

Reading the Sky Through Loops

Topological Loop-Holonomy as a Portable Instrument: Galactic Spiral Streaming and the Secular Aberration Drift, with a Mapped Competence Surface

Working Draft v0.1 · Methods / Results · Preprint, not peer reviewed · 2026

Matthew "Bear" Schulz Independent Researcher ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com License: CC-BY-4.0

Companion to the geodetic "Topological Tripwire" preprint (Schulz, 2026; DOI: 10.5281/zenodo.21287951, concept DOI 10.5281/zenodo.21287950), which validates the same estimator on GNSS and creepmeter substrates. This is a methods-and-results paper reporting one confirmed detection, one two-channel result, and an explicitly bounded map of where the instrument works and where it does not.

Abstract

We report a portable instrument for reading coherent structure in vector fields on the sphere: a loop-holonomy estimator that walks closed loops through a fixed pentagon mesh and reads the enclosed circulation as topological closure-failure. Applied to Gaia DR3 stellar kinematics (≈ 22.3 million stars after selection), the instrument recovers the Milky Way's in-plane spiral streaming field as coherent, parity-clean vorticity structure — 85 of 509 mesh rings exceeding 3σ against a pre-registered azimuth-shuffle null (chance expectation ≈ 1), organized into eight kpc-scale arcs of alternating circulation sign, peaking at 7.0σ and surviving full Bonferroni correction; the holonomy reader agrees with a standard Jeans analysis on identical stars at the 0.1% level. Applied to a quasar reference-frame sample (Gaia DR3 $\times$ ICRF3, ≈ 303 k sources), the instrument was tested against the known secular aberration drift — the $\approx 5 \mu\text{s/yr}$ dipole in quasar apparent proper motions induced by the Solar System's galactocentric acceleration. The dipole is recovered in the per-source channel at 9.4σ , amplitude consistent with the published value at 0.4σ , apex within 6° of the published direction. The loop-holonomy channel recovers the *direction* (apex within $\sim 11^\circ$, stable across all runs) but suppresses the *amplitude* by $\approx 60\%$, a deficit we localize to spatially-correlated proper-motion systematics rather than measurement noise, and which we report as a mapped instrument limitation rather than an error to be explained away. We frame the central result as a **competence surface**: loop-holonomy is a low-frequency, global-phase reader that recovers *where structure is and whether it is real* more reliably than *how much* of it there is, and that favors simple witnesses arranged in coherent fields over complex witnesses in simple ones. We give the pre-registered bars, the null batteries, the failure diagnostics, and the substrate-independence lineage (a single unmodified circulation routine spanning GPS stations, seismic creepmeters, and now galactic scales) in full.

Keywords: loop holonomy, discrete curvature, Gaia DR3, galactic kinematics, spiral streaming, secular aberration drift, ICRF3, topological data analysis, Regge calculus, robust estimation, pre-registration

1. Introduction

A great deal of astrophysical information lives not in the magnitude of a field but in its *structure* — the way vectors curl, converge, and circulate across the sky. Rotation curves, streaming motions, reference-frame drifts, and lensing deflections are all, at root, statements about how a vector field is organized. The dominant estimators for such fields are amplitude estimators: fit a model, read off a coefficient, quote an error bar. These are optimal, in the Cramér–Rao sense, against clean Gaussian noise, and nothing in this paper disputes that.

We report a *different* estimator, aimed at a different question. Instead of asking how large a field is, it asks how much it *circulates* — it walks a closed loop through the field and measures the failure of a transported frame to return to itself. This is loop holonomy, and it is old: the same object appears as the Burgers vector of a crystal dislocation (Burgers 1939), as the deficit angle of a triangulated curved surface (Regge 1961), and as the geometric phase acquired by a state carried around a cycle (Berry 1984). What is new here is not the object but its use as a *portable measurement instrument* — a single construction, applied without modification across substrates that differ by many orders of magnitude in scale, together with an honest account of the regimes in which it succeeds and fails.

This paper reports the instrument's application to two sky substrates. The first is the stellar velocity field of the Milky Way disk, from Gaia DR3, where the instrument recovers the galaxy's spiral streaming as coherent vorticity structure and agrees with a standard Jeans analysis at the sub-percent level. The second is the quasar reference frame, where the instrument is tested against a signal whose amplitude and direction are already published — the secular aberration drift — and where the result is instructive precisely because it splits: the signal is recovered cleanly in one channel and suppressed, in a diagnosable way, in another.

We treat that split not as a partial failure but as the paper's most useful result. An instrument that claims to be universally superior is not a scientific instrument; it is a marketing claim. What makes a tool trustworthy is a spec sheet — a characterization of where it reads clean and where it fights the substrate. The bulk of this paper is devoted to producing that spec sheet honestly, including reporting two clean failures and one mapped amplitude-suppression mechanism at full volume.

Throughout, we hold to a set of standing disciplines carried over from the program's geodetic predecessor: every analysis is pre-registered against a frozen charter before the data is unblinded; nulls are run first; numbers of record are computed twice-identically before being written anywhere; and failures are reported in the text as findings, never quietly dropped.

2. Origins and Methodological Lineage

The instrument did not originate in astrophysics, and its history is part of its evidentiary weight, because the central claim — that this is a *portable* reader of structural invariants — is demonstrated less by any single detection than by the fact that one unmodified construction crosses domains at all.

The lineage begins, without embarrassment, at a card table: the core routine grew from a quaternion-based dead-reckoning method for tracking state under adversarial, partially-hidden dynamics, where the quantity worth defending was not a value but an accumulated orientation that had to close consistently around a cycle of decisions. From that combinatorial problem the method moved to geophysics. In the companion geodetic work (the "Topological Tripwire" preprint), the same closure logic was applied to GNSS station networks and to the Parkfield creepmeter array, detecting episodic tremor-and-slip signatures as integer-channel anomalies in a conserved-quantity field, validated against Nevada Geodetic Laboratory data and USGS instrumentation. The routine that walks a loop of GPS stations around a fault — computing the circulation of a strain-rate field and reading its closure — is, character for character, the same routine that this paper walks around a pentagon of Voronoi cells at kiloparsec scale. The pentagon loop lineage runs continuously from a ring of five GPS stations around Parkfield to rings of mesh cells spanning the galactic disk, eight orders of magnitude apart in scale, with the estimator unchanged.

This continuity is the substrate-independence thesis in its most concrete form. The claim is not that galaxies and fault zones are physically alike; they are not. The claim is that *coherent circulation in a vector field is a substrate-independent structural fact*, and that an instrument built to read it will read it wherever it genuinely exists — and, equally importantly, will fail to read it where it does not, in ways that diagnose the mismatch rather than manufacture a false signal.

Two design commitments follow from this lineage and are load-bearing for everything below. First, the estimator reports *structure and rate*, never magnitude alone: a uniform level shift in a field is the least informative and most easily counterfeited quantity, so the instrument is built to be insensitive to it and sensitive to how the field is organized. Second, every substrate is approached under a two-regime discipline: anomalies are treated as instrument or systematic confessions until proven otherwise, never as adversarial signals, because natural substrates do not lie — they are merely read through imperfect apparatus.

3. The Instrument

3.1 The fiducial mesh

The mesh is pure geometry, constructed with no reference to the data it will later read. On the sphere (for the quasar application) or the flat Galactocentric plane (for the disk application), we lay down a fiducial tiling of cells whose vertices and edges are frozen in advance. For the disk application the tiling is a triangular lattice with a controlled set of topological defects, producing cells of five, six, and seven edges — pentagons, hexagons, heptagons — carrying discrete curvature charge $q = 6 - n$, where n is the cell's edge count. A pentagon carries charge $+1$ (a positive-curvature, gravity-well-like defect); a hexagon carries 0 (flat); a heptagon carries -1 (a saddle). The charges of interior cells must pair, and any global imbalance is absorbed by the boundary — a discrete Gauss–Bonnet ledger enforced as a construction check.

The cell census for the disk application is shown in Figure 1. That the mesh is data-independent is the load-bearing property. Stars, or quasars, are *assigned into* cells after the mesh is fixed; they never deform it. This is

what makes the instrument's split-tests meaningful: a sample can be cut by brightness, colour, source class, or distance, and every cut rides the *identical* rigging, so any difference between cuts is a statement about the data and never an artifact of a shifted geometry.

3.2 Consensus at a cell

Each cell's population of vectors is reduced to a single consensus vector by a robust estimator — a geometric median of directions, which tolerates up to half its inputs being corrupted before it can be dragged off. Crucially, the vote is on *direction*, not amplitude: the standing law that magnitude never buys influence is built into the read operation itself, so a handful of wild-amplitude outliers cannot capture a cell. Cells whose occupancy falls below a pre-registered floor are ruled data-starved and excluded, never silently filled with a spurious value.

3.3 The loop walk

Loops are closed paths through adjacent cell centres, walked along the frozen mesh edges. A frame is carried from cell to cell by parallel transport and composed around the loop; the composition, in the group of rotations, does not in general return to the identity, and the failure to close — expressed as the logarithm of the residual rotor, a pure axis-times-angle vector — is the enclosed holonomy. The order of composition matters (the group is non-abelian), so the discipline is *compose first, take the logarithm last*: the loop is walked as a product of rotations and only then reduced to a comparison vector, never flattened prematurely into an additive shadow that would discard exactly the non-commutative content the instrument exists to read. On a curved substrate the exact geometric baseline of each loop — its spherical excess, a computable quantity — is subtracted, so that the residual holonomy is the *field's* and not the sphere's.

3.4 Parity audit and nulls

Every loop admits a mirror operation: reverse its traversal. A genuine curvature signal is odd under this reversal (it flips sign cleanly); an estimator bias is even (it survives the flip). This parity mirror runs on every field and provides an artifact veto independent of any null model: circulation that is not cleanly parity-odd is not curvature. On top of this, every application runs a pre-registered null battery *first*, before the real read is unblinded — direction shuffles, stratified permutations that preserve the sample's own selection footprint, and rigging rotations that expose anything locked to the mesh rather than to the sky. The measured null distributions, not any assumed Gaussian, set the disks against which significance is judged.

4. Application I — Galactic Spiral Streaming (Gaia DR3)

4.1 Sample and baseline

From Gaia DR3 we selected a 6D-complete stellar sample (positions, parallaxes, proper motions, and radial velocities) under a frozen quality-only cut, yielding ≈ 22.3 million stars across the disk fence; a back-of-

envelope prediction of the fenced count agreed with the star-level result to within 1.3%. The analysis sample of record was constructed under a tracer-homogeneity rule selecting an old, kinematically-relaxed population, whose members respond to the smooth potential rather than carrying natal streaming — the "similar bells" principle, with its dynamical justification in the phase-mixing of hot populations. Each star's residual velocity was formed by subtracting a pre-registered axisymmetric baseline *at that star's own position*, with the baseline itself built for the selected tracer including its asymmetric drift, so that a population's mean lag could not masquerade as structure.

A note of record emerged here as a diagnostic in its own right: a rotation-curve constructed from the measured-counts density profile rose unphysically toward 250 km/s, because a distance-truncation in the count profile poisons the outer annuli, whereas a pinned exponential profile gave the flat ≈ 232 km/s the literature reports. This is an in-crate demonstration of *why* published analyses pin an external density profile, discovered by declining to choose a convention silently.

4.2 Verdict K1 — agreement with Jeans analysis

On identical stars, the holonomy reader and a standard Jeans analysis were compared over the fence. The median absolute deviation across the annuli was $\approx 0.7\%$ on the circular speed and $\approx 1.4\%$ on the enclosed mass — two orders of magnitude inside the pre-registered bars, in both velocity-convention treatments. Two fence-edge annuli produced large deviations traceable to one-sided coverage where the annulus fell half outside the live-cell footprint; the median statistic, pre-registered for exactly this reason, is robust to them. We record K1 as **PASS**, of validation class only: agreement with a standard method establishes that the instrument reads quantitatively true, and we explicitly note that this agreement is *not* independent evidence for any dark-structure claim, because the zeroth-order circulation equals the rotation curve by construction and is quarantined under the tautology guard.

4.3 Verdict K2 — recovery of spiral streaming

The residual velocity field, read as loop holonomy across the mesh, is not featureless. Against the pre-registered azimuth-shuffle null, 85 of 509 mesh rings exceed 3σ , where chance expectation is ≈ 1 (Figure 2). These rings are not scattered: they organize into eight coherent patches, kpc-scale arcs of alternating circulation sign, with a positive-vorticity arc near $R \approx 6.4$ kpc spanning $\approx 19^\circ$ of azimuth and peaking at 6.9σ , a neighbouring negative arc near $R \approx 6.7$ kpc, and further alternating arcs out to $R \approx 10$ kpc, at amplitudes of ± 5 – 15 km/s/kpc. Neighbouring-ring correlation is 0.589 where the null gives zero. The circulation is parity-odd to 6×10^{-15} — pure curvature class, with no measurable bias component. The peak rings survive full Bonferroni correction, including the 7.0σ ring surviving alone.

This morphology and amplitude match the known Gaia in-plane spiral streaming field — arm-scale velocity waves — now read out through loop holonomy. We record K2 as **PASS via the K2b bar**: the pre-registered bar for K2b explicitly counts recovery of known kinematic structure (spiral or bar signatures) as a pass, and every clause is met — a $\geq 3\sigma$ cell-level feature against the frozen null, surviving the parity-mirror audit. We are careful about the verdict's reach. A stricter claim — that the loop-holonomy channel carries information a

standard abelian (Oort-tensor) analysis lacks — is a *different* claim, and it is not demonstrated: the same patches are recoverable in the abelian channel, and any non-abelian excess is at the 0.1% level. That stricter question is retained, unfolded, as open. And the tautology guard holds: spiral streaming is structure *beyond* the axisymmetric baseline, not the rotation curve itself.

The K2b result was hardened against a pinned-constants perturbation row: perturbing R_0 and the solar motion within their errors left all 85 patches standing (indeed, injecting ± 3 km/s of fake solar motion spawned only 5–7 marginal rings while the 85 real ones sat unmoved — the perturbation test doubling as a measured sensitivity to the very systematic it probes). A residual axisymmetric offset in the null was diagnosed as finite-annulus-width leakage, a systematic that cannot manufacture alternating $\pm$ patches; the patches stand regardless.

4.4 The clean failures

Two further routes were tested and failed cleanly, and we report them as deliverables. K2a asked whether the mesh route reduces variance against the Jeans baseline; the confidence-interval ratio was 1.00 against a ≤ 0.80 bar — the mesh buys no variance advantage, the information living in the stars and no aggregation path changing that. K2c asked whether the field carries non-trivial integer winding; the measured winding was zero across all rings. A third channel, K3, probed a stellar-stream substructure and returned data-starved with a printed count (1,047 band members), with its reopen trigger named. Each of these is worth more as an honest closed door than as a softened maybe; the verdict document is stronger for two of three secondary routes having died.

5. Application II — The Secular Aberration Drift (Gaia DR3 $\times$ ICRF3)

5.1 The test and why it is a good one

The Solar System accelerates toward the Galactic centre, and this acceleration induces a small, coherent dipole in the apparent proper motions of distant quasars — the secular aberration drift — of magnitude ≈ 5 $\mu\text{s}/\text{yr}$ toward the Galactic centre, a quantity Gaia has measured directly (published value 5.05 ± 0.35 $\mu\text{s}/\text{yr}$ toward $\alpha \approx 269.1^\circ$, $\delta \approx -31.6^\circ$; Gaia Collaboration / Klioner et al. 2021). It is an ideal instrument test: a coherent, sky-fixed, published signal against which a new reader can be scored blind. On a quasar mesh the authorship of any coherent apparent-motion signal is sharply constrained — the sources cannot move, so a dipole is the observer's own acceleration, a quadrupole-and-higher pattern would be the astrometric gravitational-wave channel (reported here as an upper limit only), and anything else is an instrument confession.

The instrument required no modification; the quasars simply anchor sightlines whose apparent motions cannot be the sources moving. The published aberration value was pulled from the literature with citation and printed *before* any unblinded read, and all nulls were run first.

5.2 First light — a wrong-pointing field, and its diagnosis

The first real run did not recover the dipole, and we report this at full volume, because the sequence that followed is the instrument working as designed. The loop-holonomy arms read a coherent glide of 6.45 ± 1.03 $\mu\text{as/yr}$ pointing 128° away from the published apex — coherent (3.5σ above the measured shuffle-null disk), deterministic, and wrong. A coherent wrong answer is a systematic, not noise, and the instrument localized it three ways. First, the wrong glide lived in the anchor class: a rostering defect had caused the instrument to draw its reference quasars (VLBI-anchored ICRF3 sources) into every ring, so the whole instrument was reading a single class of source known to carry radio–optical photocentre jitter. Second, the true aberration signal was present but attenuated in the per-source field, and a magnitude-split mirror localized the attenuation to the bright end — the pre-registered bright magnitude cut had optimized the sample *away* from the faint regime where the published signal lives. Third, the verdict grammar itself was found defective: apex-error tolerances had inflated to a bar that could not fail, and those verdict strings were re-read as uninformative. Each finding was localized, not mysterious; the miss was a probe.

5.3 Recovery in the per-source channel

A second pull extended the sample into the faint regime ($\approx 303\text{k}$ sources of record after de-duplication against the anchors), with the rostering defect repaired, an informativeness gate installed on the verdict grammar, and the magnitude/class/redshift structure of every source pre-computed into balanced bins. The pre-registered prediction — that recovery should improve monotonically from bright to faint — was tested tranche by tranche.

The dipole is recovered (Figure 3). The clean per-source vector-spherical-harmonic fit gives 4.83 ± 0.51 $\mu\text{as/yr}$ toward $(248.2^\circ, -21.7^\circ)$ — 9.4σ from zero, amplitude consistent with the published 5.05 ± 0.35 at 0.4σ , apex 21° off (1.9σ combined). The faintest tranche alone ($G \in (19.25, 20]$) gives 6.10 ± 1.03 $\mu\text{as/yr}$ toward $(269^\circ, -26^\circ)$ — 6.1° from the published apex. The first-light diagnosis was correct: the signal lives faint, and the initial cut had starved it. The pre-registered apex axis is strictly monotone across tranches (separation from the published apex of $62^\circ \rightarrow 27^\circ \rightarrow 24^\circ \rightarrow 6^\circ$, bright to faint), confirming the jitter/bright-wall story on direction. The amplitude axis broke monotonicity at the final step, but at only 1.0σ on the thinnest tranche; the locked verdict is recorded as non-monotone as pre-registered, with the break read as tranche noise rather than a deeper effect.

5.4 The holonomy channel — a mapped amplitude suppression

The per-source channel recovers the signal; the loop-holonomy channel does not fully, and the way it fails is the paper's competence-map result. After transfer-function correction (calibrating out a deterministic ≈ 0.80 aperture-smoothing dilution, which arises because a dodecahedral mesh vertex sits $\approx 37^\circ$ from its face centre, so a ring reads an aperture-smoothed sample of a dipole shrunk by $\approx \cos 37^\circ \approx 0.80$ — a construction constant, measured blind and confirmed in both arms), the holonomy arm recovers the dipole's *direction* well (apex within $\approx 11^\circ$ of published, stable across every run, persisting under rigging-rotation nulls as a sky-fixed signal) but its *amplitude* at roughly half strength.

We tested our own leading explanation and it failed, which is the useful part. The transfer battery was homoscedastic, while the real sample spans per-source errors from ≈ 60 to ≈ 1200 $\mu\text{as/yr}$; we hypothesized that an unweighted median vote over that spread suppresses amplitude, and pre-registered the prediction that a heteroscedastic battery drawn from the real error distribution would restore the amplitude. It did not: the corrected transfer coefficient was statistically unchanged, and the corrected amplitude remained 4.5σ low. The prediction's failure is itself informative — white-noise synthetics at the real error spread reproduce only $\approx 20\%$ amplitude loss, while the real data shows $\approx 60\%$, so the missing $\approx 40\%$ lives in *structure the synthetics do not contain*. The leading suspect is spatially-correlated proper-motion systematics within cells (scanning-law correlations), to which a median-vote-then-compose construction responds very differently than to independent noise. Per the pre-registered terms, the holonomy amplitude channel is flagged for redesign against correlated systematics, not re-tuned; no third attempt was taken.

A companion decimation test (comparing the holonomy reader against a plain global vector-spherical-harmonic fit across progressively smaller subsamples) found no crossover: the global fit dominates on amplitude recovery at every source count, in both clean and jitter-inclusive subsets. The economy claim — that robust loops reach frame accuracy with fewer sources — fails cleanly on this substrate. What the loops buy is not amplitude economy but direction, localization, and audited degradation.

6. Discussion — The Competence Surface

The two applications, read together, describe the instrument's character (summarized in Figure 4), and we state that character as a set of falsifiable hypotheses rather than settled laws, each earned on the substrates tried here and each testable on the next.

The instrument is a low-frequency, global-phase reader. A loop closure is fundamentally a phase measurement — how much a transported frame rotated — not an amplitude measurement. This places the instrument on the conjugate side of a Fourier-style trade from amplitude estimators: it reads coarse, global, coherent structure (the galactic-scale glide, the kpc streaming arcs) and goes blind on fine, local structure (one quasar's photocentre weather). It is therefore not a competitor to amplitude estimators such as the global harmonic fit; it is their complement. Where the two were compared head-to-head on amplitude, the global fit won at every sample size — as it should. Where the instrument is strong is orthogonal: *which* structure is where, and whether it is real.

The instrument favours simple witnesses in complex configurations, and fights complex witnesses in simple ones. The Gaia stars — enormous numbers of individually simple test particles arranged in a complex field — are a home-team substrate, and the instrument read them at the 0.1% level. The bright VLBI quasars — individually complex, structured, jittering objects — are an away game, and the instrument's amplitude channel suppressed their signal. The faint quasars sit closer to the instrument's focal plane, where the sources are nearly ideal fixed points and the signal lives in their arrangement.

The instrument needs at least two vantages to close a loop. Holonomy requires a path along which state is transported and around which closure can fail. A single steady source with no differential — one pin, no loop

— offers nothing to walk. This is a loop-existence requirement, not a distance limit, and it delimits where the instrument can be pointed at all.

A topological mismatch is a statement about fit, not about quality. When a substrate offers no genuine transported state with non-trivial loop closure, the instrument reads nothing — and this is not a failure of the tool but a mismatch of data to tool. Examining what a substrate does under the instrument, even in mismatch, can reveal that a field carries structure its own community had not catalogued. The program's deliverable is therefore not a claim that loop-holonomy reads everything, but a mapped competence surface: the honest spec sheet of a characterized instrument.

We put this plainly, as the organizing principle behind every pre-registration and null in this paper: an instrument is a partner in a dance, not a harpoon for a white whale. It works when the substrate carries real invariants to walk on, when the signal is of the coherent, global-phase class the instrument reads, and when there are at least two vantages to close a loop against. Choosing the substrate, the signal, and the estimator to match the field's geometry is not a limitation to be apologized for; it is the whole of using the instrument well.

7. Reproducibility and Discipline

Every number of record in this paper was computed twice-identically before being written, and each analysis stage carries a content hash so that a reader with the code and staged data can confirm byte-for-byte agreement. Data was staged to persistent, human-controlled storage with per-chunk checksums and a frozen query manifest; no analysis stage consumed data that had not passed a custody check. Every verdict was pre-registered against a frozen charter, with both pass and fail conditions declared, before the relevant data was unblinded; nulls were run before signals; and constants were pinned from machine-readable sources at execution rather than from memory. The failures reported here — the two clean secondary-route nulls in Application I, the first-light wrong-glide and its diagnosis in Application II, and the mapped amplitude suppression of the holonomy channel — are in the text by policy, as findings.

8. Future Work

The holonomy amplitude channel is flagged for redesign against spatially-correlated systematics; a construction that reads phase differently from the current median-vote-then-compose scheme is the natural next step. The stricter non-abelian question left open in Application I — whether loop-holonomy carries information a standard abelian analysis cannot — is retained as a target for a substrate where the two channels can be cleanly separated.

On the atmospheric side, the same estimator ports naturally to rotational storm fields, and preliminary work is encouraging where the substrate supplies clean anchor points: rotating cyclonic systems present exactly the coherent, well-sampled circulation the instrument reads best. Storms whose rotation builds from the bottom up are harder, not because the tool fails but because the available composite radar products (NEXRAD-class) mix

scales in a way that must be decomposed before honest anchor points can be recovered — a substrate-conditioning problem ahead of the instrument, and a live line of work. In parallel, the program is returning to its geodetic roots to test new ideas around strain-accumulation precursors, where the anchor structure is best understood.

The broader claim implicit throughout — that loop-holonomy is a portable reader of structural invariants wherever a genuine transported frame closes around a loop — is being pursued on additional substrates under the same discipline: match the tool to the field, pre-register the bars, and report the failures at full volume. Those results will be reported separately as each clears its own validation pass.

Acknowledgements

This work is the sky application of a cross-domain program whose geodetic validation is reported in the companion "Topological Tripwire" preprint (Schulz 2026, DOI 10.5281/zenodo.21287951). Analysis, drafting, and adversarial review were carried out in extended collaboration with Anthropic's Claude, used as a reasoning and engineering partner under the author's direction and ratification; all scientific claims, pre-registrations, and verdicts are the author's own. The author thanks the Gaia Collaboration and the maintainers of the ICRF3 and Nevada Geodetic Laboratory data products, without whose open data none of this would be possible.

Data availability: Gaia DR3 and ICRF3 are public. Staging manifests, frozen charters, analysis scripts, and per-stage checksums are available on request and will accompany the deposit.

Figures




Figure 1 — The fiducial mesh census (Gaia disk). The frozen 5/6/7-cell alphabet (pentagons $q=+1$, hexagons $q=0$, heptagons $q=-1$) and the discrete Gauss–Bonnet ledger: interior charge sums to +9 with the boundary

absorbing -9 . The mesh is constructed with no reference to the data; stars are assigned into cells afterward. 1.97M stars occupy live cells (occupancy ≥ 200).




Figure 2 — Galactic spiral streaming, read as loop holonomy. Left: the z -score of every mesh ring against the pre-registered azimuth-shuffle null; 85 of 509 rings exceed 3σ (chance ≈ 1), colored by vorticity sign, peaking at 7.0σ and parity-odd to 6×10^{-15} . Right: the ring z -distribution against a unit Gaussian, showing the two-tailed excess that constitutes the detection.




Figure 3 — Secular aberration drift: the two-channel result. Left: dipole amplitude — the published value (Gaia/Klioni), the per-source fit (recovered, 0.4σ from published), and the transfer-corrected holonomy channel (4.5σ low, flagged for redesign). Right: apex direction recovered monotonically as the sample goes

faint ($62^\circ \rightarrow 6^\circ$ from the published apex, bright to faint), confirming the bright-jitter diagnosis on the direction axis.




Figure 4 — The competence surface. Where loop-holonomy reads clean and where it fights the substrate, across the results reported here: Gaia stars (home team), faint quasars (recovered), bright quasars (direction only, amplitude suppressed), and topological mismatch. Direction is holonomy's channel; amplitude is the global fit's.

References

(To be finalized at deposit; DOIs to be verified per standing discipline.)

- Berry, M. V. (1984). Quantal phase factors accompanying adiabatic changes. *Proc. R. Soc. Lond. A* 392, 45.
- Burgers, J. M. (1939). Some considerations on the fields of stress connected with dislocations in a regular crystal lattice. *Proc. K. Ned. Akad. Wet.* 42.
- Gaia Collaboration, Klioner, S. A., et al. (2021). Gaia Early Data Release 3: Acceleration of the Solar System from Gaia astrometry. *A&A* 649, A9. arXiv:2012.02036.
- Gaia Collaboration (2023). Gaia DR3 (various data-model and validation papers). *A&A*.
- Eilers, A.-C., et al. (2019). The circular velocity curve of the Milky Way from 5 to 25 kpc. *ApJ* 871, 120.
- Regge, T. (1961). General relativity without coordinates. *Nuovo Cimento* 19, 558.

- Schulz, M. (2026). Topological Tripwire: A Verification Architecture Detecting Change Through Topological Invariants [geodetic preprint]. *Zenodo*. DOI: 10.5281/zenodo.21287951 (v1); concept DOI 10.5281/zenodo.21287950.
- Schulz, M. (2026). Turtles, or Crystals [framework preprint]. *Zenodo*, DOI 10.5281/zenodo.20484438.

Preprint, not peer reviewed. Draft v0.1, 2026. Shared under CC-BY-4.0 for the benefit of anyone it may help.