived so that changes in rules are not mistaken for changes in data quality (Earth Lab, 2026w; Earth Lab, 2026x). The supplied bulk fit status ok is not this combined QA status.

Recorded R10C diagnostic	Warn threshold	Fail threshold
Within-footprint valid fraction	below 0.90	below 0.70
Negative-reflectance fraction	above 0.01	above 0.05
Usable-band fraction above 1.2	above 0.01	above 0.05
Absolute correction difference, 99th percentile	above 0.20	above 0.50
Chunk-seam score	above 1.5	above 2.5
Spectral-response coverage	below 0.98	below 0.90

Table 5: Provisional R10C screening rules, not independently validated ecological decision limits (Earth Lab, 2026x). The supplied bulk tables do not establish that these checks were run on every input.

8.1 Validation tiers and installed-artifact checks

The inspected documentation separates validation into two tiers. Tier A builds the wheel and source distribution, installs the exact artifact into a clean environment outside the source checkout, blocks network access, and runs small deterministic smoke fixtures through the NEON, drone, and bulk entry points. This bounded strategy addresses a common scientific-software problem: many computations lack a complete test oracle even when their software contracts can be tested (Kanewala and Bieman, 2014). The public validation page reports a separate offline campaign with 40 of 40 passing variations: five variations for each of eight modules. It records revision `bebf707`, marks the working tree as dirty, and gives 14 August 2026 as its last update. This supports the recorded module checks but does not uniquely identify a clean source state (Earth Lab, 2026al). The 4 September audit reports successful clean wheel checks on local Python 3.10, 3.11, 3.12, and 3.14 interpreters and a source-distribution check on Python 3.10. The tag-release workflow has the narrower supported gate of the exact wheel on Linux with Python 3.10–3.12 and the source distribution on Python 3.10 (Earth Lab, 2026ap; Earth Lab, 2026v; Earth Lab, 2026k).

The NEON smoke fixture uses an $8 \times 8 \times 32$ synthetic cube and exercises geometry normalization, configured target products, Parquet conversion and merging, QA, and restart. Drone smoke tests exercise full-image and polygon routes, and bulk smoke tests exercise census, relationship fitting, coefficient export, leave-one-site-out evaluation, optional materialization, and restart. The release-candidate README also identifies spectral-library reporting as part of the bounded exact-artifact smoke (Earth Lab, 2026aj). These tests provide evidence for packaging, stage connectivity, artifact readability, and bounded restart behavior. They do not test production scale or scientific accuracy.

Tier B consists of selected real-flightline records such as R10C and is the appropriate tier for operational scaling and data-dependent QA evidence (Earth Lab, 2026t; Earth Lab, 2026v).

The following sections use those distinctions consistently. The drone campaign examines campaign-scale orchestration and completion; R10C records a completed flightline with warnings; WREF preserves a more limited visual QA record; and the bulk exports support collection-level analysis of target-band products that had already been made.

9 First production drone campaign attempt

The supplied production inventory covers 43 historical flight packages: 17 from 2023 and 26 from 2024. Source discovery and staging produced 43 working HDF5 observations and approximately 7.15 GB of output. The inventory does not identify the exact production commit, resolved configuration, manifest, file checksums, per-stage timings, or a corrective revision; we leave those fields unspecified rather than inferring them (SpectralBridge project team, 2026).

Table 6: Observed inventory after the first 43-flight drone production attempt. Counts describe files found by post-run inspection, not products inferred from run status.

Artifact class	Observed count	Interpretation
Working HDF5	43	Persisted intermediate observation state
JSON	215	Machine-readable run or QA records; detailed roles not supplied
QA PDF	43	Visual QA artifacts found
QA PNG	43	Visual QA artifacts found
ENVI header	0	Expected downstream representation absent
ENVI image	0	Expected downstream representation absent
Parquet	0	Expected extracted analysis product absent

The orchestration summary classified all 43 flights as processed and none as failed. Yet the public output contract identifies a readable merged Parquet table with matching QA, together with configured sensor products when requested, as the primary completion signals (Earth Lab, 2026af). Those products were absent. The run therefore shows that the HDF5 route can discover, stage, and preserve intermediate states across a substantial campaign, but it does not demonstrate corrected drone outputs, configured Landsat-like variables, completed extraction, or drone/Landsat agreement.

We treat the mismatch as an orchestration and completion-accounting defect. Named artifact classes made the incomplete state visible, but the orchestration did not use that same record correctly when assigning completion. No root cause is assigned. The 43 working HDF5 files may support restart after correction, but their practical reuse and downstream completion must be shown by replay.

A publication replay should retain both records: the original incomplete attempt and the corrected run. The new record should identify its software revision, resolved configuration, artifact contract, input and output checksums, per-stage status, any reused HDF5 files, and the rule that prevents completion when required products are absent. The original attempt should not be overwritten or counted retroactively as complete.

10 R10C real-flightline demonstration

The checked-in example processes the flightline

```
NEON_D10_R10C_DP1_L002-1_20210915_directional_reflectance.
```

The workflow produced all of its declared artifacts, but the combined interpretation remained WARN (Earth Lab, 2026w; Earth Lab, 2026y). The 2,532,249,598-byte source HDF5 decoded to 1,115 lines by 5,351 samples by 426 bands. The run wrote approximately 21 GB of local output, while the retained validation-report bundle occupied 4.6 MB.

The archived reports identify package version 2.2.0 and the following generating revision:

```
d88e8d0947718d1e33b534080c224b703987e536.
```

The inspected repository commit, `f7e7ddd`, archives those reports but did not generate them. R10C also predates the 9 September release-preparation snapshot. We therefore use it as a record of that historical run, not as an exact-artifact replay (Earth Lab, 2026t).

The combined report uses schema version 1.3 and plot-contract version 1.1. Its declared `pipeline_evolution.png` uses fixed display ranges for cross-report comparability; those limits are explicitly display-only, and values outside them remain in the machine-readable metrics. The corresponding figure in the main manuscript summarizes the combined JSON rather than reproducing that original report plot; the tables below retain the detailed numerical record without duplicating that figure (Earth Lab, 2026w; Earth Lab, 2026y).

Stage	Status
Source acquisition	PASS
Input reflectance	WARN
Correction parameters	WARN
BRDF and topographic correction	WARN
Spectral convolution (target-band weighting)	PASS
Analysis tables	PASS

Table 7: Stage statuses in the R10C combined QA record (Earth Lab, 2026w).

Several apparently conflicting values become clear once their denominators are stated. The 0.571 bounding-box footprint fraction describes how much of the rectangular raster is occupied by the swath; the 1.0 within-footprint valid fraction says that all pixels inside that footprint passed the relevant check. Likewise, the provisional threshold in Table 5 applies to the usable-band fraction above 1.2, 1.41814×10^{-5} , rather than the all-evaluated-band fraction, 0.0593153. The report also records zero negative-reflectance fractions for the highlighted input, geometry-normalized, and target-band states. These are descriptive checks, not evidence that the transformations improved agreement with another sensor (Earth Lab, 2026w; Earth Lab, 2026y).

The 68 known poor-quality wavelength bands were classified in `report_only_no_masking` mode, and the report records `data_modified=false`: the bands were labeled but neither masked nor replaced. Four persisted geometry summaries fell outside their expected radian ranges. The walkthrough

Diagnostic	Recorded value
Package-declared poor-quality bands retained	68
Classification behavior for poor-quality bands	Report only; no masking or replacement
Bounding-box footprint fraction	0.571275
Within-footprint valid fraction	1.0
Input median reflectance	0.250400
Input negative-reflectance fraction	0.0
Fraction of evaluated input-band values above 1.2	0.0593153
Fraction of usable input-band values above 1.2	1.41814×10^{-5}
Combined-correction median reflectance	0.250400
Combined-correction negative-reflectance fraction	0.0
Fraction of evaluated corrected values above 1.2	0.0593168
Combined-correction absolute-difference 99th percentile	0.00389395 reflectance
Convolved median reflectance	0.144450
Convolved negative-reflectance fraction	0.0
Fraction of evaluated convolved values above 1.2	1.43012×10^{-5}
Readable Parquet outputs	18

Table 8: Selected machine-readable diagnostics from the R10C combined QA JSON (Earth Lab, 2026w). Values describe this run and are not estimates of performance across sites.

suspects that no-data values affected these summaries, but the archived evidence does not establish the cause or show whether the issue was confined to persisted summaries rather than fitted correction terms. The files were not silently repaired. The combined report summarizes stage status and two pipeline-level findings; the fuller list of unevaluated questions comes from the stage reports and real-data walkthrough. Those questions include separate attribution of topographic and BRDF effects, per-output-band weight support, held-out affine-relationship residuals and cycle consistency, and direct Landsat Collection 2 NBAR comparability (Earth Lab, 2026w; Earth Lab, 2026x; Earth Lab, 2026y).

R10C therefore demonstrates end-to-end artifact completion for its declared historical route, stage-specific QA, and preservation of warnings. It does not measure correction efficacy or cross-sensor agreement (Earth Lab, 2026w; Earth Lab, 2026y). Read beside the drone campaign, it shows why a completed artifact chain answers a different question from a successful process return code.

11 WREF visual QA example

A separate four-page visual QA report for flightline

NEON_D16_WREF_DP1_L040-1_20230625_directional_reflectance.

shows the products in more detail than the combined R10C status summary. It includes every generated ENVI representation, pre- and post-correction diagnostics, wavelength and header checks, target-band locations, filtering summaries, and the Parquet merge inventory. Because the PDF

identifies neither its generating SpectralBridge commit nor its QA schema, we treat it as a visual inspection record rather than a machine-readable collection record.

The complete four-page report is retained unchanged as a supplied source artifact. Here, its correction plots are reflowed for legibility and selected tabular-inventory values are presented as a visual summary. Its product overview preserves the narrow across-track geometry and displays every target representation, but repeated short labels for the final MicaSense products are ambiguous; future reports should include unique product identifiers or output paths.

The correction diagnostics in Figure 1 show nearly overlapping aggregate pre- and post-correction histograms and wavelength-dependent changes that are small in median but heterogeneous across pixels. The same page also displays mean slope, aspect, sensor-zenith, and sensor-azimuth values far outside physically meaningful radian ranges. The visual report therefore reinforces the need to retain the geometry warning; it does not establish whether the out-of-range values affected correction fitting or only the displayed summaries.

The spectral page, retained in the supplied source report but not reproduced here, confirms a monotonic, finite 426-band wavelength vector spanning approximately 384–2512 nm and reports that the ENVI header lacks band names. It also contrasts full-cube with filtered-table summaries that use different observational supports and therefore should not be compared as if they described the same statistical population.


The tabular inventory in Figure 2 reports 18 Parquet files totaling 32.9 MB and a merged table containing 151 rows and 964 columns. It also exposes an internal reporting inconsistency: the merge-status panel counts 16 resampled inputs, while the quality-check panel states that no resampled products were found. This disagreement should be resolved in the report generator before the visual page is used for automated cohort interpretation.

Small aggregate changes do not settle geometry QA


WREF / D16 · L040 · 25 June 2023 · original diagnostic plots

The wavelength-change panel excludes —9999 nodata values.


A / REFLECTANCE DISTRIBUTION




B / CHANGE BY WAVELENGTH



C / TOPOGRAPHIC GEOMETRY



D / VIEW AND ILLUMINATION




READ THE LOWER PANELS WITH CAUTION

Persisted geometry means are outside expected radian ranges.
The report does not establish whether correction fits were affected.

Figure 1: Original WREF diagnostic plots from report page 2, reflowed with new panel headings and a reading guide; plotted data are unchanged. (A) Pre- and post-correction distributions nearly overlap. (B) Wavelength-resolved changes retain pixel-level spread; IQR denotes the interquartile range. (C, D) Persisted geometry means lie outside expected radian ranges. These summaries require investigation and do not establish whether correction fits were affected.

A readable table, an unresolved report conflict


WREF / D16 · L040 · 25 June 2023 · selected values from report page 4



A / WHAT THE TABLE CONTAINS



B / WHAT THE REPORT CANNOT RECONCILE



Both messages are retained. This figure summarizes the report; it does not repair it.

Figure 2: Selected WREF inventory values from report page 4. (A) Column groups sum to the reported merged total of 964; strip widths are proportional to counts. (B) The merge-status and quality-check panels give contradictory resampled-product statements, retained here rather than reconciled. Inventory totals and merge-input counts describe different scopes. The displayed values were checked against the checksum-verified source PDF.

12 Bulk analysis: collection and variable definitions

The preceding records ask whether requested artifacts existed and what their QA reported. The bulk analysis begins later in the chain. It starts with eligible target-band products and asks how their same-source relationships vary across the accepted collection.

The available bundle contains a visual report and CSV exports. Table 14 identifies the report and the five primary numerical tables used here. They share the following analysis run identifier:

986e495c2c8e9dda579528b41df44cb24f670055558c85f54941341403a48489

This identifies the analysis, not an executed Git revision. No matching run manifest was supplied (*SpectralBridge bulk evidence bundle* 2026).

12.1 Collection accounting and acquisition coverage

The tables contain 122 distinct flightline identifiers and 18 pair-band combinations per flightline, giving 2,196 local fits. They also contain 54 candidate coefficient records—18 relationships under three weightings—plus 54 site fits and 54 held-out-site evaluations. Every exported fit status is `ok`, but that label says nothing about upstream QA or unrepresented attempts. Because the selection rule, attempted-input denominator, rejected-input inventory, and processing-partition metadata are unavailable, these counts describe accepted records rather than a completion rate or a completed version of the earlier 100-flightline protocol.

Summing the largest exported sample count for each flightline gives 899,692,166 selected observations. For each pair-band, the flightline counts sum to the pooled count, which is either 899,686,966 or 899,692,166. We retain the 5,200-observation difference because the tables do not reveal which masks or input conditions caused it. These observations are not independent ecological replicates and should not be multiplied by the 18 relationships. The repository’s description of the staged archive and checkpoints is an operator-reported production record, not a rerun performed for this paper (Earth Lab, 2026u).

Table 9: Acquisition years parsed from distinct exported flightline identifiers. Dates require confirmation against source metadata and the run manifest.

Site	2020	2023	Total
NIWO	23	9	32
WREF	0	40	40
YELL	0	50	50
Total	23	99	122

Site and acquisition year are unevenly represented (Table 9). Holding out NIWO removes every 2020 observation from training, whereas holding out WREF or YELL leaves mixed-year training data. The held-out scores therefore describe the observed combinations of site and acquisition rather than a pure site effect.

The available per-flightline exports permit a bounded year sensitivity for the focal TM matched band-3 diagnostic, but not a year-matched refit of the pooled, balanced, or held-out models. Table 10

separates NIWO’s two filename-derived years and summarizes the other 2023 groups. NIWO medians are similar in 2020 and 2023, while the lower WREF and YELL distributions remain visible in 2023. This descriptive check shows that NIWO’s local median is not created solely by combining its two years; it does not separate site, acquisition year, atmospheric conditions, phenology, or processing history. A 2023-only model refit would contain 99 flightlines, including only nine at NIWO, and has not been performed. The supplied models are therefore not relabeled as year-controlled evidence.

Table 10: Filename-derived year sensitivity for within-flightline R^2 of TM matched band 3. Values are recalculated from the supplied local-fit export; they are descriptive ranges, not uncertainty intervals (*SpectralBridge bulk evidence bundle 2026*).

Site	Year	Flightlines	Minimum	Median	Maximum
NIWO	2020	23	0.6938	0.7547	0.8296
NIWO	2023	9	0.7285	0.7596	0.8317
WREF	2023	40	0.4287	0.5280	0.7017
YELL	2023	50	0.2623	0.6124	0.7888

12.2 What exactly is compared?

The exported equations use x for a matched MicaSense-derived product variable and y for a configured Landsat-derived product variable, with fitted target $\hat{y} = ax + b$. The four product pairs are listed in Table 11. Their names identify computational products rather than four sets of independently acquired satellite observations. Both sides derive from corrected NEON spectra. The table preserves exported numbering without assigning native sensor band names or wavelength intervals that the run evidence does not establish.

Table 11: Exact exported product identifiers and matched indices. Source identifiers are defined immediately below the table.

Source family	Target identifier	Matched indices
A	Landsat_5_TM	1–4
A	Landsat_7_ETM+	1–4
B	Landsat_8_OLI	1–5
B	Landsat_9_OLI-2	1–5

Family A is MicaSense_to-match_TM_and_ETM+; family B is MicaSense_to-match_OLI_and_OLI-2. Source and target index fields agree in every exported row. In this manuscript, “TM matched band 3” means family A index 3 against Landsat_5_TM index 3. It is not an assertion about native Landsat or MicaSense band numbering.

The numerical variable definition is recoverable; the run-specific physical definition is not yet complete. Required records are the source wavelengths, actual response arrays on that grid, excluded-band treatment, brightness coefficients and their order of application, and stored-value scale, offset, units, and no-data rules for both products. Generic packaged centers, widths, and defaults do not demonstrate which definitions generated this run. The audited streaming reader uses stored values; it does not establish physical scaling for them (Earth Lab, 2026an). We therefore retain slopes, correlations, and descriptive scores while treating intercepts, RMSE, bias, and reference displacements as quantities in exported units. We do not apply an assumed scale factor, infer

physical reflectance outliers from large stored numbers alone, or attribute the TM result to a particular spectral region.

12.3 Numerical summaries and their estimands

The primary collection summaries use 18 relationships within each weighting scheme (Table 12). Pooling emphasizes pixels; the balanced schemes change group contributions. The schemes reuse observations and are sensitivity descriptions, not independent replications. No confidence interval is estimated by treating the 54 records as a sample of 54 independent relationships.

Table 12: Separate summaries across the same 18 pair-band relationships. TM values refer to matched band 3.

Weighting	Relationships	Median R^2	TM R^2
Pixel pooled	18	0.9944	0.5944
Flightline balanced	18	0.9909	0.5950
Site balanced	18	0.9948	0.6043

For the TM diagnostic, the exported within-flightline R^2 agrees numerically with squared correlation. Correlation-based local fit is invariant to a nonzero uniform affine rescaling of either variable, but RMSE and intercepts are not scale-invariant. This algebraic property does not rule out spatially or spectrally varying corrections, inconsistent encodings across inputs, or nonlinear effects. A high R^2 is not agreement with the identity line; agreement between measurement methods requires examination of their differences and scale, not correlation alone (Bland and Altman, 1986). A lower R^2 is likewise not by itself an unacceptable physical error. Scores depend on the response distribution and require an application-specific interpretation.

Table 13: Different diagnostics for TM matched band 3. Local ranges summarize separate within-flightline fits. The site-fit column is in-sample; the held-out column evaluates a model trained on the other two sites. These columns are not a controlled comparison of predictive performance.

Site	Local median	Local range	Site-fit R^2	Held-out R^2
NIWO	0.7560	0.6938–0.8317	0.7810	0.5529
WREF	0.5280	0.4287–0.7017	0.5041	0.3305
YELL	0.6124	0.2623–0.7888	0.6187	0.5898

Eight of the 122 TM local fits exceed 0.8; the maximum is 0.8317. No scientific acceptance threshold is implied. The descriptive conclusion is that changing the supplied collection weights does not remove the lower fit and that departures from a close linear relationship also occur within flightlines. This does not isolate within- and between-group contributions or exclude additional pooling effects. The diagnostic was selected after inspection of the report and should be understood as exploratory, not a pre-specified hypothesis test.

WREF’s held-out score of 0.3305 is separate from its in-sample site-fit score of 0.5041. Their difference is not an estimate of a geographic transfer penalty. A controlled comparison would reserve the same WREF test observations for both models, fit the local baseline on other WREF spatial units, fit the external model without WREF, and evaluate both in identical verified units. Spatial blocking, comparable training information, and year-matched sensitivity would help distinguish

transfer from fitting optimism and acquisition composition (Roberts et al., 2017). The supplied three-site collection does not establish a general geographic domain of applicability (Meyer and Pebesma, 2021).

12.4 Reference displacement is not prediction error

The supplied displacement table evaluates each equation at the median of the per-flightline source means for its pair-band. If $\bar{x}_f$ is a flightline source mean, its definitions are

$$x_{\text{ref}} = \text{median}_f \bar{x}_f, \quad \Delta = ax_{\text{ref}} + b - x_{\text{ref}}, \quad p = 100\Delta/x_{\text{ref}}. \quad (3)$$

The same reference is used under all three weightings. The field named **absolute_correction** stores signed Δ (24 of 54 entries are negative); the median of $|p|$ is 7.54%, and its maximum is 110.92%. The maximum occurs for the pooled TM matched band-3 equation, evaluated at $x_{\text{ref}} = 357.971255$ exported units. These are reference-dependent equation displacements, not residual prediction error, an accuracy improvement, or the amount of upstream brightness correction.

The results implementation inspected at **c86da6a** uses a pixel-pooled source-mean reference by default (Earth Lab, 2026g). The supplied export is described by Equation 3, not silently recomputed under that later default. Neither definition establishes a physically meaningful correction percentage without verified value definitions and a reason to compare source and target to the identity line. Accordingly, displacement is retained here for report reconstruction, not used as a principal scientific performance claim.

12.5 Executable audit and file provenance

Run `python scripts/audit_bulk.py` from the manuscript directory using Python with NumPy and pandas. The audit verifies file hashes, run identifiers, unique scientific keys, counts, separate weighting summaries, filename-derived years, the descriptive year sensitivity, local and held-out diagnostic values, inline figure coordinates, and the displacement arithmetic. It reads the supplied exports without modifying them or downloading data. This is a reproducible compact-table example; it does not execute SpectralBridge or verify the reduction of the original raster archive.

Several local convenience exports reorganize the same supplied records:

```
all_candidate_coefficients.csv, correction_magnitude_all.csv,
weakest_translation_cases.csv, weakest_20_fits.csv,
largest_translation_corrections.csv, largest_20_corrections.csv,
pooling_strategy_comparison.csv, and visual_report_summary.csv.
```

These files are derived views for review and presentation, not independent observations or additional production evidence. The five primary numerical exports and visual report identified below remain the evidential inputs to the audit.

Table 14: SHA-256 checksums of primary supplied evidence used in this revision. Long filenames wrap; checksums identify bytes, not scientific validity.

File	SHA-256
SpectralBridge_Bulk_Results_Visual_Report.pdf	ac597ef955813205689177f806a87e31ffa3195bc1af681160681279cc043e66
translation_coefficients.csv	899eaa37521be89fb8aa521528962c807cd122770c5d606eab03540212a95382
per_flightline_regressions.csv	4911fb363bdbbae37c121c5c0c511b9fba05839d9f762e84e0f900e259d4e6e1
per_site_regressions.csv	8b43ae837cf8e1ecf0efd09c4658ccaacb25942f7cb1383c397c211a319e8cd7
leave_one_site_out.csv	a8b95002fb51ac74d7a050f5803fa6528633008922dd6eb525b1d41cf39e7f3a
translation_correction_magnitude.csv	ee0f60cf319f5e31c798ce596059410407054f66b14996769e4d721a17493a2d

13 What remains to be established

The current records answer useful questions, but they do not yet reconstruct a production run. The bulk tables lack the executed revision, manifest, input checksums, configuration, product headers, and performance measurements. The drone inventory likewise lacks its production revision, resolved configuration, checksums, and verified completion record. A reproducible example needs a manifest that connects each run identifier to those materials, together with a bounded input subset, an executable command, and expected outputs. The bulk audit verifies the exported tables; the drone inventory verifies the observed file census. R10C and WREF preserve historical warnings but are not correction references.

Table 15 keeps records generated under different revisions separate. A “not supplied” entry identifies exactly what a publication replay or archive still needs to add.

Table 15: Lineage of the evidence used in this manuscript. Missing fields are recorded rather than inferred.

Evidence	Input or scope	Generating revision	Configuration and QA identity	Archive evidence
Installed-package fixtures	Tiny deterministic normal, drone, and bulk inputs	Revisions identified by the validation and release-preparation records	Fixture commands and expected outputs are recorded; the release workflow’s supported gate covers Python 3.10–3.12 for wheels and 3.10 for the source distribution.	Exact-distribution smoke records; not production data (Earth Lab, 2026ap; Earth Lab, 2026v; Earth Lab, 2026l).
R10C flightline	Named NEON R10C HDF5; bounded fingerprint and file size recorded	d88e8d0 (full identifier in the R10C section); reports later archived at f7e7ddd	Package version 2.2.0 and stage reports recorded; a single bundle containing command, resolved configuration, environment, full input checksum, schema/threshold identifier, and report checksum was not supplied.	Checked-in machine-readable reports and cited walkthrough (Earth Lab, 2026w; Earth Lab, 2026x; Earth Lab, 2026y).
WREF visual report	Separately supplied WREF report and displayed product inventory	Not identified in the PDF	Generating configuration, QA schema, and threshold identifier not supplied.	Source PDF retained; checksum and PDF metadata recorded in the provenance JSON. Used only as a historical visual artifact.
Bulk analysis	122 accepted exported flightline identities and 18 pair-band relationships	Executed software revision not supplied; analysis run ID is recorded in Section 12	Manifest, resolved configuration, response arrays, scaling, input checksums, and performance record not supplied.	Report plus five primary CSVs with recorded SHA-256 hashes; compact-table audit is reproducible (<i>SpectralBridge bulk evidence bundle</i> 2026).
Drone campaign	43 historical flight packages: 17 from 2023 and 26 from 2024	Not supplied	Resolved configuration, completion contract, stage timing, checksums, and QA schema identity not supplied.	Post-run inventory records working HDF5, JSON, PDF, and PNG states and absent ENVI/Parquet states; no permanent archive identifier supplied (Spectral-Bridge project team, 2026).

Table 16: Outstanding checks and the claims they would enable. None is reported as completed by this table audit.

Question	Required evidence or comparison	Claim enabled
Did the drone route complete?	Replay the staged campaign after orchestration correction; verify configuration-dependent ENVI, configured Landsat-like, extraction, Parquet, and QA artifacts with a completion-enforcing manifest.	End-to-end artifact completion of the campaign, not sensor validation.
Which affine relationship was applied?	Archive the reviewed registry, registry version and checksum, source and target wavelength mapping, estimator identity, coefficients, confidence state, and resolved strictness policy.	Reconstruct the configured Landsat-like drone variables without implying empirical sensor calibration.
How does a completed drone product compare with Landsat?	Follow UAV surface-reflectance validation practice; record the Landsat Collection 2 product identifier, acquisition time, scaling, masks, cache or query record, common-grid definition, resampling, overlap, temporal offset, reference characterization, and uncertainty analysis (Origo et al., 2025; Origo et al., 2020; U.S. Geological Survey, 2026).	A bounded comparison with an independent observation, not universal agreement.
What do target variables represent?	Archive actual wavelength grids, response weights, excluded bands, brightness settings, scale/offset, and no-data rules for the executed run.	Physical interpretation of band support and error units.
Did geometry warnings affect corrections?	Trace source geometry through unit conversion and fitted terms; compare corrected results after any necessary repair.	Distinguish faulty summaries from altered computation.
Did retained poor-quality bands influence outputs?	Calculate their nonzero target-weight contributions and compare results under documented exclusion rules.	Assess spectral-support sensitivity.
What changes during brightness adjustment?	Compare pure weighting and adjusted products from identical inputs, preserving both states and coefficients.	Attribute the adjustment separately from target-band weighting.
How well does the equation transfer?	Evaluate local and external models on the same spatially held-out observations, including year-matched sensitivity.	Estimate a defined predictive comparison rather than an in-sample/held-out difference.
What resources does processing require?	Measure wall time, process-tree peak RAM, bytes read/written, concurrency, and hardware on a bounded reproducible workload; use matched conditions for any comparator.	Workload-specific resource estimates, not a hardware-independent speedup.

14 Claims and their boundaries

Table 16 lists the work still needed. The next two tables summarize what the present records already support and where each claim stops.

Table 17: Claims supported by the evidence reviewed here and their boundaries.

Claim	Evidence	Boundary
The inspected production pathway does not require HyTools at runtime.	Package dependencies and internal pipeline imports at the inspected commit; the 27–28 October 2025 transition history (Earth Lab, 2026ag; Earth Lab, 2026i; Earth Lab, 2025b).	This describes runtime dependency, not source or intellectual independence from HyTools.
Inspected components retain documented HyTools lineage.	Provenance notices in <code>neon.cube.py</code> , <code>envi.writer.py</code> , and <code>corrections.py</code> (Earth Lab, 2026n; Earth Lab, 2026ae; Earth Lab, 2026z).	The relationship differs by component: adapted code, inherited algorithms, architectural influence, and SpectralBridge-specific development should not be conflated.
The intended drone pipeline supports full-image and polygon extraction, reviewed registry-based affine variables, and optional Landsat QA while recording correction availability and provenance.	Drone workflow documentation, installed-artifact smoke tests, and relevant UAV calibration/validation practice (Earth Lab, 2026m; Earth Lab, 2026aj; Earth Lab, 2026t; Cao et al., 2019; Origo et al., 2025).	The 43-flight campaign did not produce the expected downstream ENVI or Parquet artifacts; documented capability is not evidence that translation or external QA ran on real campaign data.
The pipeline preserves named intermediate products.	Canonical path registry, pipeline source, and documented outputs (Earth Lab, 2026e; Earth Lab, 2026i; Earth Lab, 2026s).	Some realized parameters and exact weights are not fully serialized.
Configured target-band weighting executes the implemented discrete weighted average.	Source inspection and five controlled validation variations (Earth Lab, 2026aa; Earth Lab, 2026ar).	Does not validate weight provenance or satellite agreement; brightness-adjusted outputs are a subsequent state.
Stages can recover from interruption by reusing artifacts that pass their format-specific checks; QA reuse additionally matches schema, parameters, thresholds, version, and artifact fingerprints.	Restart implementation, controlled save/restart tests, and the public output contract (Earth Lab, 2026f; Earth Lab, 2026aq; Earth Lab, 2026af).	Structural checks and matching QA fingerprints do not by themselves establish that every configuration-dependent artifact was requested and produced.
Opt-in across-track processing creates independent left and right pipeline units from one source HDF5.	Half-flight implementation and pipeline integration (Earth Lab, 2026o; Earth Lab, 2026i).	Memory benefit and numerical agreement with a full-flightline run have not yet been quantified.

The present records do not yet support broad claims of universal sensor agreement, direct drone-to-Landsat calibration, equivalent behavior across ecosystems, or improved ecological inference. The immediate software need is a replay of the drone route with its artifact contract enforced. Scientific agreement then requires paired observations, explicit spatial aggregation, held-out evaluation, uncertainty propagation, and characterized independent references (Origo et al., 2025; Origo et al., 2020).

Table 18: Collection and execution evidence (continued from Table 17).

Claim	Evidence	Boundary
Bulk analysis provides grouped fits for 122 flightlines across three sites.	Supplied CSVs and operator-reported production record (<i>SpectralBridge bulk evidence bundle</i> 2026; Earth Lab, 2026u).	Shared-source synthetic products, three sites, and incomplete exact-run provenance limit calibration and generalization claims.
TM matched band 3 has lower fit under all supplied weightings and within individual flightlines.	Reanalysis of 122 local fits and three weighting schemes (<i>SpectralBridge bulk evidence bundle</i> 2026).	Descriptive group-level localization; no physical cause, accuracy threshold, or isolated pooling effect is established.
The exact built distribution can execute bounded normal and drone workflows and the bulk pipeline after clean installation. The recorded offline validation campaign reports 40 of 40 passing module variations.	Tier A wheel and source-distribution smoke records (Earth Lab, 2026ap; Earth Lab, 2026v). Public validation tables: eight modules with five variations each (Earth Lab, 2026al).	Tiny deterministic fixtures do not establish production-scale performance or scientific accuracy. The record identifies revision bebf707 but also a dirty working tree; it is bounded software-contract evidence, not an exactly reconstructable clean revision or scientific validation.
Spectral-library preflight and reporting are documented public routes and are included in the bounded exact-artifact smoke. The real drone route discovered and staged 43 observations and persisted working HDF5 and QA artifacts.	Release-candidate README and public interface description (Earth Lab, 2026aj). Post-run campaign inventory: 17 packages from 2023, 26 from 2024, and approximately 7.15 GB of output (SpectralBridge project team, 2026).	No executed spectral-library analysis was supplied as manuscript evidence, so no scientific result from that route is claimed here. Orchestration reported 43 processed and zero failed despite absent ENVI and Parquet products; complete production traversal is not established.
QA preserves warnings and unevaluated questions.	Machine-readable R10C combined report and documented QA semantics (Earth Lab, 2026w; Earth Lab, 2026y).	Statuses apply only to the diagnostics that ran.
The documented R10C flightline completes end to end.	Checked-in stage artifacts and combined QA report (Earth Lab, 2026w; Earth Lab, 2026y).	One flightline is not a multi-site performance estimate.
Candidate study flightlines span approximately 2.5–16 GB in source-file size.	Recorded candidate source sizes and the pre-specified study resource record (Earth Lab, 2026t).	Variation in source size does not yet establish the corresponding variation in peak memory, wall time, or other resource requirements.

15 Reproducibility checklist

The checklist below turns those boundaries into a practical archive specification. For every included flightline, record the following items, consistent with guidance for reproducible computational research, FAIR workflows, and packaged research objects (Sandve et al., 2013; Wilson et al., 2017; Goble et al., 2020; Soiland-Reyes et al., 2022):

1. the full input identifier, product code, site, date, input checksum, and any embedded identity metadata used to verify the filename-derived identity;
2. source file size in bytes;
3. the complete command and resolved configuration, including full- or half-flight mode and window;
4. the required artifact contract implied by that configuration and the completion rule applied to it;
5. the SpectralBridge commit, exact wheel or source-distribution checksum, and dependency environment;
6. worker count, host information, wall time, and peak resident memory;
7. hashes for source-product, corrected, convolved, tabular, and QA outputs;
8. stage-level and combined QA JSON records;
9. the exact spectral-response definitions and brightness coefficients used;
10. for drone affine products, the reviewed registry version and checksum, estimator identity, wavelength mapping, coefficient identifiers and values, confidence state, and strictness policy;
11. for optional external Landsat QA, the product identifier, acquisition time, masks, query or cache record, common grid, resampling, overlap rule, and temporal offset;
12. any manual exclusions, reruns, or interpretation decisions;
13. both the original and corrected run records when a completion-accounting defect is discovered, including which intermediate artifacts were reused;
14. links between every reported study statistic and the contributing source-flightline records; and
15. the bulk run identifier, manifest, source catalog, duplicate decisions, weighting scheme, stored-value scaling, reference-value definition, and streaming, compatibility, or materialization mode when bulk analysis is used.

The R10C artifact records its generating commit and bounded input fingerprint, but it does not package the command, resolved configuration, dependency environment, full input checksum, and report checksum as one replay record. A publication archive should close that gap and should either rerun R10C under the publication commit or retain the revision difference explicitly (Earth Lab, 2026w; Earth Lab, 2026y). The drone archive has the complementary gap: it preserves intermediate campaign output but does not identify the production revision or demonstrate fulfillment of the requested downstream contract (SpectralBridge project team, 2026).

16 Figure and table provenance

The conceptual sensor figure is drawn directly in `main.tex`. Its Landsat 8 OLI and MicaSense red-band centers and widths come from `figures/sensor_band_example.json`, which records the pinned repository source (Earth Lab, 2026q). The native drawing uses centers of 654.6 and 650 nm and full widths at half maximum of 37.5 and 16 nm, respectively. The spectrum and Gaussian response profiles are illustrative and prominently labeled as approximations rather than measured response curves. Landsat 8 measured relative spectral responses are documented independently (Barsi et al., 2014); a future revision could replace the Gaussian illustration with versioned measured curves if those data are incorporated with matching provenance. No such substitution is made here.

The R10C summary figure is drawn directly in `main.tex`, with labels and rounded values checked against `figures/r10c_qa_extract.json`. Stage statuses and highlights in that extract are transcribed from the combined QA JSON; unevaluated questions are drawn from the cited stage records and walkthrough. The extract records both the archive commit and the reports’ generating revision (Earth Lab, 2026w; Earth Lab, 2026x; Earth Lab, 2026y). Before submission, the extract should be regenerated automatically from checksum-verified copies of those records or independently checked against each one.

The three-pipeline diagram, completion-gate diagram, and bulk diagnostic figure are drawn directly in `main.tex`. The pipeline’s bulk row follows the September streaming revision; its drone row represents the intended contract and is explicitly distinguished from the observed campaign inventory. The completion-gate figure is conceptual and contains no measured quantities. The bulk diagnostic’s three observed flightline ranges and medians are checked against the supplied CSVs by `scripts/audit_bulk.py`. Results, definitions, and primary file hashes are included directly in this supplement. The original report and exports are unchanged. The audit checks numerical consistency and does not rerun the raster workflow.

Typography and colors are defined in the preambles; drawings compile as vector graphics without Python execution. Status labels distinguish warnings from passing checks without relying on color alone. In the bulk figure, range lines and filled median markers represent local fits; ranges are not confidence intervals. Held-out scores are reported separately so the figure does not suggest a paired predictive comparison.

The WREF figures are composed directly in `supplement.tex` from the retained original PDF, whose checksum was verified against `figures/wref_visual_qa_provenance.json`. Original vector plots from page 2 are cropped and reflowed with new headings, top-row axis labels, and a caution note; the plotted data, tick values, and legends are unchanged. The nodata-exclusion note is retained in the new reading guide. Selected page 4 inventory values were checked against the PDF text: 18 files, 32.9 MB, 151 merged rows, and 964 columns comprising 426 raw, 426 corrected, 58 convolved, and 54 metadata columns. The merge panel’s 16 resampled inputs and the quality check’s conflicting absence message are both retained. This derivative does not reproduce every field on the original page or resolve its conflicting messages. The complete source report remains unchanged, including the product overview and spectral checks not reproduced here. Its filename, checksum, PDF creation metadata, and evidence limitations are recorded in the provenance JSON. Values from this visual report are not used as cohort observations.

A Software lineage and implementation provenance

SpectralBridge began as an effort to use HyTools directly for NEON HDF5 and ENVI input/output, image and chunk iteration, topographic and BRDF correction, spectral resampling, and output writing (Chlus et al., 2025; Earth Lab, 2025a). As sensor-specific, operational, and evidential requirements accumulated, the team repeatedly modified, specialized, and replaced parts of that pathway. The current package is therefore not a thin interface around HyTools. HyTools nevertheless remains its closest implementation predecessor and an important part of its scholarly provenance.

The relationship is not captured by a single label. It includes published scientific methods implemented by both projects, source explicitly adapted from HyTools, historical HyTools source retained outside the active package tree, interoperability work during the transition, and workflow components for which the reviewed evidence identifies no material HyTools implementation lineage. Because the available audit does not quantify line-by-line derivation across the complete codebase, the manuscript does not convert the team’s qualitative experience of an extensive rewrite into a percentage-of-code claim.

These categories answer different questions. A scientific citation identifies the origin and interpretation of a method. A source notice identifies implementation lineage. A dependency declaration states what must be installed at runtime. A historical repository tree records development history. None substitutes for the others. In particular, implementing the same published equation does not by itself establish code adaptation, while removing HyTools from runtime dependencies does not remove adaptation already present in source.

A.1 Evidence and revision boundary

The file-level claims below are bounded to explicit notices in the source files pinned for the 4 September audit and to the repository history already cited in this supplement. The later source audit supplied for this manuscript additionally reports original HyTools source and notices under `deprecated/`; the current repository exposes that historical tree, but an exact commit for the completed provenance audit was not supplied (Earth Lab, 2026ak; Earth Lab, 2026ac). No standalone `HYTOOLS_PROVENANCE.md` or provenance manifest accompanied the manuscript materials. The table therefore uses “uncertain” where the supplied evidence does not establish a stronger classification, and it does not attribute later provenance documentation backward to R10C, f7e7ddd, or c86da6a.

The runtime transition is better established. October 2025 history records direct compatibility work followed by the addition of internal NEON access, ENVI writing, correction, and resampling components (Earth Lab, 2025a; Earth Lab, 2025b). At both inspected manuscript snapshots, HyTools is absent from the declared runtime dependencies (Earth Lab, 2026ag; Earth Lab, 2026l). That supports the statement that the inspected production pathway does not require HyTools to be installed. It does not support a statement that SpectralBridge contains no HyTools-derived implementation.

A.2 Component-level provenance

Tables 19–21 separate scientific-method credit from implementation credit. “Adapted from HyTools” is reserved for source files that say so. “No material lineage detected” means only that the reviewed notices and audit did not identify such lineage; it is not a universal claim of independent creation.

Table 19: Scientific-method and implementation provenance for correction components. Classifications are evidence-bounded rather than claims of originality.

SpectralBridge component	Scientific method or source	HyTools relationship	SpectralBridge role	Attribution evidence
Solar/incidence geometry	Solar incidence on sloped terrain; established remote-sensing geometry (Teillet et al., 1997).	<code>calc_cosine_i</code> is explicitly adapted from HyTools.	Supplies pixelwise illumination geometry for topographic fitting and application.	Method literature plus source notice (Earth Lab, 2026z).
Topographic-correction framework	Terrain-illumination normalization in the remote-sensing literature (Teillet et al., 1997; Soenen et al., 2005).	Portions of the correction architecture and geometry implementation retain documented HyTools lineage.	Adds SpectralBridge masks, streamed flightline statistics, neutral fallbacks, named artifacts, and optional tiled fitting.	HyTools and pinned correction sources (Chlus et al., 2025; Earth Lab, 2026z; Earth Lab, 2026ao).
SCS+C	Sun-canopy-sensor plus C formulation (Soenen et al., 2005).	Shared published method implemented inside a module with documented HyTools adaptation; HyTools is not credited as the origin of SCS+C.	Estimates and applies per-band topographic terms in the configured processing mode.	Original method and software lineage cited separately.
BRDF correction and fitting	Kernel-driven normalization and FlexBRDF practice (Roujean et al., 1992; Wanner et al., 1995; Schaaf et al., 2002; Queally et al., 2022).	Adapted and specialized from HyTools correction and kernel code.	Fits or applies NEON-focused coefficients, reference geometry, masks, fallbacks, and stage diagnostics.	Method papers, HyTools release, and source notice (Chlus et al., 2025; Earth Lab, 2026z; Earth Lab, 2026ao).
Ross volumetric kernel	Ross model family and later kernel formulations (Ross, 1981; Roujean et al., 1992; Wanner et al., 1995).	<code>calc_volume_kernel</code> is explicitly adapted from HyTools.	Evaluates the volume-scattering term used in BRDF fitting and normalization.	Original models plus function notice (Earth Lab, 2026z).
Li geometric kernel	Li geometric-optical model family and reciprocal kernels (Li and Strahler, 1992; Wanner et al., 1995).	<code>calc_geom_kernel</code> is explicitly adapted from HyTools.	Evaluates the geometric-scattering term used in BRDF fitting and normalization.	Original models plus function notice (Earth Lab, 2026z).

A.3 Historical source, packaging, and licensing

The source under `deprecated/` must be distinguished from both active SpectralBridge modules and installed dependencies. Its presence documents historical reuse and preserves original HyTools notices; it does not by itself show that the active pipeline imports that tree or that users must install HyTools. Conversely, describing the tree as historical does not erase the adapted portions present in active files.

The inspected package metadata searches for importable packages under `src/` and does not declare

Table 20: Scientific-method and implementation provenance for input/output and sensor-translation components.

SpectralBridge component	Scientific method or source	HyTools relationship	SpectralBridge role	Attribution evidence
NEON HDF5 reading and ancillary access	NEON product structure and metadata (National Ecological Observatory Network, 2026).	Portions of NeonCube are explicitly adapted and simplified from HyTools open_neon and related access logic.	Provides bounded NEON array, metadata, mask, wavelength, and ancillary access for staged processing.	File notice and pinned source (Earth Lab, 2026n).
ENVI header and writing	ENVI representation and NEON-derived metadata handling; not a new scientific algorithm.	Portions of header and writer logic are explicitly adapted from HyTools.	Writes canonical source-product and corrected raster/header pairs with SpectralBridge naming and stage contracts.	File notice and pinned source (Earth Lab, 2026ae).
Target-band weighting / spectral resampling	Normalized response-weighted spectral integration (Chander et al., 2013; Roy et al., 2016).	Early workflows used HyTools resampling. The reviewed resample.py citation does not itself document current adapted code; material current lineage is therefore uncertain.	Implements the tested discrete weight-matrix contract and records configured target products and subsequent adjustments.	History and current source cited without an independence claim (Earth Lab, 2025a; Earth Lab, 2025b; Earth Lab, 2026aa).
Sensor translation	Empirical affine relationships fitted among same-source configured representations; not a HyTools scientific method.	No material HyTools implementation lineage was detected in the reviewed evidence.	Maintains reviewed coefficient registries and applies configured affine relationships as a distinct state.	SpectralBridge documentation and bulk records (Earth Lab, 2026m; Earth Lab, 2026b).
Drone pipeline	UAV calibration and correction requirements derive from the UAV remote-sensing literature (Aasen et al., 2018; Mamaghani and Salvaggio, 2019; Cao et al., 2019).	Correction functions may traverse the HyTools-derived correction components above; no material HyTools lineage was detected for drone discovery, staging, translation, extraction, or campaign orchestration.	Preserves drone identity, correction availability, native and configured products, extraction, and QA.	Component-specific boundary (Earth Lab, 2026m; Earth Lab, 2026z).

Table 21: Implementation provenance for workflow, analysis, and historical-source components.

SpectralBridge component	Scientific method or source	HyTools relationship	SpectralBridge role	Attribution evidence
Pipeline contracts and restart	File-mediated scientific workflow design; no single correction-method source.	Earlier orchestration interacted directly with HyTools; current stage architecture was informed by that history. No claim of source independence is made.	Defines canonical paths, required artifacts, validation, reuse, and restart boundaries.	Transition history and architecture (Earth Lab, 2025b; Earth Lab, 2026p; Earth Lab, 2026i; Earth Lab, 2026f).
QA and provenance	Scientific validation and research-object principles (Morissette et al., 2006; Goble et al., 2020; Soiland-Reyes et al., 2022).	No material HyTools implementation lineage was detected in the reviewed evidence.	Produces machine-readable stage and combined reports, diagnostic figures, evidence levels, and artifact contracts.	SpectralBridge QA and validation records (Earth Lab, 2026s; Earth Lab, 2026ap).
Bulk analysis	Grouped regression, balanced summaries, and held-out site analysis as specified in this study.	No material HyTools implementation lineage was detected in the reviewed evidence.	Streams completed products into per-flightline statistics, pooled and balanced fits, site summaries, and leave-one-site-out diagnostics.	Bulk source and documentation (Earth Lab, 2026b; Earth Lab, 2026c; Earth Lab, 2026g).
Historical deprecated/ source	Original HyTools software and notices, not a distinct scientific method.	Historical/vendored HyTools source with original authorship, copyright, and GPL information retained.	Development and archival context; not treated here as an active runtime component.	Later repository audit and current tree (Chlus et al., 2025; Earth Lab, 2026ac).

HyTools as a dependency (Earth Lab, 2026ag; Earth Lab, 2026l). However, the manuscript materials do not include an unpacked wheel and source distribution from the later provenance-audit revision. This supplement therefore does not make a final claim about whether every file under **deprecated/** is absent from both distribution formats. That packaging question should be closed by recording the exact audit commit and listing the contents of the built wheel and source archive.

HyTools 1.6.1 is distributed under GPL-3.0, and the inspected SpectralBridge revisions declare GPLv3 or GPL-3.0-or-later. Active files containing adapted portions retain HyTools attribution notices (Chlus et al., 2025; Earth Lab, 2026ag; Earth Lab, 2026l; Earth Lab, 2026n; Earth Lab, 2026ae; Earth Lab, 2026z). This is a concise record of the licenses and notices present in the cited software, not a legal interpretation.

A.4 SpectralBridge contribution boundary

The package’s contribution does not depend on presenting established correction equations or adapted low-level components as newly invented. The evidence supports a contribution at the level of organization and application: explicit sensor and product handling; named intermediate artifacts; target-band construction and recorded translation states; drone preparation and extraction; restart and completion contracts; QA and provenance records; and campaign- and collection-scale analysis (Earth Lab, 2026p; Earth Lab, 2026i; Earth Lab, 2026m; Earth Lab, 2026s; Earth Lab, 2026b). These components combine established environmental remote-sensing methods and attributed prior software into an inspectable system built for a specific cross-sensor measurement problem.

This distinction reinforces the manuscript’s broader evidential argument. Inspectable scientific software should expose not only input provenance, parameters, transformations, QA, and outputs, but also the lineage of important implementations used to perform those transformations. A reader can then distinguish the origin of a scientific method, the ancestry of its code, the package required at runtime, and the workflow contribution made by SpectralBridge without collapsing them into a single novelty claim.

B Relationship to neighboring software

SpectralBridge overlaps several established software families, but the resemblance occurs at different layers. Table 22 distinguishes similarity in algorithms, scientific objective, data access, orchestration, and scalable analysis. This comparison is important because the package’s contribution is architectural and evidential rather than a claim to have invented the underlying normalization or weighting methods.

Table 22: Relationship between SpectralBridge and neighboring correction, harmonization, and retrieval software.

System	Important resemblance	Boundary relative to SpectralBridge
HyTools	The enabling predecessor and closest methodological relative: early SpectralBridge used it directly, and current components retain documented adapted and algorithmic lineage (Chlus et al., 2025; Earth Lab, 2026n; Earth Lab, 2026ae; Earth Lab, 2026z).	A general hyperspectral-processing library. Current SpectralBridge does not require it at runtime; the divergence is a specialized file contract, explicit measurement states, restart pathway, tabular products, and staged QA evidence, not an absence of HyTools-derived code.
Harmonized Landsat and Sentinel-2 (HLS)	Applies atmospheric correction, cloud and shadow masking, co-registration, common gridding, BRDF normalization, and spectral bandpass adjustment to operational Landsat and Sentinel-2 observations (Claverie et al., 2018; Ju et al., 2025).	An operational harmonized satellite data product at common spatial support. Its multi-step design and evaluated bandpass methods clarify why SpectralBridge’s configured representations or affine variables are not, by themselves, a complete harmonized satellite observation (Claverie, 2023).
Sen2Like	Generates Sentinel-2-like harmonized or fused products from Landsat and Sentinel-2 using geometric correction, atmospheric correction, spectral band adjustment, BRDF normalization, and optional spatial fusion (Saunier et al., 2022).	A satellite virtual-constellation processor whose goal includes common geometry and resolution. SpectralBridge instead uses NEON as a translating reference and deliberately leaves spatial reconciliation to downstream analysis.
FORCE	Mass-processes Landsat and Sentinel-2 archives into analysis-ready data and supports large-area time-series analysis, including radiometric, topographic, and BRDF-related processing (Frantz, 2019).	An operational archive-processing and data-cube system. SpectralBridge is smaller in scope and emphasizes persisted intermediate measurement states for diagnosing a cross-sensor comparison.
ISOFIT	Uses optimal estimation and radiative-transfer modeling to estimate surface reflectance and atmospheric state from imaging-spectrometer radiance (Thompson et al., 2018).	Solves an upstream atmospheric-retrieval problem. SpectralBridge inherits the reflectance state of the NEON directional-reflectance product rather than retrieving reflectance from radiance.

SpectralBridge is thus closest to HyTools in methods and portions of code, to HLS and Sen2Like

Table 23: Relationship between SpectralBridge and neighboring toolbox, workflow, and data-infrastructure software.

System	Important resemblance	Boundary relative to SpectralBridge
EnMAP-Box	Provides a broad environment for visualizing and processing multi- and hyperspectral imagery, spectral profiles, and derived analyses (Jakimow et al., 2023).	A general QGIS-based imaging-spectroscopy toolbox. It supplies interactive algorithms and applications rather than SpectralBridge’s opinionated cross-sensor sequence and evidence vocabulary.
NEON utilities	Tools such as <code>neonUtilities</code> support discovery, download, stacking, and preparation of NEON data products (Lunch et al., 2026).	Data-access and preparation utilities do not themselves define the geometry normalization, target-band construction, staged artifacts, or cross-sensor QA used here.
Snakemake and Nextflow	Express file- or data-dependent computational workflows, manage ordered execution, and support reproducible restart across computing environments (Köster and Rahmann, 2012; Di Tommaso et al., 2017).	Domain-general workflow engines do not supply reflectance semantics, sensor transformations, spatial-support distinctions, or scientifically interpreted QA. They could orchestrate SpectralBridge but do not replace its domain model.
Open Data Cube and Pangeo	Organize and analyze large Earth-observation or geoscience collections and provide scalable access to multidimensional data (Killough, 2018; Pangeo Community, 2026).	Data infrastructure addresses cataloging, storage, and computation at scale. It does not specify how cross-sensor reflectance comparability should be constructed or which intermediate measurement states should be retained.

in scientific motivation, to FORCE in operational organization, and to workflow systems in its file-mediated execution. It does not replace any of them. Its narrower role is to preserve the stages of a cross-sensor construction, keep band space distinct from spatial support, and show where an unexpected pattern first appears. It is not a new correction algorithm, atmospheric-retrieval system, global harmonized product, spatial-fusion method, general workflow engine, or direct drone-to-Landsat calibration.

C Prospective QA and resource protocol

An earlier study plan specified 100 NEON flightlines. We retain it as a prospective protocol rather than applying it retroactively to either the 122 accepted flightlines in the bulk collection or the separate 43-package drone attempt. Before NEON collections are combined, the selection rule, attempted-input list, configurations, and missing QA and resource records must be reconciled. These runs answer different questions and should remain separate records.

The study treats each original source flightline as one operational unit for completion, QA prevalence, and resource reporting. Pixels are not pooled across flightlines for the primary summaries, because cube dimensions differ and pooled pixels would allow larger flightlines to dominate. Left and right processing halves, where used, remain paired partitions of one source flightline and are not counted as independent flightlines. Site is the highest unit for statements about out-of-site operational or diagnostic stability (Hurlbert, 1984).

The study record should define the sampling frame and selection rule; represented sites and years; attempted, completed, excluded, and repeated flightlines; processing partitions; publication commit and resolved configuration; execution environment; and missing or unreadable QA reports. Missing

reports remain in the denominator. Runs combined in one summary must use the same response and brightness definitions, QA schema, and thresholds. If a run is replaced, the original attempt and the reason for replacement should remain in the archive.

Completion and status outcomes will be reported as counts and percentages; continuous diagnostics will be summarized by median, interquartile range, and range. Balanced summaries will remain distinct from pooled pixel descriptions. The completed synthetic-translation results do not supply the missing upstream QA and resource measurements.

C.1 Resource-use analysis

Hyperspectral processing is computationally demanding because hundreds of co-registered bands create high-dimensional data volumes, and useful analysis must also manage noise, calibration, dimensionality, and algorithmic scaling (Plaza et al., 2009; Bioucas-Dias et al., 2013). SpectralBridge adds correction models, multiple target products, tabular extraction, and QA records to that baseline workload, making resource behavior part of the reproducibility evidence rather than a purely operational detail.

Candidate source HDF5 files recorded for the study range from approximately 2.5–16 GB. Repository guidance notes that some production workflows may require roughly 250 GB of RAM, but that value is an operational planning estimate rather than a study result (Earth Lab, 2026t; Earth Lab, 2026ai). On-disk HDF5 size is only one candidate predictor because compression, decoded dimensions, band count, data type, processing mode, and concurrency can alter memory demand. Resource reporting will therefore avoid a single unqualified memory requirement. For each run, the analysis will retain source size in bytes; raster lines, samples, bands, and data type; available compression metadata; complete-process-tree peak resident memory; wall time; worker count; backend; full- or half-flight mode and window; completion status; and host CPU and RAM information.

Peak memory will be sampled for the parent process and descendants throughout each run, with the sampling interval and aggregation rule recorded. The primary resource figure will plot peak resident memory against both source file size and decoded cube size, identify failed or incomplete runs, and distinguish worker count and full- versus half-flight processing. Associations will be stratified or adjusted for these execution settings rather than pooled without qualification. Memory-aborted runs will be reported as right-censored operational observations, not treated as completed peak-memory measurements or silently excluded. The analysis will report rank associations with each size measure; a regression will be added only if its form and residual behavior are defensible. This describes operational scaling for the evaluated environment, not a hardware-independent complexity guarantee.

For the updated bulk route, profile selected ENVI bytes and decoded dimensions, eligible observations, requested pair-bands, window size, concurrent readers, storage characteristics, process-tree peak RAM, and checkpoint/temporary/output bytes. Time cataloging, the statistics pass, aggregation, the second error-evaluation pass, and compact report generation separately. Compatibility and explicit dataset-materialization modes require their own measurements. A bounded representative workload can establish resource use for that workload; a claim of superior performance requires a stated comparator under matched conditions. The staged archive’s size does not establish a speedup or measured peak-memory reduction (Earth Lab, 2026c; Earth Lab, 2026u).

D Availability and citation status

The historical 4 September metadata snapshot reports version 2.2.0 and Python 3.10 or newer. Its license and CI/release configuration are described by the pinned sources (Earth Lab, 2026ag; Earth Lab, 2026ab; Earth Lab, 2026h; Earth Lab, 2026k). Those records remain the basis of the historical flightline source audit, not proof of the software configuration that generated either the supplied bulk run or the drone campaign attempt.

The 4 September hardening audit identified metadata, licensing, provenance, dependency, and publication-configuration issues (Earth Lab, 2026v). The 9 September commit `c86da6a` prepares 2.3.0rc1, synchronizes metadata, and reports source and exact-distribution smoke checks. External publisher configuration, tag creation, and publication remain separate owner-controlled steps (Earth Lab, 2026l). As of 25 September 2026, the repository README documented candidate installation but the public GitHub release page listed no 2.3.0rc1 release; a live check on 12 October lists `v1.0.0` as latest (Earth Lab, 2026aj; Earth Lab, 2026ah). The exact revision used for the later 43-flight drone production attempt was not supplied, so neither `f7e7ddd` nor `c86da6a` is assigned to it. The completion-accounting defect is under investigation, but this manuscript cites no corrective commit and makes no claim that a fix has passed human review or campaign replay (SpectralBridge project team, 2026).

The public GitHub Pages validation record was last updated on 14 August 2026 and should not be read as a live summary of the later release-preparation or drone-campaign evidence. Its landing page also uses broader language about harmonization, Landsat referencing, and model transfer than the evidence reviewed here establishes. This supplement therefore uses the site for documented interfaces and recorded tests while retaining the manuscript’s narrower distinctions among configured representations, artifact completion, QA, and independent scientific validation (Earth Lab, 2026al; Earth Lab, 2026ad).

The DOI `10.5281/zenodo.11167877` identifies the 9 May 2024 pre-rename `cross-sensor-cal` archive (Earth Lab, 2026j). Submission should use a permanent post-rename software archive and a packaged results record linking the run manifest, inputs, software, configuration, executable examples, outputs, and exact numerical definitions (Smith et al., 2016; Soiland-Reyes et al., 2022). The local bulk-bundle citation explicitly records that no permanent public identifier was supplied.

References

- Aasen, Helge, Eija Honkavaara, Arko Lucieer, and Pablo J. Zarco-Tejada (2018). “Quantitative remote sensing at ultra-high resolution with UAV spectroscopy: A review of sensor technology, measurement procedures, and data correction workflows”. In: *Remote Sensing* 10.7, p. 1091. DOI: [10.3390/rs10071091](https://doi.org/10.3390/rs10071091).
- Barsi, Julia A., Kenton Lee, Geir Kvaran, Brian L. Markham, and Jeffrey A. Pedelty (2014). “The Spectral Response of the Landsat-8 Operational Land Imager”. In: *Remote Sensing* 6.10, pp. 10232–10251. DOI: [10.3390/rs61010232](https://doi.org/10.3390/rs61010232).
- Bioucas-Dias, José M., Antonio Plaza, Gustavo Camps-Valls, Paul Scheunders, Nasser M. Nasrabadi, and Jocelyn Chanussot (2013). “Hyperspectral remote sensing data analysis and future challenges”. In: *IEEE Geoscience and Remote Sensing Magazine* 1.2, pp. 6–36. DOI: [10.1109/MGRS.2013.2244672](https://doi.org/10.1109/MGRS.2013.2244672).

- Bland, J. Martin and Douglas G. Altman (1986). “Statistical methods for assessing agreement between two methods of clinical measurement”. In: *The Lancet* 327.8476, pp. 307–310. DOI: [10.1016/S0140-6736\(86\)90837-8](https://doi.org/10.1016/S0140-6736(86)90837-8).
- Cao, Shuang, Benjamin Danielson, Stephanie Clare, Shawn Koenig, Carlos Campos-Vargas, and Arturo Sanchez-Azofeifa (2019). “Radiometric calibration assessments for UAS-borne multispectral cameras: Laboratory and field protocols”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 149, pp. 132–145. DOI: [10.1016/j.isprsjprs.2019.01.016](https://doi.org/10.1016/j.isprsjprs.2019.01.016).
- Chander, Gyanesh, Nischal Mishra, Dennis L. Helder, David B. Aaron, Amit Angal, Taeyoung Choi, Xiaoxiong Xiong, and David R. Doelling (2013). “Applications of spectral band adjustment factors (SBAF) for cross-calibration”. In: *IEEE Transactions on Geoscience and Remote Sensing* 51.3, pp. 1267–1281. DOI: [10.1109/TGRS.2012.2228007](https://doi.org/10.1109/TGRS.2012.2228007).
- Chlus, Adam, Zhiwei Ye, Ting Zheng, Natalie Queally, Evan Greenberg, and Philip A. Townsend (2025). *HyTools: Hyperspectral image processing library*. Version 1.6.1. Software release. DOI: [10.5281/zenodo.16900026](https://doi.org/10.5281/zenodo.16900026).
- Claverie, Martin (2023). “Evaluation of surface reflectance bandpass adjustment techniques”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 198, pp. 210–222. DOI: [10.1016/j.isprsjprs.2023.03.011](https://doi.org/10.1016/j.isprsjprs.2023.03.011).
- Claverie, Martin, Junchang Ju, Jeffrey G. Masek, Jennifer L. Dungan, Eric F. Vermote, Jean-Claude Roger, Sergii V. Skakun, and Christopher Justice (2018). “The Harmonized Landsat and Sentinel-2 surface reflectance data set”. In: *Remote Sensing of Environment* 219, pp. 145–161. DOI: [10.1016/j.rse.2018.09.002](https://doi.org/10.1016/j.rse.2018.09.002).
- Di Tommaso, Paolo, Maria Chatzou, Evan W. Floden, Pablo Prieto Barja, Emilio Palumbo, and Cedric Notredame (2017). “Nextflow enables reproducible computational workflows”. In: *Nature Biotechnology* 35.4, pp. 316–319. DOI: [10.1038/nbt.3820](https://doi.org/10.1038/nbt.3820).
- Dungan, Jennifer L., James N. Perry, Mark R. T. Dale, Pierre Legendre, Sheila Citron-Pousty, Marie-Josée Fortin, Anna Jakomulska, Maria Miriti, and Michael S. Rosenberg (2002). “A balanced view of scale in spatial statistical analysis”. In: *Ecography* 25.5, pp. 626–640. DOI: [10.1034/j.1600-0587.2002.250510.x](https://doi.org/10.1034/j.1600-0587.2002.250510.x).
- Earth Lab (2025a). *SpectralBridge history: HyTools compatibility and dependency work, 8–23 October 2025*. <https://github.com/earthlab/spectralbridge/commits/main/?since=2025-10-08&until=2025-10-23>. Includes commits 89178b6, 92c2bc7, 447d237, e3d3a41, 53e8206, 8deb979, a6656ea, and 4642eba; inspected 4 September 2026.
- (2025b). *SpectralBridge history: transition to internal NEON processing components, 27–28 October 2025*. <https://github.com/earthlab/spectralbridge/commits/main/?since=2025-10-27&until=2025-10-28>. Includes commits fc370c4, b5969f5, fa0235a, e73b739, 5677078, 6942cdb, e88473d, 9aa0159, 328922b, a241cb1, 38c42dc, and 1c6a55f; inspected 4 September 2026.
- (2026a). *Build a bulk cross-run analysis*. Source snapshot inspected 9 September 2026. URL: <https://github.com/earthlab/spectralbridge/blob/c86da6a642a8a0bef37108df31af1fed1a882bd2/docs/vignettes/bulk-analysis.md>.
- (2026b). *Build a production bulk population analysis*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/vignettes/bulk-analysis.md>. Accessed 4 September 2026.
- (2026c). *Bulk pipeline implementation: September snapshot*. Source snapshot inspected 9 September 2026. URL: <https://github.com/earthlab/spectralbridge/blob/c86da6a642a8a0bef37108df31af1fed1a882bd2/src/spectralbridge/pipelines/bulk.py>.

- (2026d). *Bulk population-analysis pipeline source*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/pipelines/bulk.py>. Accessed 4 September 2026.
- (2026e). *Canonical path registry*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/paths.py>. Accessed 4 September 2026.
- (2026f). *Carry On My Wayward Son: resume a run*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/vignettes/carry-on-wayward-son.md>.
- (2026g). *Compact bulk results interpretation implementation*. Source snapshot inspected 9 September 2026. URL: <https://github.com/earthlab/spectralbridge/blob/c86da6a642a8a0bef37108df31af1fed1a882bd2/src/spectralbridge/bulk/results.py>.
- (2026h). *Continuous-integration workflow*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/.github/workflows/ci.yml>. Accessed 4 September 2026.
- (2026i). *Core NEON pipeline source*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/pipelines/pipeline.py>. Accessed 4 September 2026.
- (2026j). *DOI and Zenodo status*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/dev/doi-zenodo.md>. Reviewed 3 June 2026.
- (2026k). *Exact-artifact release workflow*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/.github/workflows/release.yml>. Accessed 4 September 2026.
- (2026l). *Finalize SpectralBridge 2.3.0rc1 preparation*. Commit inspected 9 September 2026; preparation is distinct from publication. URL: <https://github.com/earthlab/spectralbridge/commit/c86da6a642a8a0bef37108df31af1fed1a882bd2>.
- (2026m). *Module vignette 6: process drone imagery*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/vignettes/drone-processing.md>.
- (2026n). *NeonCube and NEON ancillary-data implementation*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/neon_cube.py. Source identifies portions adapted and simplified from HyTools under GPLv3; accessed 4 September 2026.
- (2026o). *Opt-in across-track half-flight processing*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/half_flight.py. Implementation and pipeline documentation inspected 4 September 2026.
- (2026p). *Package architecture*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/dev/architecture.md>.
- (2026q). *Packaged Landsat and MicaSense band definitions*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/data/landsat_band_parameters.json. Accessed 4 September 2026.
- (2026r). *Packaged Landsat brightness adjustments derived from Table Mountain HLS medians*. <https://github.com/earthlab/spectralbridge/tree/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/data/brightness>. Immutable coefficient-directory snapshot inspected 4 September 2026.

- (2026s). *Pipeline outputs and file structure*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/pipeline/outputs.md>.
- (2026t). *Production validation record*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/dev/production-validation-record.md>. Accessed 4 September 2026.
- (2026u). *Production validation record: September snapshot*. Operator-reported production evidence; inspected 9 September 2026. URL: <https://github.com/earthlab/spectralbridge/blob/c86da6a642a8a0bef37108df31af1fed1a882bd2/docs/dev/production-validation-record.md>.
- (2026v). *PyPI publication-hardening audit*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/dev/publication-hardening-audit-2026-09-04.md>. Review dated 4 September 2026.
- (2026w). *R10C combined QA record*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/validation/artifacts/r10c-1002-20210915/qa/combined/combined_qa.json. Archived at commit f7e7ddd; report provenance records generating commit d88e8d0947718d1e33b534080c224b703987e536.
- (2026x). *R10C source-acquisition QA record and threshold provenance*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/validation/artifacts/r10c-1002-20210915/qa/stages/00_acquisition/stage_qa.json. Records generating commit d88e8d0947718d1e33b534080c224b703987e536.
- (2026y). *Real flightline QA walkthrough: R10C, D10, L002, 2021-09-15*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/validation/real-data-example.md>.
- (2026z). *SCS+C and kernel-driven BRDF correction implementation*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/corrections.py>. Adapted from HyTools; accessed 4 September 2026.
- (2026aa). *Spectral resampling source*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/resample.py>. Accessed 4 September 2026.
- (2026ab). *SpectralBridge citation metadata*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/CITATION.cff>. Accessed 4 September 2026; author block remains marked for maintainer replacement.
- (2026ac). *SpectralBridge deprecated historical source tree*. <https://github.com/earthlab/spectralbridge/tree/main/deprecated>. Current repository tree inspected 26 September 2026; the exact commit for the later provenance audit was not supplied with the manuscript materials.
- (2026ad). *SpectralBridge documentation landing page*. Documentation inspected 25 September 2026; public scientific framing is broader than the bounded evidence evaluated in this manuscript. URL: <https://earthlab.github.io/spectralbridge/>.
- (2026ae). *SpectralBridge ENVI header and writer implementation*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/envi_writer.py. Source identifies portions adapted from HyTools under GPLv3; accessed 4 September 2026.
- (2026af). *SpectralBridge outputs and file structure*. Documentation inspected 25 September 2026; defines the public per-flightline artifact contract, success signals, restart checks, and QA reuse fingerprints. URL: <https://earthlab.github.io/spectralbridge/pipeline/outputs/>.

- (2026ag). *SpectralBridge package metadata*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/pyproject.toml>. Accessed 4 September 2026.
 - (2026ah). *SpectralBridge public GitHub releases*. Inspected 12 October 2026; v1.0.0 was listed as latest; no public 2.3.0rc1 GitHub release was listed. URL: <https://github.com/earthlab/spectralbridge/releases>.
 - (2026ai). *SpectralBridge README and command-line quickstart*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/README.md>.
 - (2026aj). *SpectralBridge repository README: September release-candidate snapshot*. Pinned source inspected 25 September 2026; documents the affine drone-translation registry, optional Landsat QA, compact bulk analysis, spectral-library reporting, and exact-artifact smoke scope. URL: <https://github.com/earthlab/spectralbridge/blob/c86da6a642a8a0bef37108df31af1fed1a882bd2/README.md>.
 - (2026ak). *SpectralBridge source repository, commit f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705*. <https://github.com/earthlab/spectralbridge/tree/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705>. Accessed 4 September 2026.
 - (2026al). *SpectralBridge validation evidence*. Documentation inspected 25 September 2026; page reports 40 of 40 offline contract variations passing at revision bebf707 from a dirty working tree and identifies 14 August 2026 as its last update. URL: <https://earthlab.github.io/spectralbridge/validation/>.
 - (2026am). *SpectralBridge: Start Here*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/START_HERE.md.
 - (2026an). *Streaming bulk statistics implementation*. Source snapshot inspected 9 September 2026. URL: <https://github.com/earthlab/spectralbridge/blob/c86da6a642a8a0bef37108df31af1fed1a882bd2/src/spectralbridge/bulk/streaming.py>.
 - (2026ao). *Topographic and BRDF implementation*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/src/spectralbridge/brdf_topo.py. See also [corrections.py](#); accessed 4 September 2026.
 - (2026ap). *Validation evidence index*. <https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/validation/index.md>. Updated 4 September 2026.
 - (2026aq). *Validation—save and restart behavior*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/validation/save_restart.md.
 - (2026ar). *Validation—sensor convolution*. https://github.com/earthlab/spectralbridge/blob/f7e7ddd2e01ee5f7cd74688eb479fd12fbc0b705/docs/validation/sensor_convolution.md.
- Frantz, David (2019). “FORCE—Landsat + Sentinel-2 analysis ready data and beyond”. In: *Remote Sensing* 11.9, p. 1124. DOI: [10.3390/rs11091124](https://doi.org/10.3390/rs11091124).
- Goble, Carole, Sarah Cohen-Boulakia, Stian Soiland-Reyes, Daniel Garijo, Yolanda Gil, Michael R. Crusoe, Kristian Peters, and Daniel Schober (2020). “FAIR computational workflows”. In: *Data Intelligence* 2.1–2, pp. 108–121. DOI: [10.1162/dint_a_00033](https://doi.org/10.1162/dint_a_00033).
- Goetz, Alexander F. H., Gregg Vane, Jerry E. Solomon, and Barrett N. Rock (1985). “Imaging spectrometry for Earth remote sensing”. In: *Science* 228.4704, pp. 1147–1153. DOI: [10.1126/science.228.4704.1147](https://doi.org/10.1126/science.228.4704.1147).
- Huang, Chengquan, John R. G. Townshend, Shunlin Liang, Satish N. V. Kalluri, and Ruth S. DeFries (2002). “Impact of sensor’s point spread function on land cover characterization: Assessment and

- deconvolution”. In: *Remote Sensing of Environment* 80.2, pp. 203–212. DOI: [10.1016/S0034-4257\(01\)00298-X](https://doi.org/10.1016/S0034-4257(01)00298-X).
- Hurlbert, Stuart H. (1984). “Pseudoreplication and the design of ecological field experiments”. In: *Ecological Monographs* 54.2, pp. 187–211. DOI: [10.2307/1942661](https://doi.org/10.2307/1942661).
- Jakimow, Benjamin, Andreas Janz, Fabian Thiel, Akpona Okujeni, Patrick Hostert, and Sebastian van der Linden (2023). “EnMAP-Box: Imaging spectroscopy in QGIS”. In: *SoftwareX* 23, p. 101507. DOI: [10.1016/j.softx.2023.101507](https://doi.org/10.1016/j.softx.2023.101507).
- Ju, Junchang, Qiang Zhou, Brian Freitag, David P. Roy, Hankui Zhang, Madhu Sridhar, John Mandel, Saeed Arab, Gail L. Schmidt, Christopher J. Crawford, Ferran Gascon, Peter A. Strobl, Jeffrey G. Masek, and Christopher S. R. Neigh (2025). “The Harmonized Landsat and Sentinel-2 version 2.0 surface reflectance dataset”. In: *Remote Sensing of Environment* 324, p. 114723. DOI: [10.1016/j.rse.2025.114723](https://doi.org/10.1016/j.rse.2025.114723).
- Kampe, Thomas U., Brian R. Johnson, Michele Kuester, and Michael Keller (2010). “NEON: The first continental-scale ecological observatory with airborne remote sensing of vegetation canopy biochemistry and structure”. In: *Journal of Applied Remote Sensing* 4.1, p. 043510. DOI: [10.1117/1.3361375](https://doi.org/10.1117/1.3361375).
- Kanewala, Upulee and James M. Bieman (2014). “Testing scientific software: A systematic literature review”. In: *Information and Software Technology* 56.10, pp. 1219–1232. DOI: [10.1016/j.infsof.2014.05.006](https://doi.org/10.1016/j.infsof.2014.05.006).
- Killough, Brian (2018). “Overview of the Open Data Cube initiative”. In: *2018 IEEE International Geoscience and Remote Sensing Symposium*, pp. 8629–8632. DOI: [10.1109/IGARSS.2018.8517694](https://doi.org/10.1109/IGARSS.2018.8517694).
- Koontz, Michael J., Victoria M. Scholl, Adam I. Spiers, Megan E. Cattau, Justin Adler, Joe McGlinchy, Tessa Goulden, Brett A. Melbourne, and Jennifer K. Balch (2022). “Democratizing macroecology: Integrating unoccupied aerial systems with the National Ecological Observatory Network”. In: *Ecosphere* 13.8, e4206. DOI: [10.1002/ecs2.4206](https://doi.org/10.1002/ecs2.4206).
- Köster, Johannes and Sven Rahmann (2012). “Snakemake—a scalable bioinformatics workflow engine”. In: *Bioinformatics* 28.19, pp. 2520–2522. DOI: [10.1093/bioinformatics/bts480](https://doi.org/10.1093/bioinformatics/bts480).
- Li, Xiaowen and Alan H. Strahler (1992). “Geometric-optical bidirectional reflectance modeling of the discrete crown vegetation canopy: Effect of crown shape and mutual shadowing”. In: *IEEE Transactions on Geoscience and Remote Sensing* 30.2, pp. 276–292. DOI: [10.1109/36.134078](https://doi.org/10.1109/36.134078).
- Lunch, Claire, Christine Laney, Nathan Mietkiewicz, Eric Sokol, Kaelin Cawley, and National Ecological Observatory Network (2026). *neonUtilities: Utilities for working with NEON data*. R package version 4.0.1. DOI: [10.32614/CRAN.package.neonUtilities](https://doi.org/10.32614/CRAN.package.neonUtilities).
- Mamaghani, Baabak and Carl Salvaggio (2019). “Multispectral sensor calibration and characterization for sUAS remote sensing”. In: *Sensors* 19.20, p. 4453. DOI: [10.3390/s19204453](https://doi.org/10.3390/s19204453).
- Meyer, Hanna and Edzer Pebesma (2021). “Predicting into unknown space? Estimating the area of applicability of spatial prediction models”. In: *Methods in Ecology and Evolution* 12.9, pp. 1620–1633. DOI: [10.1111/2041-210X.13650](https://doi.org/10.1111/2041-210X.13650).
- Morissette, Jeffrey T. et al. (2006). “Validation of global moderate-resolution LAI products: A framework proposed within the CEOS land product validation subgroup”. In: *IEEE Transactions on Geoscience and Remote Sensing* 44.7, pp. 1804–1817. DOI: [10.1109/TGRS.2006.872529](https://doi.org/10.1109/TGRS.2006.872529).
- National Ecological Observatory Network (2026). *Spectrometer orthorectified surface directional reflectance—flightline, RELEASE-2026*. Data product DP1.30006.001. Accessed 4 September 2026. DOI: [10.48443/jhbd-n155](https://doi.org/10.48443/jhbd-n155).
- Origo, Niall et al. (2025). *Good Practice Guidelines for UAV-Based Surface Reflectance Validation*. Tech. rep. Version 1. Version 1; official CEOS-hosted guideline; report author list is abbreviated here. Committee on Earth Observation Satellites Working Group on Calibration and Validation.

- URL: https://ceos.org/document_management/Working_Groups/WGCV/GPG_UAV_SR_CEOS_version1.pdf.
- Origo, Niall, Javier Gorroño, James Ryder, Joanne Nightingale, and Agnieszka Bialek (2020). “Fiducial Reference Measurements for validation of Sentinel-2 and Proba-V surface reflectance products”. In: *Remote Sensing of Environment* 241, p. 111690. DOI: [10.1016/j.rse.2020.111690](https://doi.org/10.1016/j.rse.2020.111690).
- Pangeo Community (2026). *Pangeo: A community for open, reproducible, scalable geoscience*. <https://pangeo.io/>. Accessed 4 September 2026.
- Plaza, Antonio, Jón Atli Benediktsson, Joseph W. Boardman, Jason Brazile, Lorenzo Bruzzone, Gustavo Camps-Valls, Jocelyn Chanussot, Mathieu Fauvel, Paolo Gamba, Anthony Gualtieri, Mattia Marconcini, James C. Tilton, and Giovanna Trianni (2009). “Recent advances in techniques for hyperspectral image processing”. In: *Remote Sensing of Environment* 113.Supplement 1, S110–S122. DOI: [10.1016/j.rse.2007.07.028](https://doi.org/10.1016/j.rse.2007.07.028).
- Queally, Natalie, Zhiwei Ye, Ting Zheng, Adam Chlus, Fabian Schneider, Ryan P. Pavlick, and Philip A. Townsend (2022). “FlexBRDF: A flexible BRDF correction for grouped processing of airborne imaging spectroscopy flightlines”. In: *Journal of Geophysical Research: Biogeosciences* 127.1, e2021JG006622. DOI: [10.1029/2021JG006622](https://doi.org/10.1029/2021JG006622).
- Roberts, David R., Volker Bahn, Simone Ciuti, Mark S. Boyce, Jane Elith, Gurutzeta Guillera-Arroita, Severin Hauenstein, José J. Lahoz-Monfort, Boris Schröder, Wilfried Thuiller, David I. Warton, Brendan A. Wintle, Florian Hartig, and Carsten F. Dormann (2017). “Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure”. In: *Ecography* 40.8, pp. 913–929. DOI: [10.1111/ecog.02881](https://doi.org/10.1111/ecog.02881).
- Ross, Juhan (1981). *The Radiation Regime and Architecture of Plant Stands*. Vol. 3. Tasks for Vegetation Science. The Hague: Dr W. Junk Publishers. DOI: [10.1007/978-94-009-8647-3](https://doi.org/10.1007/978-94-009-8647-3).
- Roujean, Jean-Louis, Marc Leroy, and Pierre-Yves Deschamps (1992). “A bidirectional reflectance model of the Earth’s surface for the correction of remote sensing data”. In: *Journal of Geophysical Research: Atmospheres* 97.D18, pp. 20455–20468. DOI: [10.1029/92JD01411](https://doi.org/10.1029/92JD01411).
- Roy, David P., Valerii Kovalsky, Hankui K. Zhang, Eric F. Vermote, Lin Yan, S. S. Kumar, and Alexander Egorov (2016). “Characterization of Landsat-7 to Landsat-8 reflective wavelength and normalized difference vegetation index continuity”. In: *Remote Sensing of Environment* 185, pp. 57–70. DOI: [10.1016/j.rse.2015.12.024](https://doi.org/10.1016/j.rse.2015.12.024).
- Sandve, Geir Kjetil, Anton Nekrutenko, James Taylor, and Eivind Hovig (2013). “Ten simple rules for reproducible computational research”. In: *PLOS Computational Biology* 9.10, e1003285. DOI: [10.1371/journal.pcbi.1003285](https://doi.org/10.1371/journal.pcbi.1003285).
- Saunier, Sophie, Bringfried Pflug, Ignacio M. Lobos, Belen Franch, Jerome Louis, Raquel De Los Reyes, Virginie Debaecker, Elisabetta G. Cadau, Valentina Boccia, Ferran Gascon, and Sultan Kocaman (2022). “Sen2Like: Paving the way towards harmonization and fusion of optical data”. In: *Remote Sensing* 14.16, p. 3855. DOI: [10.3390/rs14163855](https://doi.org/10.3390/rs14163855).
- Schaaf, Crystal B., Feng Gao, Alan H. Strahler, Wolfgang Lucht, Xiaowen Li, Trevor Tsang, Nicholas C. Strugnell, Xiaoyang Zhang, Yufang Jin, Jan-Peter Muller, Philip Lewis, Michael Barnsley, Paul Hobson, Mathias Disney, Gareth Roberts, Michael Dunderdale, Christopher Doll, Robert P. D’Entremont, Baoxin Hu, Shunlin Liang, Jeffrey L. Privette, and David Roy (2002). “First operational BRDF, albedo nadir reflectance products from MODIS”. In: *Remote Sensing of Environment* 83.1–2, pp. 135–148. DOI: [10.1016/S0034-4257\(02\)00091-3](https://doi.org/10.1016/S0034-4257(02)00091-3).
- Schneider-Zapp, Korbinian, Marta Cubero-Castan, Da Shi, and Christoph Strecha (2019). “A new method to determine multi-angular reflectance factor from lightweight multispectral cameras with sky sensor in a target-less workflow applicable to UAV”. In: *Remote Sensing of Environment* 229, pp. 60–68. DOI: [10.1016/j.rse.2019.04.007](https://doi.org/10.1016/j.rse.2019.04.007).

- Smith, Arfon M., Daniel S. Katz, Kyle E. Niemeyer, and FORCE11 Software Citation Working Group (2016). “Software citation principles”. In: *PeerJ Computer Science* 2, e86. DOI: [10.7717/peerj-cs.86](https://doi.org/10.7717/peerj-cs.86).
- Soenen, Scott A., Derek R. Peddle, and Craig A. Coburn (2005). “SCS+C: A modified sun-canopy-sensor topographic correction in forested terrain”. In: *IEEE Transactions on Geoscience and Remote Sensing* 43.9, pp. 2148–2159. DOI: [10.1109/TGRS.2005.852480](https://doi.org/10.1109/TGRS.2005.852480).
- Soiland-Reyes, Stian, Peter Sefton, Mercè Crosas, Leyla J. Castro, Frederik Coppens, José M. Fernández, Daniel Garijo, et al. (2022). “Packaging research artefacts with RO-Crate”. In: *Data Science* 5.2, pp. 97–138. DOI: [10.3233/DS-210053](https://doi.org/10.3233/DS-210053).
- SpectralBridge bulk evidence bundle* (2026). *SpectralBridge bulk results visual report and numerical exports*. Author-supplied local evidence bundle. Report and CSV exports reviewed 9 September 2026. Primary filenames, analysis identifier, and SHA-256 checksums are recorded in the supplement. No permanent public identifier was supplied.
- SpectralBridge project team (2026). *2023–2024 drone campaign production inventory*. Author-supplied production inventory. Reviewed 18 September 2026. Records 43 staged working HDF5 observations, approximately 7.15 GB of output, QA and JSON artifacts, absent ENVI and Parquet products, and orchestration status of 43 processed and zero failed. Exact production revision, manifest, checksums, resolved configuration, root cause, and corrective revision were not supplied.
- Teillet, Philippe M., Karl Staenz, and Darrel J. Williams (1997). “Effects of spectral, spatial, and radiometric characteristics on remote sensing vegetation indices of forested regions”. In: *Remote Sensing of Environment* 61.1, pp. 139–149. DOI: [10.1016/S0034-4257\(96\)00248-9](https://doi.org/10.1016/S0034-4257(96)00248-9).
- Thompson, David R., Vijay Natraj, Robert O. Green, Michael C. Helmlinger, Bo-Cai Gao, and Michael L. Eastwood (2018). “Optimal estimation for imaging spectrometer atmospheric correction”. In: *Remote Sensing of Environment* 216, pp. 355–373. DOI: [10.1016/j.rse.2018.07.003](https://doi.org/10.1016/j.rse.2018.07.003).
- U.S. Geological Survey (2026). *Landsat 8–9 Collection 2 Level-2 Science Product Guide*. U.S. Geological Survey product documentation. Accessed 25 September 2026. URL: <https://www.usgs.gov/media/files/landsat-8-9-collection-2-level-2-science-product-guide>.
- Wang, Changhai and Soe W. Myint (2015). “A simplified empirical line method of radiometric calibration for small unmanned aircraft systems-based remote sensing”. In: *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 8.5, pp. 1876–1885. DOI: [10.1109/JSTARS.2015.2422716](https://doi.org/10.1109/JSTARS.2015.2422716).
- Wanner, Wolfgang, Xiaowen Li, and Alan H. Strahler (1995). “On the derivation of kernels for kernel-driven models of bidirectional reflectance”. In: *Journal of Geophysical Research: Atmospheres* 100.D10, pp. 21077–21089. DOI: [10.1029/95JD02371](https://doi.org/10.1029/95JD02371).
- Wilson, Greg, Jennifer Bryan, Karen Cranston, Justin Kitzes, Lex Nederbragt, and Tracy K. Teal (2017). “Good enough practices in scientific computing”. In: *PLOS Computational Biology* 13.6, e1005510. DOI: [10.1371/journal.pcbi.1005510](https://doi.org/10.1371/journal.pcbi.1005510).
- Woodcock, Curtis E. and Alan H. Strahler (1987). “The factor of scale in remote sensing”. In: *Remote Sensing of Environment* 21.3, pp. 311–332. DOI: [10.1016/0034-4257\(87\)90015-0](https://doi.org/10.1016/0034-4257(87)90015-0).
- Wulder, Michael A., David P. Roy, Volker C. Radeloff, Thomas R. Loveland, Martha C. Anderson, David M. Johnson, Sean Healey, Zhe Zhu, Theodore A. Scambos, Nima Pahlevan, Matthew Hansen, Noel Gorelick, Christopher J. Crawford, Jeffrey G. Masek, Txomin Hermosilla, Joanne C. White, Alan S. Belward, Crystal Schaaf, Curtis E. Woodcock, Justin L. Huntington, Leo Lymburner, Patrick Hostert, Feng Gao, Alexei Lyapustin, Jean-François Pekel, Peter Strobl, and Bruce D. Cook (2022). “Fifty years of Landsat science and impacts”. In: *Remote Sensing of Environment* 280, p. 113195. DOI: [10.1016/j.rse.2022.113195](https://doi.org/10.1016/j.rse.2022.113195).