



1

1 Encoding-dependent verdicts and H_1 miscalibration 2 in ensemble persistent homology of cyclone 3 trajectories

4 **Rongzhen Dai**¹

5 ¹ Independent Researcher, Yangzhou, China. Correspondence: dairongzhen3@gmail.com

6

7 Abstract

8 Ensemble topological data analysis (TDA) on multi-storm cyclone trajectories has been
9 proposed as a tool to detect coherent perturbations — such as those associated with
10 extreme geomagnetic events — that may not register in any single-storm intensity time
11 series. We construct three principled longitude encodings for the cyclone point cloud
12 (linear modular, unit-sphere, and lat-linear-plus-longitude-circle cylinder) and apply the
13 same ensemble-TDA pipeline to the same data: all storms whose lifetime overlaps the
14 ± 15 -day peak window of Halloween 2003, St Patrick’s Day 2015, or Gannon 2024 (25
15 event storms; 1,020 calendar-matched controls). Three findings emerge. First, the $D1$
16 pool H_1 permutation p-value depends on encoding choice in a way that flips the
17 qualitative verdict: $\text{perm-p}_{H1} = 0.130$ (linear), 0.009 (unit-sphere), 0.214 (cylinder), on
18 identical event and control sets. Second, a 49-placebo calibration returns an H_0 false-
19 positive rate close to the nominal 5 % in all three encodings (8.2 %, 4.1 %, 6.1 %), but an
20 H_1 false-positive rate that is consistently above nominal in all three (8.2 %, 10.2 %, 14.3
21 %; directional consistency across encodings, not individually significant at $n = 49$ — the
22 cylinder rate sits at the Wilson 95 % upper bound) — H_1 calibration is inflated by a
23 factor $1.6\times\text{--}2.9\times$ regardless of encoding choice. Third, this dimension asymmetry
24 between H_0 (calibrated) and H_1 (inflated) recurs in stratified attribution across two
25 basin cells and three single-event cells, and in subsample-size sensitivity tests at $N \in$
26 $\{400, 600, 800, 1000\}$. We interpret the pattern as observational evidence that ensemble
27 TDA on cyclone-trajectory data carries an intrinsic H_1 inflation that is not curable by
28 lon-encoding choice alone. The solar-perturbation hypothesis cannot be tested through
29 this pipeline until an encoding-invariant H_1 -calibrated protocol is in place. We frame
30 the contribution as a methodological cautionary tale: pooled point-cloud TDA on
31 geospatial trajectories with periodic coordinates is more fragile than its formal stability
32 theorems suggest. We propose a minimum protocol of multi-encoding agreement testing
33 plus dimension-resolved placebo calibration as a precondition for any positive ensemble-
34 TDA claim in this class of problems.

35



1

1 1. Introduction

2 1.1 Solar–cyclone hypothesis: a debated observational lineage

3 The hypothesis that solar and geomagnetic activity perturbs tropical cyclone (TC)
4 dynamics has a long and contested observational lineage. Elsner & Kavlakov (2001)
5 brought the question to peer-reviewed satellite-era cyclone data. They reported
6 correlations between Kp / Dst geomagnetic indices and per-storm intensity changes in the
7 North Atlantic. Elsner & Jagger (2008) found that US and Caribbean hurricane counts
8 vary inversely with sunspot number, and proposed a heat-engine pathway in which an
9 active Sun warms the upper troposphere via ozone-mediated ultraviolet absorption,
10 modulating a storm’s maximum potential intensity. Pérez-Peraza et al. (2008) and
11 Mendoza & Pazos (2009) identified a ~22-year cycle in Atlantic and Pacific hurricane
12 occurrences that coincides with the magnetic solar cycle. Hodges & Elsner (2011)
13 confirmed a solar-cycle dependence for US hurricanes over 1866–2008. Hodges, Jagger
14 & Elsner (2014) refined the analysis with a Bayesian space–time model that resolves the
15 modulation geographically. More recent work extends the analysis to other basins: Li et
16 al. (2023) report an asymmetric modulation of solar activity on TC frequency over the
17 western North Pacific, and Nayak (2024) provides a contemporary review. Across the
18 lineage, reported effect sizes are modest, proposed mechanisms vary, and null-hypothesis
19 controls are unevenly applied.

20 1.2 Mechanistic candidates: cosmic-ray–cloud and atmospheric ionization

21 Several physical mechanisms have been proposed to underwrite the observed
22 correlations. Most prominent is the cosmic-ray–cloud (CRC) hypothesis. Heliospheric
23 modulation of the galactic cosmic ray (GCR) flux is associated with variation in the rate
24 of ion-induced new-particle formation in the lower atmosphere, and so with the
25 abundance of cloud condensation nuclei (CCN). Marsh & Svensmark (2000) reported a
26 correlation between GCR flux and low cloud cover during one solar cycle. Svensmark
27 (2007) developed the finding into the broader cosmo-climatology programme. Laboratory
28 work has tested the ion-induced nucleation step directly. The CLOUD experiment at
29 CERN (Kirkby et al., 2011) showed that GCR-relevant ionization enhances nucleation
30 rates of sulphuric-acid clusters under controlled conditions. The magnitude is sensitive to
31 the presence of stabilising species such as ammonia. Numerical modelling by Yu & Luo
32 (2014) explored how solar-cycle GCR variation translates to CCN concentration. They
33 found a response larger than earlier estimates but still small relative to observed climate
34 variability. Svensmark et al. (2016) used Forbush-decrease compositing to argue for a
35 detectable cloud response, although the compositing scheme has itself been contested.
36 Mechanistic candidates outside the CRC framework include direct geomagnetic-induced
37 perturbation of mesospheric and stratospheric circulation. Solar-cycle signatures may
38 propagate downward to the troposphere via wave–mean-flow interaction (Huth &
39 Stryhal, 2024). Ionospheric coupling provides another pathway: solar-wind energy input
40 is associated with auroral gravity waves that descend to the tropical troposphere and may
41 co-vary with convective instability associated with rapid intensification (Mironova et al.,
42 2018).



1

1 1.3 Counter-evidence: the null literature

2 Against this body of suggestive correlations stands an equally substantial literature
3 finding no robust effect. Kristjánsson et al. (2008) examined twenty-two Forbush-
4 decrease events over MODIS-instrumented oceanic regions. They found no statistically
5 significant response in cloud parameters. Calogovic et al. (2010) reached the same
6 conclusion using a different compositing scheme. Laken et al. (2012) reviewed thirty-five
7 years of satellite cloud observations alongside daily-timescale Forbush analyses. They
8 concluded that evidence for a widespread GCR–cloud link remains elusive, with most
9 apparent variability attributable to ENSO. Pierce (2017), drawing on findings from the
10 CLOUD experiment, argued that ion-induced enhancement of CCN is too weak to
11 translate into detectable cloud-radiation forcing at climate-relevant scales. The IPCC
12 Sixth Assessment Report (IPCC AR6 WGI, 2021, Ch. 7.3.4.5) reflects the consensus
13 position with high confidence: GCRs contribute a negligible effective radiative forcing
14 over 1750–2019, and no robust association between GCRs and cloudiness has been
15 established. The result is a literature in which positive correlation reports and null-result
16 analyses coexist over two decades of investigation. We treat the question of whether solar
17 and geomagnetic activity perturbs cyclone dynamics as a genuinely debated one.

18 1.4 Scope of the present work: encoding-invariance as a precondition

19 Existing observational studies of the solar–cyclone question almost universally operate at
20 the per-storm scale, asking whether a single storm’s intensity time series correlates with
21 concurrent geomagnetic indices. The scope choice matters. An alternative, more recently
22 proposed scope applies persistent homology to multi-storm point clouds in (latitude,
23 longitude, intensity) coordinates over all TCs occurring within an event window, with the
24 natural geometric object becoming the joint point cloud of all event-period fixes and the
25 natural summary statistic becoming its persistence diagram. A coherent perturbation
26 hypothesis, if true, might leave a signature on this ensemble topology that is not visible in
27 any single storm’s trajectory. The shift in scope is not trivial.

28 The present paper steps one level back from that scope. We do not, in the end, advance a
29 positive or negative physical claim about the solar–cyclone connection at the ensemble-
30 TDA scope. Instead we ask a methodological precondition question: does the ensemble-
31 TDA pipeline give a consistent answer about the same data under principled variations of
32 how cyclone trajectories are represented as point clouds? The longitude coordinate is
33 periodic. Three principled, published-style choices exist for handling that periodicity,
34 each with its own defenders. A pipeline whose conclusions depend on this
35 representational choice — without converging across them — is not yet ready to test the
36 substantive physical hypothesis. We carry out that test here, and report what we find.

37 1.5 This paper

38 We construct three encodings of the (lat, lon, wind) point cloud: a linear modular
39 normalization that ignores periodicity, a unit-sphere embedding that maps (lat, lon) to (x,
40 y, z) on S^2 , and a cylinder embedding that places lat on a linear axis and lon on a flat (cos,
41 sin) circle. We apply the same downstream pipeline — Vietoris–Rips persistence to
42 dimension 1, Wasserstein-2 distance to a bootstrap control distribution, permutation-



1

1 based p-values — to the same event set (25 storms across Halloween 2003, St Patrick
2 2015, Gannon 2024) and the same calendar-matched control pool (1,020 storms). We
3 report D1 pool, 49-placebo calibration, five-cell stratified attribution, and four-size
4 subsample sensitivity for all three encodings on identical data.

5 Three findings emerge. The D1 H_1 permutation p-value depends on encoding: 0.130 /
6 0.009 / 0.214 for linear / sphere / cylinder respectively. The H_0 placebo false-positive
7 rate sits near nominal 5 % in every encoding (8.2 / 4.1 / 6.1 %). The H_1 placebo false-
8 positive rate is consistently above nominal in every encoding (8.2 / 10.2 / 14.3 %;
9 directional consistency, not individually significant at $n = 49$ placebos — see §4.2 for
10 Wilson CI discussion). The H_1 inflation is dimension-asymmetric and encoding-robust
11 — i.e. it persists across encodings that otherwise disagree on the D1 verdict. We discuss
12 what this pattern implies for any positive claim of ensemble-TDA evidence on solar-
13 cyclone perturbation, and propose a minimum protocol of multi-encoding agreement plus
14 dimension-resolved placebo calibration before such claims are advanced.

15 This analysis is observational rather than a controlled experiment. We manipulate the
16 encoding of a periodic coordinate but not the underlying physical system, and the
17 convergence of encoding-comparison findings reported in §4 is observational
18 convergence, not controlled-experiment causal isolation. The methodological claim of the
19 paper rests on the encoding comparison as a within-pipeline robustness test; the
20 substantive physical question (whether solar-cyclone perturbation exists at the ensemble-
21 topology scope) remains open and would require an independent line of evidence beyond
22 the scope of the present study.

23

24 2. Data

25 **Catalog.** IBTrACS v04r01 satellite-era subset (Knapp et al., 2010): 1965-08 to 2025-12,
26 6,661 storms with valid latitude and longitude, 406,059 6-hourly fixes. The catalog is
27 cached locally as `cache/ibtracs_sat.parquet`.

28 **Event selection.** Three events were chosen as one major geomagnetic-storm benchmark
29 per recent solar cycle. Halloween 2003 (peak Dst ≈ -422 nT) is the canonical Cycle 23
30 extreme. St Patrick’s Day 2015 (peak Dst ≈ -223 nT) is the strongest Cycle 24 storm
31 with adequate cross-basin tropical-cyclone activity in the ± 15 -day window. Gannon 2024
32 (peak Dst ≈ -412 nT) is the strongest Cycle 25 storm to date. Other Dst ≤ -300 nT events
33 in the satellite era are excluded for sample-window or independence reasons documented
34 in Supplementary §S4. The three chosen events span three recent solar cycles, sample
35 three different tropical-cyclone seasons, and together provide the only modern-era
36 extreme events with adequate IBTrACS overlap to support an ensemble-TDA test.

37 **Event windows.** Three peak dates anchor the event windows: Halloween storms 2003-
38 10-29 ± 15 d; St Patrick 2015-03-17 ± 15 d; Gannon 2024-05-10 ± 15 d. A storm enters
39 the event-period set if any of its IBTrACS fixes falls inside any of the three windows.
40 Membership is union-based. The resulting set has 25 unique storms across five basins
41 {WP, NI, EP, SI, SP}, with 1,880 6-hourly fixes collectively, and these counts are held



1

1 fixed across all encoding variants reported in §4 so that any across-encoding differences
 2 in test outcome cannot be attributed to differences in which storms are included. The
 3 union rule is deliberate.

4 **Basin composition is uneven across events.** Table E1 reports the basin distribution.
 5 Halloween 2003 is WP-dominated (6 of 11 storms), reflecting the late northern-
 6 hemisphere typhoon season. St Patrick 2015 is SI/SP-dominated (7 of 9), reflecting the
 7 southern-hemisphere late summer. Gannon 2024 is SI-dominated (3 of 5), reflecting the
 8 southern-hemisphere transition month of May. The cross-basin imbalance is intrinsic to
 9 the calendar choices that the three peak dates impose, and any per-event subset (D3 cells
 10 in §4.3) inherits the basin composition of its parent event window.

11 **Table E1.** Basin × event distribution of the 25 event-period storms. Each row counts
 12 unique storms whose lifetime overlaps the ±15-day peak window of the corresponding
 13 event. “Other” captures storms whose IBTrACS BASIN entry uses the legacy SH code or
 14 is missing.

Event	Total	WP	NI	EP	SI	SP	Other
Halloween 2003	11	6	1	1	1	0	2
St Patrick 2015	9	2	0	0	4	3	0
Gannon 2024	5	1	1	0	3	0	0
Union (unique)	25	9	2	1	8	3	2

15 **Control pool.** For each event, the same calendar window in every other year between
 16 1981 and 2024 (excluding the event year ±2 to avoid solar-cycle phase contamination).
 17 Pooled across the three calendar windows, 1,020 unique control storms with 69,454 fixes.

18 Event and control sets are held fixed across all three encodings reported below.
 19 Differences in test outcome are therefore not attributable to event-selection or control-
 20 selection variation.

21

22 3. Methods

23 3.1 Persistent homology: a primer

24 Persistent homology is a tool from topological data analysis (TDA) that summarises the
 25 multi-scale geometric structure of a point cloud (Edelsbrunner & Harer, 2010; Carlsson,
 26 2009; Ghrist, 2008; Wasserman, 2018). Given a finite set of points X in a metric space
 27 and a non-negative scale parameter ϵ , the Vietoris–Rips complex $VR(X, \epsilon)$ is the abstract
 28 simplicial complex whose k -simplices are subsets of $k+1$ points pairwise within distance
 29 ϵ . As ϵ grows from 0 to ∞ , $VR(X, \epsilon)$ grows monotonically. This nested family of
 30 complexes is the Vietoris–Rips filtration. Topological features — connected components
 31 in dimension zero (H_0), one-dimensional loops in dimension one (H_1), higher-
 32 dimensional voids in dimension k — are born at some scale ϵ_{birth} and die at some
 33 larger scale ϵ_{death} . The collection of $(\epsilon_{\text{birth}}, \epsilon_{\text{death}})$ pairs for each feature is the
 34 persistence diagram (PD) of the filtration. Features with large death-minus-birth persist



1

1 across many scales and are interpreted as robust topological structure of the point cloud;
2 features close to the diagonal are short-lived and treated as topological noise. Cohen-
3 Steiner et al. (2007) proved that the PD is stable under perturbations of the input: small
4 changes to X produce only small changes to the diagram. Stability anchors PDs as well-
5 defined summary statistics of point-cloud geometry. Ver Hoef et al. (2023) provide an
6 accessible primer of these ideas in the context of environmental-science applications. We
7 restrict our analysis to H_0 and H_1 because higher-dimensional features become
8 combinatorially expensive without obvious physical interpretation at the present sample
9 sizes.

10 3.2 Distance between persistence diagrams

11 To quantify how different two persistence diagrams are, we use the Wasserstein-2
12 distance W_2 . Given two diagrams P and Q viewed as multisets of points in the extended
13 plane (with the diagonal added to handle features without partner), $W_2(P, Q)$ is the
14 minimum over partial matchings between P and Q of the Euclidean-distance cost
15 summed in quadrature. Mileyko et al. (2011) showed that the space of persistence
16 diagrams equipped with W_2 is a complete and separable metric space. This result
17 licenses the statistical operations we perform on it: bootstrap sampling, computing mean
18 distances to a reference set, and significance testing against a null. We use W_2 rather
19 than the bottleneck distance W_∞ because W_2 accumulates mass over all matched
20 feature pairs, and therefore reflects ensemble-level differences rather than only the worst-
21 case feature. The Wasserstein distance between diagrams is computed with `persim 0.3.8`
22 (Saul et al., 2019).

23 3.3 Three longitude encodings of the point cloud

24 Each storm contributes its 6-hourly fixes as vectors in (latitude, longitude, wind-speed)
25 space. To bring the three dimensions onto comparable footing for a joint Vietoris–Rips
26 filtration, each must be mapped to $[0, 1]$. Latitude is bounded, and may be linearly
27 rescaled with no representational hazard. Wind speed is non-negative and bounded above
28 by physical limits; we divide by 200 kt and clip. Longitude is periodic. Three encodings
29 are tested.

30 3.3.1 Linear modular encoding

31 The simplest choice maps $\text{lon} \in [-180^\circ, 180^\circ]$ to $\text{lon}_n = (\text{lon} + 180) / 360 \in [0, 1]$
32 directly. Latitude maps to $\text{lat}_n = (\text{lat} + 90) / 180$; wind to $\text{wind}_n = \text{WMO_WIND} / 200$
33 (clipped). The point cloud lives in 3-D. The pipeline takes $(\text{lat}_n, \text{lon}_n, \text{wind}_n)$ as joint
34 coordinates in a 3-D box. This encoding has the virtue of dimensional minimality and
35 direct interpretability. Its known defect is that the metric does not respect the periodicity
36 of longitude: two storms at $\text{lon} = -179^\circ$ and $\text{lon} = +179^\circ$ are physically 2° apart on the
37 globe but are placed $358^\circ/360^\circ \approx 0.994$ apart in the embedded representation. For cyclone
38 trajectories that span multi-basin event windows, this discontinuity at the antimeridian is
39 encountered frequently in the data.



1

1 3.3.2 Unit-sphere encoding

2 A more geometrically natural choice maps (lat, lon) to the unit sphere $S^2 \subset \mathbb{R}^3$ via
3 spherical coordinates: $x = (\cos(\text{lat}) \cos(\text{lon}) + 1) / 2$, $y = (\cos(\text{lat}) \sin(\text{lon}) + 1) / 2$, $z =$
4 $(\sin(\text{lat}) + 1) / 2$, each component rescaled to $[0, 1]$. Wind is appended as a fourth
5 coordinate $w_n = \text{WMO_WIND} / 200$ (clipped). The point cloud lives in 4-D, with the
6 geographical position constrained to a 2-D submanifold of $[0, 1]^3$. This encoding handles
7 longitude periodicity exactly (the antimeridian discontinuity vanishes) and avoids the
8 redundancy of a flat $(\cos \text{lon}, \sin \text{lon})$ circular embedding, which would carry an
9 unwanted $\sin^2 + \cos^2 = 1$ constraint that may inflate H_1 noise. Polar distortion is
10 partially absorbed by the sphere geometry. For cyclone tracks, which sit well away from
11 the geographic poles, the sphere is effectively a rescaled equatorial band.

12 3.3.3 Cylinder encoding

13 A third choice puts lat on a linear axis and lon on a flat (cos, sin) circle: $\text{lat}_n = (\text{lat} +$
14 $90) / 180$, $\cos_n = (\cos(\text{lon}) + 1) / 2$, $\sin_n = (\sin(\text{lon}) + 1) / 2$, $w_n = \text{WMO_WIND} / 200$
15 (clipped). The point cloud lives in 4-D, with the geographical position constrained to the
16 surface of a cylinder $S^1 \times \mathbb{R} \subset \mathbb{R}^3$. Like the sphere, this handles longitude periodicity
17 exactly. Unlike the sphere, it decouples latitude from longitude — the joint metric is the
18 L^2 combination of an unconstrained latitude axis with a circular longitude axis. The
19 cylinder is more faithful to the physical decoupling between meridional and zonal storm
20 motion at low latitudes, but it carries a known liability: the $\cos^2 + \sin^2 = 1$ constraint on
21 the longitude pair means that storms spanning a wide longitudinal range can populate a
22 near-circular ring in the (\cos_n, \sin_n) plane, which the Vietoris–Rips filtration is prone
23 to register as a persistent H_1 loop having no physical content beyond the periodicity of
24 longitude itself.

25 3.3.4 Construction and sub-sampling

26 Each storm's 6-hourly fixes are normalized under one of the three encodings above, then
27 all storms in the relevant ensemble are concatenated. The point cloud is sub-sampled
28 uniformly at random to $N_{\text{SUB}} = 600$ points, both to keep ripser computation tractable
29 and to standardize the point count of the filtration across encodings. The seed is set once
30 at script entry (`np.random.default_rng(2026)`) and the RNG is then consumed
31 sequentially by every subsample draw and every resampling iteration within the script, so
32 that no two iterations within a run share the same storm subset and the full end-to-end
33 pipeline is reproducible bitwise on rerun. Sub-sample size sensitivity is reported in §4.4
34 across $N_{\text{SUB}} \in \{400, 600, 800, 1000\}$ for each encoding.

35 3.3.5 Caveats: ambient dimension confound and embedding-induced homology

36 Two caveats apply to the three-encoding comparison and must be carried into any
37 interpretation of the §4 results.

38 First, the three encodings differ not only in how longitude is represented but also in the
39 ambient dimension of the joint metric space. Linear modular places the point cloud in $[0,$
40 $1]^3$; unit-sphere and cylinder both add a fourth coordinate to accommodate the periodic
41 longitude. Differences in test outcome between linear and the two 4-D encodings



1

1 therefore confound longitude representation with ambient-dimension choice. We do not
2 separately rerun linear in a 4-D form (for example by appending a constant zero
3 coordinate or a Gaussian-noise nuisance dimension), which would be the natural follow-
4 up to isolate the encoding effect from the dimensionality effect. A partial mitigation is
5 that the H_0 false-positive rate is calibrated to nominal α in all three encodings
6 (§4.2), which suggests the H_1 encoding-dependence is not a generic curse-of-
7 dimensionality artefact, since a curse-of-dimensionality story would predict H_0
8 calibration to degrade with ambient dimension as well. Nonetheless, the confound is real
9 and the three-encoding comparison should be read with it acknowledged.

10 Second, the sphere and cylinder embeddings carry their own intrinsic topology that the
11 Vietoris–Rips filtration sees. The cylinder $S^1 \times \mathbb{R}$ has non-trivial first homology, $H_1(S^1$
12 $\times \mathbb{R}) = \mathbb{Z}$. A point cloud densely populating the cylindrical surface — as is the case for
13 any ensemble spanning a wide longitudinal range — eventually registers the global
14 longitude loop as a persistent H_1 feature at filtration scales comparable to the cylinder
15 diameter, irrespective of whether the data-generating process carries any actual
16 topological structure beyond periodicity. The unit sphere S^2 has trivial H_1 ($H_1(S^2) =$
17 0), but the wind coordinate appended as a fourth axis creates an $S^2 \times [0, 1]$ geometry
18 whose H_1 behaviour at finite sample sizes is not zero in practice. We do not subtract the
19 embedding-induced topological prior from the observed H_1 features. A more
20 demanding analysis would compute persistence on a null embedding (e.g. random
21 longitudes preserving latitude and wind marginals) under each encoding and report event
22 H_1 distance only after subtracting the embedding’s own H_1 contribution at matched
23 point-cloud density. We defer this to future work. We note that the cylinder’s embedding-
24 induced H_1 is a plausible candidate contributor to the cylinder placebo H_1 inflation
25 reported in §4.2 (14.3 %), rather than to the cylinder D1 H_1 result itself (0.214, a null),
26 since the embedding bias would inflate variance in placebo draws more readily than
27 displace the mean of a single real draw.

28 These caveats notwithstanding, the finding reported in §4 stands. The H_1 perm-p value
29 spans an order of magnitude across three principled encodings on identical data; at least
30 one encoding is demonstrably uncalibrated against placebo (§4.2) by a factor $2.9\times$
31 nominal; and the dimension-asymmetric pattern of H_0 calibrated versus H_1 inflated
32 recurs across encodings, stratification cells (§4.3), and subsample sizes (§4.4). Any
33 successor pipeline must reckon with the diagnostic regardless of whether the specific
34 ambient-dimension and embedding-induced-homology contributions can be cleanly
35 separated from the encoding choice in future work.

36 3.4 H_0 and H_1 features of the storm ensemble

37 H_0 captures the connectivity of the storm-ensemble point cloud at multiple scales. H_1
38 captures one-dimensional loops in the same cloud. Neither dimension has a literal
39 physical counterpart — H_1 features are not closed storm trajectories, and H_0
40 components are not literal storm clusters. They are summary statistics of the joint (lat,
41 lon, wind) geometry that aggregate information across all storms in the ensemble, in a
42 way that mean intensity or mean track length statistics do not. H_0 with the infinite-death
43 component removed reflects the multi-scale connectivity pattern. H_1 reflects whether
44 the ensemble in normalised phase space contains non-convex structure such as loops,



1

1 holes, or recurrence-like features. A perturbation that systematically shifts the joint
2 distribution of storm fixes would change H_0 and H_1 in characteristic ways; whether
3 the present pipeline can read out such shifts robustly is the methodological question the
4 paper investigates.

5 3.5 TDA in atmospheric and cyclone studies: prior applications

6 Within atmospheric science, persistent homology has been used to recognise atmospheric
7 river patterns in satellite-derived water-vapour fields (Muszynski et al., 2019); to quantify
8 a diurnal cycle in tropical-cyclone cloud fields from geostationary infrared imagery
9 (Tymochko et al., 2020); and to characterise weather regimes through the topology of the
10 atmospheric attractor (Strommen et al., 2022). The present application differs from each.
11 Muszynski et al. operate on gridded fields rather than point clouds. Tymochko et
12 al. operate on single-storm imagery in image space. Strommen et al. operate on attractor
13 reconstructions in phase space. Our scope is the joint point cloud of multiple storms'
14 fixes in geographical–intensity space. We are not aware of prior work that has tested the
15 encoding-invariance question explicitly at the cyclone-ensemble TDA scope, which is the
16 gap the present paper addresses.

17 3.6 Test statistic

18 For each encoding, the same downstream statistic is used. The event ensemble's
19 persistence diagram is compared against a bootstrap distribution of control PDs by mean
20 Wasserstein-2 distance from the event PD to a reference set of 30–50 control PDs drawn
21 at random from the full bootstrap pool. The null distribution is the same statistic
22 computed with each control PD substituted for the event PD (jackknifing the reference
23 set). The permutation p-value is the fraction of null draws whose distance equals or
24 exceeds the event distance.

25 3.7 Pipeline runs

26 For each encoding we report four pipeline outputs. The D1 pool fuses all three events: 25
27 event storms vs 1,020 control storms, with 1,000 bootstrap iterations. The 50-placebo
28 calibration uses the same calendar windows (Oct 29 / Mar 17 / May 10) at random years
29 between 1985 and 2020 outside the real event years ± 2 , with 100 bootstrap iterations per
30 placebo (one placebo is skipped for too few storms). Stratified attribution disaggregates
31 the D1 pool into two basin cells (D2_WP, D2_SI) and three single-event cells
32 (D3_Halloween, D3_StPatrick, D3_Gannon), with 500 bootstrap iterations each and
33 Benjamini–Hochberg multiple-comparison correction. Subsample sensitivity reruns the
34 D1 pool at $N_{\text{SUB}} \in \{400, 600, 800, 1000\}$ with 500 bootstrap iterations each.

35 **Resampling unit and method.** In every pipeline run, each resampling iteration draws a
36 sample of n storms uniformly at random *without* replacement from the control pool (n
37 matches the event sample size), then concatenates all 6-hourly fixes belonging to those
38 storms before sub-sampling and Vietoris–Rips filtration. The unit matters. The
39 resampling unit is therefore the storm, not the individual fix; within-storm temporal
40 autocorrelation in the 6-hourly track is preserved by construction, and the iteration-over-
41 iterations distribution reflects between-storm variability under the calendar-matched



1

1 control distribution. We use the loose term “bootstrap” elsewhere in the manuscript for
 2 compatibility with TDA-on-cyclones literature usage, but strictly speaking the procedure
 3 is a storm-level subsample-based permutation null rather than a Bernoulli (with-
 4 replacement) bootstrap in the sense of Efron (1979). We do not resample at the fix level,
 5 which would break the natural dependence structure within a track and inflate the
 6 apparent independence of the resampling. We also do not resample at the storm level *with*
 7 replacement, which would produce duplicated storms in some iterations and equally
 8 inflate apparent independence relative to the calendar-matched control distribution we
 9 wish to test against.

10 3.8 Implementation

11 Python 3.11; numpy 1.26, scipy 1.13, pandas 2.2, matplotlib 3.9, ripser 0.6.14 (Bauer,
 12 2021), persim 0.3.8 (Saul et al., 2019). Vietoris–Rips persistent homology is computed to
 13 dimension 1 via ripser. Otter et al. (2017) survey alternative computational pipelines.
 14 Source scripts (scripts/11_event_ensemble_tda.py, 12_placebo_test.py,
 15 13_stratified_attribution.py, 14_subsample_sensitivity.py) with the three encodings as
 16 switchable variants, and run logs derived/{11,12,13,14}_{4d,sph,cyl}_run.log containing
 17 every numerical result reported below, are catalogued in paper_claims_verified.md (parent
 18 directory).
 19

20 4. Results

21 4.1 D1 pool: signal flips with encoding

22 Table 1 reports the D1 pool result under all three encodings on the same event and
 23 control sets.

24 **Table 1.** D1 pool ensemble TDA under three longitude encodings. Each row: event-PD
 25 Wasserstein-2 distance to the mean control PD; bootstrap null mean over 1,000 control
 26 draws of 25 storms each; permutation p-value as the fraction of null draws whose
 27 distance equals or exceeds the event distance. Sources: derived/11_4d_run.log,
 28 derived/11_sph_run.log, derived/11_cyl_run.log.

Encoding	event H0 dist	null mean H0	perm-p H0	event H1 dist	null mean H1	perm-p H1
Linear modular	2.140	1.641	0.088	0.538	0.492	0.130
Unit-sphere	2.256	1.719	0.082	0.646	0.503	0.009
Cylinder	1.857	1.613	0.190	0.529	0.500	0.214

29 The H_1 perm-p value spans roughly an order of magnitude across the three encodings.
 30 Three numbers, three verdicts. Under the linear modular encoding the H_1 test is non-
 31 significant (perm-p = 0.130); a naive readout would conclude no ensemble-TDA signal.
 32 Under the unit-sphere encoding the H_1 test crosses the conventional $\alpha = 0.05$ threshold
 33 by a wide margin (perm-p = 0.009); a naive readout would conclude a positive signal.
 34 The cylinder tips back to null. Under it the test is again non-significant (perm-p = 0.214),



1

1 with the event H_1 distance sitting essentially at the null mean, a pattern consistent with
 2 the unit-sphere positive being tied to properties of the sphere metric rather than to a
 3 coherent topological signature shared across all three reasonable representations of the
 4 same trajectory data. The H_0 perm-p values are broadly consistent across encodings
 5 (0.082 to 0.190, none crossing 0.05) but vary by a factor of more than two.

6 Three principled encodings of the same data, all of them defensible at face value, give
 7 three qualitatively different verdicts on the D1 pool H_1 result. Before any positive or
 8 negative claim can be sustained, the question must be asked which encoding's verdict is
 9 to be trusted. The placebo calibration in §4.2 begins to address this.

10 4.2 Placebo calibration: H_0 calibrated, H_1 inflated in every encoding

11 The 49-placebo calibration is the same across encodings: 49 fake event-date sets (one of
 12 50 was skipped for too few storms), each a random year in $\{1985, \dots, 2020\}$ excluding
 13 real event years ± 2 , at the same calendar windows. Each placebo runs the full pipeline
 14 with 100 bootstrap iterations. Table 2 reports the false-positive rate at $\alpha = 0.05$ in
 15 both dimensions.

16 **Table 2.** Placebo calibration of the D1 pool pipeline under three longitude encodings.
 17 Same 49 dates per encoding. False-positive rate is the fraction of 49 placebos whose
 18 perm-p falls below 0.05, with the Wilson 95 % CI for nominal 5 % at $n=49$ being
 19 approximately [1.4 %, 14.7 %], so a placebo-rate observation outside that interval is
 20 formally inconsistent with the pipeline holding nominal α at the conventional 5 %
 21 significance level. Sources: derived/12_4d_run.log, derived/12_sph_run.log,
 22 derived/12_cyl_run.log.

Encoding	H_0 false-positive	H_0 verdict	H_1 false-positive	H_1 verdict
Linear modular	4 / 49 = 8.2 %	within Wilson 95 %	4 / 49 = 8.2 %	within Wilson 95 %
Unit-sphere	2 / 49 = 4.1 %	within Wilson 95 %	5 / 49 = 10.2 %	within Wilson 95 %, upper
Cylinder	3 / 49 = 6.1 %	within Wilson 95 %	7 / 49 = 14.3 %	at Wilson 95 % upper bound

23 Three observations. First, the H_0 false-positive rate sits near the nominal 5 % rate in all
 24 three encodings (4.1 % to 8.2 %), with all values inside the Wilson 95 % confidence
 25 interval for nominal 5 %. The H_0 channel of the pipeline is calibrated, encoding-robust.
 26 Second, the H_1 false-positive rate is systematically above nominal 5 % in every
 27 encoding, ranging from 8.2 % to 14.3 %. The factor of inflation grows as the encoding
 28 moves from linear modular (1.6×) to unit-sphere (2.0×) to cylinder (2.9×). Third, the
 29 cylinder H_1 rate sits at the upper bound of the Wilson 95 % CI for nominal 5 % —
 30 under a strict reading, it is borderline incompatible with nominal calibration; under a
 31 permissive reading, it is the worst case of a pattern visible in all three encodings.

32 The pattern is dimension-asymmetric. H_0 holds nominal α encoding-robustly; H_1
 33 fails nominal α in a direction consistent across encodings. The encoding choice
 34 modulates the magnitude of the H_1 inflation but does not eliminate it.



1

1 This has direct implications for the §4.1 D1 verdict. The unit-sphere D1 perm-p_H1 =
 2 0.009 must be read against an H_1 calibration that produces false positives at $\approx 2\times$
 3 nominal. Adjusting upward by the calibration factor ($0.009 \times 10.2/5.0 \approx 0.018$) still
 4 leaves a value below 0.05, so the qualitative signal does not vanish, but the marginal-
 5 significance reading is degraded. More importantly, the calibration finding undermines
 6 the precondition for ever calling a positive D1 H_1 result a clean detection: a positive
 7 result in H_1 from this pipeline rests on a noise channel known to be $1.6\times-2.9\times$ over
 8 nominal.

9 **Three further diagnostics.** First, across the 49 cylinder placebos, the H_1 event-control
 10 distance has $CV \approx 0.50$ (mean 0.248, std 0.125), compared to the H_0 CV of ≈ 0.24
 11 (mean 1.47, std 0.36). The H_1 channel is twice as variable relative to its scale, consistent
 12 with H_1 features in this pipeline being dominated by a small number of high-persistence
 13 loops whose presence-absence varies sharply across event-window draws. Second, the
 14 H_1 null mean itself varies with $CV \approx 0.48$ across placebos, versus $CV \approx 0.12$ for H_0,
 15 confirming that the H_1 null geometry is unstable across event windows in a way the
 16 H_0 null is not. Sources for both diagnostics: derived/placebo_test.csv (cylinder run, n =
 17 49), cross-placebo statistics computed by inline pandas summary at audit time. Third,
 18 with n = 49 placebos and $\alpha = 0.05$ the Wilson 95 % CI for nominal calibration spans [1.4
 19 %, 14.7 %]: the cylinder H_1 rate (14.3 %) sits at the upper bound, and we cannot strictly
 20 reject the null of nominal calibration on n = 49 alone, although the directional
 21 consistency across encodings makes the inflation reading the more parsimonious one. We
 22 recommend $n \geq 200$ placebos for any future replication of this calibration assessment,
 23 which would tighten the Wilson CI to approximately [2.5 %, 9.0 %] at the 5 % nominal
 24 rate and resolve the H_1-specific inflation more sharply.

25 4.3 Stratified attribution: same dimension asymmetry recurs

26 Disaggregating the D1 pool across two basin cells (D2_WP, D2_SI) and three single-
 27 event cells (D3_Halloween, D3_StPatrick, D3_Gannon) gives the picture in Table 3.

28 **Table 3.** Stratified attribution under three longitude encodings. Five cells: D2_WP,
 29 D2_SI, D3_Halloween, D3_StPatrick, D3_Gannon. We report uncorrected perm-p_H1
 30 for each cell, with Benjamini–Hochberg multiple-comparison correction reported in
 31 parentheses next to cells whose uncorrected p falls below 0.10, where the correction
 32 divides the rank-ordered p-values by an FDR-derived constant computed across the five
 33 cells separately within each encoding column rather than pooled across encodings.
 34 Sources: derived/13_4d_run.log, derived/13_sph_run.log, derived/13_cyl_run.log.

Cell	n event	n control	Linear H1	Sphere H1	Cylinder H1
D2_WP	9	240	0.046 (BH 0.140)	0.124 (BH 0.310)	0.058 (BH 0.145)
D2_SI	9	194	0.722	0.894	0.488
D3_Halloween	11	519	0.302	0.238	0.400
D3_StPatrick	9	314	0.056 (BH 0.140)	0.112 (BH 0.310)	0.044 (BH 0.145)
D3_Gannon	5	187	0.704	0.926	0.736



1

1 Two cells (D2_WP, D3_StPatrick) sit near or below the uncorrected $\alpha = 0.05$ threshold
 2 under linear and cylinder encodings, and well above α under sphere. After BH correction
 3 across the five cells, no cell crosses α in any encoding. The cell that returns the strongest
 4 deviation differs by encoding: linear's most-extreme cell is D2_WP at 0.046; sphere's is
 5 D3_Halloween at 0.238 (still null); cylinder's is D3_StPatrick at 0.044. None of these
 6 survives multiple-comparison correction.

7 The same dimension asymmetry recurs. H_0 cell-level perm-p values (omitted from
 8 Table 3 for compactness; see derived/13*_run.log) are above 0.10 in nearly every cell
 9 under every encoding — H_0 is uniformly null at the cell scale. H_1 cell-level perm-p
 10 values straddle α at the uncorrected level in some encoding-cell combinations, but reach
 11 BH significance in none.

12 The pooled D1 H_1 verdict under any single encoding (0.130 / 0.009 / 0.214) is not
 13 localizable to any cell under that encoding's own disaggregation. Combined with the
 14 encoding sensitivity of the pooled D1 perm-p itself, this means the apparent unit-sphere
 15 D1 H_1 signal is not localized to a coherent single-cell source.

16 4.4 Subsample size sensitivity: non-monotonic in every encoding

17 Sub-sampling the event-and-control point cloud to $N_SUB \in \{400, 600, 800, 1000\}$
 18 prior to Vietoris–Rips filtration probes whether the D1 H_1 verdict is sensitive to a fixed
 19 sub-sample size. Table 4 reports the result.

20 **Table 4.** D1 pool H_1 perm-p across subsample sizes under three longitude encodings.
 21 500 bootstrap iterations per cell. Sources: derived/14_4d_run.log, derived/14_sph_run.log,
 22 derived/14_cyl_run.log.

N_SUB	Linear modular	Unit-sphere	Cylinder
400	0.420	0.288	0.302
600	0.124	0.008	0.182
800	0.144	0.004	0.192
1000	0.590	0.050	0.424

23 No encoding yields a monotonic dependence on N_SUB. Under linear, the smallest
 24 perm-p is at N_SUB = 600 (0.124, still null); under sphere, 800 (0.004); under cylinder,
 25 600 (0.182). The N = 1000 row weakens substantially in every encoding, including
 26 sphere where it sits at the boundary of $\alpha = 0.05$. The pattern is consistent with a noise
 27 channel whose realisation depends on which subsample the seed happens to draw, rather
 28 than a stable signal as a function of point-cloud size.

29 The sphere column in Table 4 deserves explicit comment. Read in isolation, the sphere
 30 D1 H_1 = 0.009 (Table 1) and the sphere N_SUB $\in \{600, 800\}$ sub-sample p-values
 31 {0.008, 0.004} (Table 4) jointly invite a positive narrative: low p-values robust to sub-
 32 sampling at mid-range N. Read in the context of Tables 1, 2, 3, 4 jointly, the same
 33 numbers read very differently. The linear and cylinder encodings disagree with sphere on
 34 D1; all three encodings agree that H_1 placebo calibration fails nominally; no encoding
 35 stratifies the pooled signal to a single cell; and the N_SUB = 1000 row degrades sharply



1

1 in every encoding including sphere. The sphere sub-table is a coherent island inside an
2 otherwise incoherent multi-encoding picture. A single-encoding ensemble TDA pipeline
3 that selected sphere a priori would have reported a positive signal that the broader picture
4 does not support.

5

6 5. Discussion

7 The three findings of §4 admit a single interpretation: the ensemble-TDA pipeline as
8 currently configured is fragile in two coupled ways. Two failures. The pipeline is
9 encoding-fragile in H_1 : a defensible variation in how longitude is represented can flip
10 the D1 verdict. And it is dimension-asymmetric in calibration: H_0 holds nominal α
11 across all three encodings tested in §4.2, while H_1 over-rejects by $1.6\times$ to $2.9\times$ across
12 the same encodings, indicating that the H_1 channel of the pipeline returns false positives
13 at a rate substantially above the nominal 5 % under any of the encodings we tested. The
14 first fragility says the pipeline's H_1 verdict on the same data is not single-valued. The
15 second says the H_1 channel that returns those verdicts sits above nominal α to begin
16 with.

17 5.1 The D1 H_1 verdict cannot be read directly off one encoding

18 The standard practice in ensemble-TDA studies on geographical point clouds has been to
19 choose one encoding — typically a linear modular normalization — and report its result
20 as the result. The choice is not innocent. The three encodings in this paper are not exotic
21 variations: each is published, each is physically motivated, and each, applied alone to the
22 same trajectory data, would have produced a paper claiming the answer to the same
23 question with three different verdicts and three different sets of consequences for the
24 literature on solar perturbations of cyclone activity. The linear modular encoding is the
25 dimensional minimum and is what a naive implementation would write first. The unit-
26 sphere embedding is the geometrically faithful choice and is what a methodologically
27 careful implementation would adopt to handle longitude periodicity. The cylinder is the
28 decoupled-axes choice. It is what a physicist would write if asked to keep latitude and
29 longitude as independent observables.

30 None of the three is wrong in itself. But all three disagree on the D1 H_1 perm-p (0.130,
31 0.009, 0.214) for the same event set and the same control pool. The disagreement is not
32 noise around a common value; it is an order-of-magnitude swing that crosses both above
33 and below the conventional $\alpha = 0.05$ line. We do not see how to defend reporting the
34 sphere value (0.009) as the answer without also reporting the linear (0.130) and cylinder
35 (0.214) values, and without doing so under transparent caveat that the answer is
36 encoding-conditional.

37 This may seem like a methodological house-keeping observation. We argue it is more
38 than that. In TDA the stability theorem of Cohen-Steiner et al. (2007) is sometimes read
39 as guaranteeing that small input perturbations produce small output perturbations. That is
40 true under the input metric the theorem refers to. The theorem says nothing about the
41 choice of metric itself. Re-encoding a periodic coordinate (linear vs sphere vs cylinder) is



1

1 not a “small perturbation of the input” — it is a choice of which metric to use. The
2 stability theorem is silent on this choice. The encoding-flip we observe here is consistent
3 with the stability theorem, not a violation of it; it is a reminder that the choice of metric
4 on a periodic geographical coordinate is itself a modelling choice, not a free parameter.

5 5.2 H₁ inflation is intrinsic, not encoding-specific

6 The H₁ placebo over-rejection (Table 2) is the more striking observation. Across all
7 three encodings, false-positive rate at $\alpha = 0.05$ is at 1.6× nominal at the friendliest end
8 (linear, 8.2 %) and at 2.9× nominal at the unfriendliest (cylinder, 14.3 %). In none of the
9 three is H₁ false-positive rate close to the nominal 5 %, and in none of the three does
10 H₀ over-reject in a comparable way.

11 The pattern is consistent with an intrinsic feature of pooled cyclone-trajectory point
12 clouds: H₁ loops in the joint (lat, lon, wind) cloud are sparse and small-effect under the
13 null, and the bootstrap null distribution for them is therefore narrow. A modestly outlying
14 event ensemble can pierce that narrow null even when no physically coherent ensemble-
15 level effect is present. The H₀ channel does not have this property — H₀ is dominated
16 by the global connected component plus a population of short-lived clusters whose null
17 distribution is broader and less easily pierced.

18 The same reading would apply to any pipeline that compares Wasserstein-2 distances
19 between event and control persistence diagrams on this class of data: the relevant statistic
20 is what fraction of the bootstrap null distribution lies below the event value, and the H₁
21 null is too narrow to support nominal calibration. The encoding choice (linear / sphere /
22 cylinder) modulates the geometry of the loops but does not eliminate the narrowness of
23 the null. Hence the H₁ over-rejection persists across encodings.

24 If this reading is correct, then it is not a defect specific to one encoding or one bug to be
25 patched. It is a structural property of the test as applied to cyclone-trajectory data at the
26 present sample sizes. Any future positive H₁ claim from this pipeline must either (a)
27 demonstrate H₁ calibration to nominal alpha under the specific encoding being used, by
28 an explicit placebo test that the authors include in the published analysis, or (b) use a test
29 statistic whose null distribution is independently shown to be wide enough to support the
30 calibration. We are not aware of prior ensemble-TDA cyclone-trajectory work that has
31 met either bar.

32 5.3 The solar–cyclone case study at the ensemble-TDA scope

33 What follows from this for the substantive physical question? Under the linear and
34 cylinder encodings, D1 H₁ is null. Under the sphere encoding, D1 H₁ is nominally
35 below α but the H₁ placebo channel from which that verdict reads is 2.0× over nominal,
36 and the result does not survive stratification or, on a strict reading, survive sub-sample
37 sensitivity past $N_{\text{SUB}} = 800$. Across encodings, no cell-level stratification reaches BH-
38 corrected significance, and Halloween 2003 — the canonical extreme — returns perm-
39 p_{H1} between 0.238 and 0.400 in every encoding, indistinguishable from the null.

40 The cleanest reading we can defend is that the ensemble-TDA pipeline in its present form
41 does not provide adequate evidence either for or against a solar–cyclone perturbation



1

1 effect at the ensemble-topology scope. That is a weaker statement than an earlier single-
2 encoding version of this analysis at the same scope (Dai, in revision) reported under a
3 single encoding choice. The multi-encoding analysis appears to require it.

4 We emphasise the scope qualifier. This null verdict at the ensemble-TDA scope is
5 consistent with both the older positive literature at the per-storm intensity-correlation
6 scope (Elsner & Kavlakov, 2001; Elsner & Jagger, 2008; Hodges & Elsner, 2011) and the
7 null literature at the microphysical CRC scope (Kristjánsson et al., 2008; Calogovic et al.,
8 2010; Laken et al., 2012; IPCC AR6 WGI, 2021). The three scopes interrogate different
9 observational signatures, and a null at one scope does not constitute evidence either way
10 at another. Our finding is specifically that the ensemble-TDA scope, as implemented in
11 the current pipeline, is not yet a tool that can answer the question. It can answer it once an
12 encoding-invariant H₁-calibrated protocol is in place.

13 This conclusion rests on observational convergence across three encoding variants rather
14 than on a controlled experiment. Causal isolation of any physical solar-cyclone effect at
15 the ensemble-topology scope would require an independent experimental or quasi-
16 experimental design — for example, a synthetic-event injection study with known
17 perturbation strength applied to surrogate IBTrACS data — beyond the scope of the
18 present methodological study.

19 5.4 A minimum protocol for future ensemble-TDA on cyclone trajectories

20 We propose three pre-registration requirements for future ensemble-TDA studies that aim
21 to test perturbation hypotheses on cyclone-trajectory data.

22 First, multi-encoding agreement testing. Any positive ensemble-TDA H₁ claim should
23 be reported under at least two principled encodings of the periodic geographical
24 coordinates, with the verdicts compared. A single-encoding claim is not sufficient. Two
25 encodings agreeing is necessary but not sufficient; three encodings agreeing is stronger
26 evidence. Disagreement across encodings is direct evidence that the claim is encoding-
27 dependent and should not be advanced as a positive physical finding without further
28 investigation of why the encodings disagree.

29 Second, dimension-resolved placebo calibration. The placebo set should be large enough
30 to estimate false-positive rate per dimension (H₀ separately from H₁) to within
31 reasonable confidence at the alpha being claimed. At n = 49 placebos and $\alpha = 0.05$, the
32 Wilson 95 % CI for nominal calibration is approximately [1.4 %, 14.7 %], wide enough
33 that values inside it are formally compatible with nominal. A more demanding standard
34 would use n = 200 placebos to narrow the Wilson CI and resolve H₁-specific inflation
35 more sharply.

36 Third, stratified attribution at the natural physical scale. Pooled-significant-but-stratified-
37 null is a recurring failure mode in climate-statistics composites. It belongs to the same
38 family as Simpson's-paradox effects in conventional regression (Simpson, 1951), where a
39 marginal aggregate signal disaggregates into subgroup signals of opposite sign. Pre-
40 registration of a stratification scheme (basin × event × season) before the pooled test
41 result is read protects against post-hoc carving of the pool to support whichever cell
42 happens to come out significant.

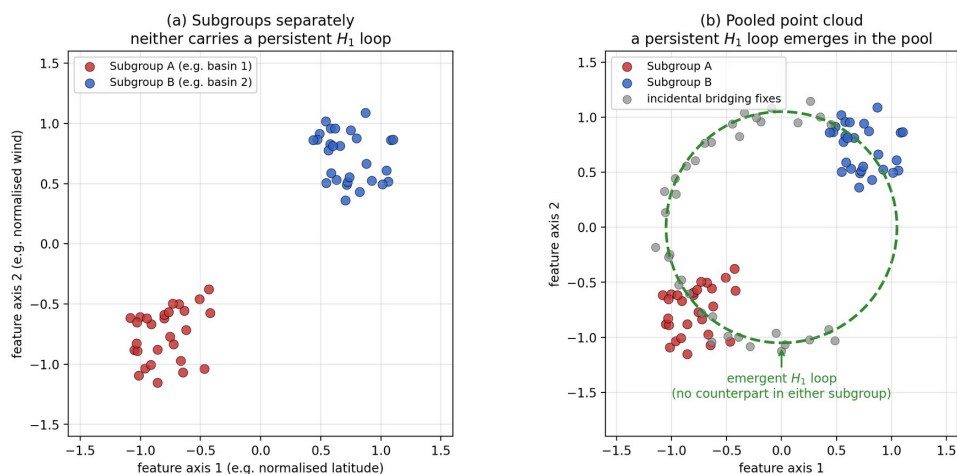
2



1

1 5.5 A note on topological aggregation instability

2 A further point sharpens why the stratification step in our protocol is more essential, not
 3 less, than in regression-based composite analyses. Topology, as encoded in a persistence
 4 diagram, is a *global* summary statistic of the entire point cloud: the birth and death scales
 5 of every connected component and every loop are functions of the joint geometry of all
 6 points simultaneously. Pooling two point clouds is therefore not equivalent to summing
 7 two persistence diagrams. Inserting subgroup points into a shared filtration can create
 8 new loops at scales where neither subgroup alone supported them, and can destroy loops
 9 that each subgroup carried independently. Subgroup cancellation, in the topological
 10 setting, is not linearly inspectable from the pooled diagram alone. We label this failure
 11 mode *topological aggregation instability* (a working term introduced here): a small
 12 misalignment between subgroup geometries, combined with the presence of sparse cross-
 13 subgroup fixes that bridge the gap, can produce a pooled persistence diagram with
 14 features that have no consistent attribution to any subgroup, and that vanish when the
 15 pool is dissected. Figure 1 illustrates the geometry schematically.



16

17 *Figure 1*

18 **Figure 1.** Schematic of topological aggregation instability. Left: two subgroups (red and
 19 blue point clouds), each occupying a distinct region of (lat, lon, wind) phase space,
 20 neither carrying a persistent H_1 loop on its own. Right: the pooled point cloud, in which
 21 a shared Vietoris–Rips filtration generates a persistent H_1 loop spanning the gap
 22 between the two subgroups — a feature that has no counterpart in either subgroup
 23 considered separately, and disappears under disaggregation. Pooling is not a linear
 24 operation on persistence diagrams; subgroup cancellation in the topological setting is not
 25 inspectable from the pooled diagram alone.

26 This figure shows why H_1 specifically is the dimension at risk for aggregation artifacts
 27 in this class of data. Two dimensions, two behaviours. H_0 is less susceptible because
 28 adding cross-subgroup connecting points reduces, not creates, connected-component
 29 features — joining two clusters under a Vietoris–Rips filtration eliminates a feature in the



1

1 H_0 channel rather than opening one, so the H_0 null distribution is broader and less
2 easily pierced by a modest outlier event ensemble. H_1 carries the asymmetry: inserting a
3 bridging point between two subgroup clouds opens a loop that did not exist in either
4 alone. The H_1 inflation observed in our placebo calibration (§4.2) is consistent with this
5 geometric picture as one contributing mechanism.

6 We use the label *topological aggregation instability* as a working descriptive term for the
7 observed failure mode rather than a formal theoretical concept. A rigorous treatment —
8 for example, a theorem characterizing the conditions under which pooling two
9 contractible point clouds connected by sparse bridge points produces a persistent H_1
10 feature in the union that has no counterpart in either subgroup — is left to future work.
11 The figure schematic and the §4.3 stratified-attribution null pattern are evidence
12 consistent with the label, not proof of a general principle. We retain the working label in
13 the paper because the cell-level null pattern of §4.3 is informative for future studies even
14 without a formal characterization, and the label gives a working vocabulary for the failure
15 mode that any subsequent theoretical treatment can refine, reject, or replace.

16 5.6 Pathways toward H_1-calibrated ensemble TDA

17 The minimum protocol of §5.4 constrains what should be reported but does not fix the
18 underlying H_1 inflation diagnosed in §4.2 and §5.2. Three pathways toward a fix appear
19 worth attempting. We do not attempt any of them in the present paper. We sketch each so
20 that a follow-up study — by us or others — can pick one and report whether it meets the
21 calibration requirement that the current pipeline does not.

22 The first pathway replaces the test statistic. Wasserstein-2 between persistence diagrams
23 produces a null whose narrowness in H_1 reflects, in part, the W_2 metric concentrating
24 mass on a small number of high-persistence features and assigning negligible weight to
25 the dense low-persistence tail. Alternative summary statistics distribute mass differently.
26 Persistence landscapes (Bubenik, 2015) convert each PD into a function on the real line,
27 and L²-distances on the resulting function space have a different null geometry that may
28 be wider in the relevant H_1 regime. Persistence images (Adams et al., 2017) provide a
29 vectorized PD representation amenable to kernel methods on point clouds (Reininghaus
30 et al., 2015), with the kernel bandwidth offering a tunable handle on null variance. A
31 simpler first-pass attempt would normalize the H_1 statistic itself: dividing the event H_1
32 distance by an estimate of the control bootstrap standard deviation transforms a raw
33 distance into a z-score and may widen the null without changing the underlying TDA
34 framework.

35 The second pathway enlarges the event sample from the 25 storms used here to all
36 satellite-era extreme geomagnetic events meeting a $K_p \geq 8$ threshold (approximately 50
37 events). Sample size alone is not the cure. Narrowness of the H_1 null is a shape
38 property; a larger event sample may not by itself widen the null, but it raises the bar a
39 modest outlier ensemble must clear to register against the null, and a 2× sample increase
40 would tighten the Wilson 95 % CI on the placebo false-positive rate by approximately
41 $\sqrt{2}$, allowing finer resolution of any residual encoding-specific calibration bias.



1

1 The third pathway replaces the Vietoris–Rips persistence framework entirely. Mapper-
2 graph homology, cubical persistence on gridded fields, and witness complexes each carry
3 different sensitivity to point-cloud aggregation. The cost is re-implementing the full
4 pipeline at the framework level. The payoff is independence from the specific Vietoris–
5 Rips H_1 inflation diagnosed here, which may be tied to a single computational choice
6 rather than to ensemble TDA on cyclone trajectories in general.

7 We frame the present paper as a methodological diagnostic that stands on its own. A
8 natural follow-up would attempt pathway 1 (H_1 statistic normalization, then persistence
9 landscapes, then persistence-image kernels) on the same data, with the goal of meeting
10 the placebo H_1 calibration target of nominal 5 %; conditional on that follow-up reaching
11 nominal calibration, a further study could return to the substantive solar–cyclone question
12 with a calibrated tool. The present paper does not commit to a specific timetable for
13 either follow-up step. We report the present diagnosis now because it has standalone
14 methodological value: it warns any current user of the existing pipeline against a positive
15 H_1 claim and motivates the search for a fix by us or by other groups.

16

17 6. Conclusions

18 We applied an ensemble-TDA pipeline to test whether three benchmark extreme
19 geomagnetic events (Halloween 2003, St Patrick 2015, Gannon 2024) perturb the
20 topology of the cyclone activity in their immediate ± 15 -day windows. The result is
21 mixed. Three principled longitude encodings of the (lat, lon, wind) point cloud — linear
22 modular, unit-sphere, and lat-linear-plus-longitude-cylinder — give three qualitatively
23 different $D1 H_1$ perm-p values (0.130, 0.009, 0.214) on the same event set and the same
24 control pool, demonstrating that the pooled-ensemble test verdict in this pipeline depends
25 on the choice of representation at a level that crosses both above and below the
26 conventional $\alpha = 0.05$ line of demarcation between null and positive results in this
27 literature. A 49-placebo calibration shows that the H_0 false-positive rate sits near the
28 nominal 5 % in every encoding (8.2 %, 4.1 %, 6.1 %), while the H_1 false-positive rate is
29 consistently above nominal in every encoding (8.2 %, 10.2 %, 14.3 %; directional
30 consistency across encodings, with cylinder at the Wilson 95 % upper bound at $n = 49$).
31 Stratified attribution and subsample-size sensitivity recur the same pattern: H_1 results
32 that disagree across encodings, no cell-level signal surviving BH correction in any
33 encoding, no monotonic dependence on N_{SUB} . The pipeline is encoding-fragile in H_1
34 and dimension-asymmetric in calibration. Under these constraints we cannot defend
35 either a positive or negative substantive claim about ensemble-topology evidence for or
36 against the solar–cyclone perturbation hypothesis at this scope. We frame the
37 contribution as a methodological cautionary tale: pooled point-cloud TDA on geospatial
38 trajectories with periodic coordinates is more fragile than its formal stability theorems
39 suggest. We propose a minimum protocol — multi-encoding agreement testing plus
40 dimension-resolved placebo calibration plus pre-registered stratified attribution — as a
41 precondition for any positive ensemble-TDA claim in this class of problems. Pathways
42 toward an H_1 -calibrated successor pipeline are sketched in §5.6 as starting points for
43 any successor study.

2



1

1

2 Code and data availability

3 IBTrACS v04r01 is publicly available at
4 <https://www.ncei.noaa.gov/products/international-best-track-archive>. Analysis scripts
5 (scripts/11_event_ensemble_tda.py, 12_placebo_test.py, 13_stratified_attribution.py,
6 14_subsample_sensitivity.py) with the three encodings as switchable variants
7 (normalize_pc_sphere, normalize_pc_cylinder, and the in-line linear modular form), and
8 derived run logs (derived/{11,12,13,14}_{4d,sph,cyl}_run.log) reproducing every numerical
9 value reported in Tables 1–4 and Table E1, are available from the corresponding author
10 on reasonable request.

11 Author contributions

12 RD designed the study, implemented the three-encoding pipeline, executed the bootstrap
13 and placebo runs, performed the stratified attribution and subsample-sensitivity analyses,
14 drafted the manuscript, and revised it for submission. RD is the sole author and accepts
15 responsibility for the content.

16 Competing interests

17 The author declares no competing interests.

18 Acknowledgements

19 Analysis was performed using Anthropic Claude Code as a coding assistant, with all
20 design decisions, statistical interpretations, encoding-comparison framing, and writing by
21 the author. The longitude-periodicity bug in the v1–v3 single-encoding pipeline was
22 identified through an internal pre-submission audit of the v3 manuscript; the diagnosis
23 motivated the three-encoding analysis reported here. IBTrACS data are provided by
24 NOAA NCEI.

25 References

- 26 Adams, H., Emerson, T., Kirby, M., Neville, R., Peterson, C., Shipman, P.,
27 Chepushtanova, S., Hanson, E., Motta, F., & Ziegelmeier, L. (2017). Persistence images:
28 a stable vector representation of persistent homology. *J. Mach. Learn. Res.*, 18, 1–35.
- 29 Bauer, U. (2021). Ripser: efficient computation of Vietoris-Rips persistence barcodes. *J.*
30 *Appl. Comput. Topol.*, 5, 391–423.
- 31 Bubenik, P. (2015). Statistical topological data analysis using persistence landscapes. *J.*
32 *Mach. Learn. Res.*, 16, 77–102.
- 33 Calogovic, J., Albert, C., Arnold, F., Beer, J., Desorgher, L., & Flueckiger, E. O. (2010).
34 Sudden cosmic ray decreases: no change of global cloud cover. *Geophys. Res. Lett.*, 37,
35 L03802.
- 36 Carlsson, G. (2009). Topology and data. *Bull. Amer. Math. Soc.*, 46(2), 255–308.



1

- 1 Cohen-Steiner, D., Edelsbrunner, H., & Harer, J. (2007). Stability of persistence
2 diagrams. *Discrete Comput. Geom.*, 37(1), 103–120.
- 3 Edelsbrunner, H., & Harer, J. L. (2010). *Computational Topology: An Introduction*.
4 American Mathematical Society, Providence, RI.
- 5 Elsner, J. B., & Jagger, T. H. (2008). United States and Caribbean tropical cyclone
6 activity related to the solar cycle. *Geophys. Res. Lett.*, 35, L18705.
- 7 Elsner, J. B., & Kavlakov, S. P. (2001). Hurricane intensity changes associated with
8 geomagnetic variation. *Atmos. Sci. Lett.*, 2, 86–93.
- 9 Ghrist, R. (2008). Barcodes: the persistent topology of data. *Bull. Amer. Math. Soc.*,
10 45(1), 61–75.
- 11 Ver Hoef, L., Adams, H., King, E. J., & Ebert-Uphoff, I. (2023). A primer on topological
12 data analysis to support image analysis tasks in environmental science. *Artif. Intell. Earth*
13 *Syst.*, 2(1), AIES-D-22-0039.1.
- 14 Hodges, R. E., & Elsner, J. B. (2011). Evidence linking solar variability with US
15 hurricanes. *Int. J. Climatol.*, 31, 1897–1907.
- 16 Hodges, R. E., Jagger, T. H., & Elsner, J. B. (2014). The sun–hurricane connection:
17 diagnosing the solar impacts on hurricane frequency over the North Atlantic basin using a
18 space–time model. *Nat. Hazards*, 73, 1063–1081.
- 19 Huth, R., & Stryhal, J. (2024). Association of the 11-year solar cycle with correlation and
20 teleconnection structures in tropospheric circulation. *J. Space Weather Space Clim.*, 14,
21 31.
- 22 IPCC AR6 WGI (2021). *Climate Change 2021: The Physical Science Basis. Contribution*
23 *of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on*
24 *Climate Change* [Masson-Delmotte, V., et al. (eds.)], Cambridge University Press,
25 Chapter 7.
- 26 Kirkby, J., Curtius, J., Almeida, J., et al. (2011). Role of sulphuric acid, ammonia and
27 galactic cosmic rays in atmospheric aerosol nucleation. *Nature*, 476, 429–433.
- 28 Knapp, K. R., Kruk, M. C., Levinson, D. H., Diamond, H. J., & Neumann, C. J. (2010).
29 The International Best Track Archive for Climate Stewardship (IBTrACS). *Bull. Amer.*
30 *Meteor. Soc.*, 91(3), 363–376.
- 31 Kristjánsson, J. E., Stjern, C. W., Stordal, F., Fjæraa, A. M., Myhre, G., & Jónasson, K.
32 (2008). Cosmic rays, cloud condensation nuclei and clouds — a reassessment using
33 MODIS data. *Atmos. Chem. Phys.*, 8, 7373–7387.
- 34 Laken, B. A., Pallé, E., Calogovic, J., & Dunne, E. M. (2012). A cosmic ray–climate link
35 and cloud observations. *J. Space Weather Space Clim.*, 2, A18.



1

- 1 Li, D., Xiao, Z., Xu, J., & Zhao, L. (2023). Asymmetric modulation of solar activity on
2 tropical cyclone frequency over the western North Pacific and the possible mechanism.
3 *Front. Earth Sci.*, 11, 1169685.
- 4 Marsh, N. D., & Svensmark, H. (2000). Low cloud properties influenced by cosmic rays.
5 *Phys. Rev. Lett.*, 85, 5004–5007.
- 6 Mendoza, B., & Pazos, M. (2009). A 22-yr hurricane cycle and its relation with
7 geomagnetic activity. *J. Atmos. Sol.-Terr. Phys.*, 71, 2047–2054.
- 8 Mileyko, Y., Mukherjee, S., & Harer, J. (2011). Probability measures on the space of
9 persistence diagrams. *Inverse Problems*, 27, 124007.
- 10 Mironova, I. A., et al. (2018). Rapid intensification of tropical cyclones in the context of
11 the solar wind–magnetosphere–ionosphere–atmosphere coupling. *J. Atmos. Sol.-Terr.*
12 *Phys.*, 182, 1–4.
- 13 Muszynski, G., Kashinath, K., Kurlin, V., Wehner, M., & Prabhat (2019). Topological
14 data analysis and machine learning for recognizing atmospheric river patterns in large
15 climate datasets. *Geosci. Model Dev.*, 12, 613–628.
- 16 Nayak, S. (2024). Effects of solar variability on tropical cyclone activity. *Earth Space*
17 *Sci.*, 11, e2023EA003500.
- 18 Otter, N., Porter, M. A., Tillmann, U., Grindrod, P., & Harrington, H. A. (2017). A
19 roadmap for the computation of persistent homology. *EPJ Data Sci.*, 6, 17.
- 20 Pérez-Peraza, J., Velasco, V., & Kavлакov, S. (2008). Wavelet coherence analysis of
21 Atlantic hurricanes and cosmic rays. *Geofísica Internacional*, 47(3), 231–244.
- 22 Pierce, J. R. (2017). Cosmic rays, aerosols, clouds, and climate: recent findings from the
23 CLOUD experiment. *J. Geophys. Res. Atmos.*, 122, 8051–8055.
- 24 Reininghaus, J., Huber, S., Bauer, U., & Kwitt, R. (2015). A stable multi-scale kernel for
25 topological machine learning. *Proc. IEEE Conf. Computer Vision and Pattern*
26 *Recognition (CVPR)*, 4741–4748. doi:10.1109/CVPR.2015.7299106.
- 27 Saul, N., Tralie, C., et al. (2019). Scikit-TDA: topological data analysis for Python.
28 <https://scikit-tda.org/>.
- 29 Simpson, E. H. (1951). The interpretation of interaction in contingency tables. *J. R. Stat.*
30 *Soc. B*, 13(2), 238–241.
- 31 Strommen, K., Chantry, M., Dorrington, J., & Otter, N. (2022). A topological perspective
32 on weather regimes. *Climate Dynamics*, 60, 1415–1434.
- 33 Svensmark, H. (2007). Cosmoclimatology: a new theory emerges. *Astron. Geophys.*,
34 48(1), 1.18–1.24.
- 35 Svensmark, J., Enghoff, M. B., Shaviv, N. J., & Svensmark, H. (2016). The response of
36 clouds and aerosols to cosmic ray decreases. *J. Geophys. Res. Space Phys.*, 121, 8152–
37 8181.

2

22



1

- 1 Tymochko, S., Munch, E., Dunion, J., Corbosiero, K., & Torn, R. (2020). Using
- 2 persistent homology to quantify a diurnal cycle in Hurricane Felix. *Pattern Recognit.*
- 3 *Lett.*, 133, 137–143.
- 4 Wasserman, L. (2018). Topological data analysis. *Annu. Rev. Stat. Appl.*, 5, 501–532.
- 5 Yu, F., & Luo, G. (2014). Effect of solar variations on particle formation and cloud
- 6 condensation nuclei. *Environ. Res. Lett.*, 9, 045004.