UNNS Substrate Research Program | Turbulence Branch | Frozen Replication Record

Reproducible Scale–Time Commutation Structure in Forced Isotropic Turbulence

Instruments: STRUC-ROUTE-I v0.1.2 · ROUTE_LADDER_EXTRACT · STRUC-I v1.0.4 · STRUC-PERC-I v2.5.0 · STITCH-MECH v0.1.1

Data: Johns Hopkins Turbulence Database, `isotropic1024coarse` forced-isotropic velocity field

Samples: two operationally independent (preregistered, non-overlapping) 256^3 cutouts of the same DNS realization (Pilot A, Pilot B)

Preregistration: eight frozen criteria R1–R8 · Classification: **Structural Partial Replication**

Archive: Zenodo DOI 10.5281/zenodo.22650769 · Date: 2026

Abstract

Reproducibility of structural claims about turbulent flow organization is rarely tested against an independent physical sample rather than an independent statistical resample of the same data. We report a preregistered two-pilot replication study of a scale–time routing signature in forced isotropic turbulence, using velocity fields from the Johns Hopkins Turbulence Database (`isotropic1024coarse`). Pilot A (frozen baseline) and Pilot B (this replication) are disjoint 256^3 spatial cutouts of the same simulation with disjoint 10-frame stored time windows: zero shared spatial grid points, zero shared stored frames. Both samples were processed through an unmodified pipeline — a physical adapter, STRUC-ROUTE-I routing analysis, deterministic ladder extraction, STRUC-I perturbative admissibility analysis, STRUC-PERC-I percolation analysis, and STITCH-MECH harder-null mechanism testing — and scored against eight criteria (R1–R8) preregistered before Pilot B acquisition.

The central quantity is the scale–time stitching defect $D_\square = \mathrm{JSD}(\mathrm{ST}, \mathrm{TS})/\ln 2$, the Jensen–Shannon divergence between scale-then-time and time-then-scale structural routing. In Pilot B, the real population mean is $D_\square = 0.014172$, against 0.996336 under a degree-preserving null (N0), 0.326204 under a distance-stratified null (N1), and 0.211452 under a distance-*and*-feature-stratified null (N2, the principal mechanism target). Pilot A shows the same qualitative hierarchy ($0.026628 \to 0.996786 \to 0.329091 \to 0.223649$), and the real/null separation is in fact *larger* in Pilot B (N1/real = 23.02×, N2/real= 14.92×) than in Pilot A (12.36×, 8.40×). Six of eight preregistered criteria pass in Pilot B (R1, R2, R4, R5, R6, R8); two fail (R3: STRUC-I $D_\square$ admissibility reaches only STRUCTURAL BOUNDARY rather than GEOMETRIC PERSISTENCE; R7: the frozen upper-decile physical-intensity localization test is degenerate because more than 90% of eligible objects sit at exactly $D_\square = 0$). The frozen classification is **Structural Partial Replication**.

We interpret this as evidence for a reproducible transformation grammar in which scale transformation and temporal evolution are, for most structural objects, nearly interchangeable in order — a largely commuting scale–time regime whose sparse non-commuting exceptions localize preferentially around branching events, in both samples. The route mechanism reproduces; the scalar admissibility strength of the same ladder and the physical-intensity localization of its extreme tail do not. This separates two questions that a single pilot could not distinguish: route-mechanism reproducibility versus full cross-chamber and localization reproducibility. The raw source cutouts and provenance/checksum records for both samples are archived at Zenodo (DOI: 10.5281/zenodo.22650769); frozen chamber result records are described with their hashes in Sec-

tion 10 and held in the project repository. A third preregistered sample (Pilot C) is proposed to adjudicate the R3/R7 divergence; it is future work, not a precondition for the present claims.

Contents

1 Introduction

1.1 The reproducibility problem in structural turbulence signatures

Turbulent flow is routinely described through statistical quantities — energy spectra, structure functions, intermittency exponents [2, 3, 4] — whose reproducibility across independent realizations is well established. Claims about *structural organization*, by contrast, are harder to replicate in the ordinary sense. A structural claim ("coherent structures persist across scale," "vortical objects nest in a particular hierarchy") is typically evaluated once, on one simulation snapshot or one experimental run, and its statistical significance is assessed against a null model computed *from that same sample.* This tests whether the observed structure is unlikely under a null, but it does not test whether the same structure reappears in a second, physically independent sample.

This is not a minor distinction. A result can be highly significant against an in-sample null while still being an artifact of that particular sample's boundary conditions, forcing realization, or finite-time sampling. The appropriate check is a preregistered replication: a second sample, selected by rule before its data are inspected, processed through the identical frozen pipeline, and scored against criteria fixed in advance.

1.2 Prior result: a scale–time stitching defect

A predecessor study ("Pilot A," frozen prior to this work) analyzed a 256^3 cutout of Johns Hopkins Turbulence Database (JHTDB) forced isotropic turbulence (`isotropic1024coarse`) and reported a routing signature we term *scale–time commutation*: the ancestry of a structural object constructed by first coarsening scale and then advancing time ("ST") substantially agrees with the ancestry constructed by first advancing time and then coarsening scale ("TS"). The disagreement between these two routes is quantified by a stitching defect

$$D_{\square} \;=\; \frac{\mathrm{JSD}(\mathrm{ST},\mathrm{TS})}{\ln 2} \;\in\; [0,1], \tag{1}$$

the normalized Jensen–Shannon divergence [10] between the ST and TS routing distributions for a given structural object. $D_{\square} = 0$ indicates perfect route agreement (commutation); $D_{\square} = 1$ indicates maximal disagreement.

Pilot A's headline result was that the real population mean of $D_{\square}$ is far below what a degree-preserving null predicts, that this gap survives progressively harder graph-local geometric controls, and that the ladder forms a well-connected structural population under percolation analysis. Pilot A alone, however, could not distinguish a reproducible physical mechanism from a single-sample artifact.

1.3 This work: a preregistered independent-sample replication

We report Pilot B: a second 256^3 JHTDB cutout, selected by a preregistered rule to share zero spatial grid points and zero stored time frames with Pilot A, processed through the identical frozen pipeline, and scored against eight criteria (R1–R8) fixed in a preregistration document before Pilot B's data were acquired. This manuscript reports that frozen, completed record.

Central claim. *A low scale–time stitching defect and a constrained-routing mechanism reproducibly emerge in two operationally independent forced-isotropic JHTDB samples and survive increasingly local graph-geometric null controls, while scalar cross-chamber admissi-*

bility and physical-intensity localization show sample dependence.

This claim is deliberately narrower than "turbulence is structurally organized." It separates *route-mechanism reproducibility* from *full cross-chamber and localization reproducibility* — a distinction that a single pilot cannot express, because a single pilot has no second sample against which to ask whether every component of its signature reappears.

What is new. The novelty here is not a new vortex taxonomy, nor the recovery of a known spectral scaling law. It is the replicated observation that two operations which need not commute in general — coarse-graining across scale and advancing time — behave, at the level of structural routing, as if they approximately commute for most structural objects in forced isotropic turbulence, with departures from that compatibility concentrated around branching events. Coherent-structure and intermittency research have long documented multiscale organization [6, 5] and scale-dependent departures from self-similarity [4]; commutation errors between filtering and differentiation are likewise a known concern in the large-eddy simulation literature [8]. What this study adds is a preregistered, criterion-by-criterion replication showing that a specific, quantitative route-commutation signature — not merely a qualitative multiscale regularity — reappears across two operationally independent samples of the same flow.

Organization. Section 2 describes the JHTDB source, the independence design, the frozen pipeline, the null hierarchy, and the R1–R8 preregistration. Section 3 reports the complete Pilot B result set. Section 4 gives the direct Pilot A↔Pilot B criterion-by-criterion comparison. Section 5 presents the N0→N1→N2→real null hierarchy. Section 6 examines the two failing criteria, R3 and R7, in detail. Section 7 interprets the two-pilot record. Section 8 states limitations. Section 9 proposes a preregistered Pilot C. Section 10 documents data and reproducibility artifacts, and Section 11 concludes.

2 Methods

2.1 Data source

Both samples are drawn from the JHTDB `isotropic1024coarse` dataset, a forced isotropic turbulence direct numerical simulation on a 1024^3 periodic grid with grid spacing $\Delta x = 2\pi/1024 \approx 0.0061359$. The two pilots were acquired by different routes, both yielding the same validated HDF5 cutout format: Pilot A's frozen baseline cutout was obtained from a Hugging Face mirror of the dataset, while Pilot B's cutout was retrieved directly via the JHTDB `getCutout` service (authenticated, registered token; serial z-slab queries, slab depth 30). Acquisition identity for each sample is fixed by source hash (Section 10). Although the underlying DNS domain is periodic, extracted 256^3 cutouts are *not* treated as periodic computational domains: no opposite-face wrapping is applied, derivatives are boundary-safe, and boundary-truncated structures are rejected or explicitly controlled.

2.2 Independent-sample design

Pilot A (frozen baseline) uses the cutout $x, y, z \in$ [1:256] (1-based), stored frames [1:10], stored $\Delta t = 0.002$, physical time window $[0.000, 0.018]$.

Pilot B was selected under a preregistered independence rule requiring zero shared spatial grid points and, where possible, zero shared stored frames with Pilot A. The realized selection is a half-domain translation of the Pilot A origin in x, y, and z: cutout $x, y, z \in [513{:}768]$ (1-based; [512:767] zero-based), stored frames [501:510], physical time window $[1.000, 1.018]$. Table 1 records the independence audit. The selection was hash-locked (SHA-256 `385cca0a...a40875e`) before acquisition, and the acquired source file (2,015,302,728 bytes, 1.8769 GiB) was hashed after download (SHA-256 `977e6ab3...ca0b969f`); HDF5 structural validation passed against all expected `Velocity_0501`–`0510` datasets, shape $256 \times 256 \times 256 \times 3$, axis order `zyxc`.

Throughout this manuscript we use *operationally independent* to denote exactly this preregistered zero-overlap relation between two samples: disjoint spatial grid points and disjoint stored time frames, audited before either sample is used in downstream analysis. Pilot A and Pilot B are two operationally independent, preregistered, non-overlapping samples of the *same* DNS realization — not two separate DNS realizations, and not statistically independent draws in the sense of independent turbulent flow instances. All subsequent uses of "independent" in this manuscript refer to this operational sense unless stated otherwise.

Criterion 2.1 (Independence rule)**.** A sample may be labeled Pilot B only if it is physically independent of Pilot A to the extent supported by the source: zero spatial grid-point overlap is required; zero stored-frame overlap is preferred. A same-cube time subset or a spatial subdivision of the Pilot A cube would be labeled `jhtdb_pilot_a_internal_time` / `..._internal_space`, not Pilot B.

2.3 Frozen analysis pipeline

Both samples are processed by the identical, version-frozen pipeline:

1. **Physical adapter** (`JHTDB_ROUTE_ADAPTER v0.1.0`): converts the raw velocity cutout into a population of structural objects and candidate scale/time relations, at frozen scale factors $\{1, 2, 4, 8, 16\}$ (full reproducibility detail in Section 2.4).
2. **STRUC-ROUTE-I v0.1.2**: constructs a directed, layered scale/time graph over the object population and evaluates route persistence and the stitching defect $D_\square$ against a degree-preserving endpoint-swap null (N0).
3. **ROUTE_LADDER_EXTRACT**: deterministically extracts three one-column ladders from the frozen STRUC-ROUTE-I output — `D_STITCH` (per-object stitching defect on stitching-eligible sources), `P_TIME` (temporal route persistence), and `P_SCALE` (scale route persistence) — preserving duplicates, applying no jitter, smoothing, normalization, or rescaling.
4. **STRUC-I v1.0.4**: perturbative admissibility analysis of each ladder (defined below).
5. **STRUC-PERC-I v2.5.0**: percolation/connectivity analysis of each ladder (defined below).
6. **STITCH-MECH v0.1.1**: a harder mechanism-test null hierarchy (N1, N2; defined below) and object-level physical/restructuring association tests.

No frozen chamber, adapter, or threshold is modified between Pilot A and Pilot B. Sample-specific execution wrappers were used only to redirect I/O paths and sample identity; the frozen scientific configuration (thresholds, null counts, seeds, block factors, segmentation rules, mechanism permutation counts, regime boundaries) is unchanged.

2.4 The physical adapter: object segmentation and candidate relations

This subsection gives the reproducibility detail an external reader needs to reconstruct the object population and the ST/TS routing distributions from the raw velocity field, beyond the summary in Section 2.3. All values below are frozen in the adapter's per-pilot configuration file and are identical for Pilot A and Pilot B except sample identity and I/O paths.

Scale construction. Coarse-graining is exact, non-overlapping block averaging, not a spectral or opaque server-side filter: a block factor f replaces every $f \times f \times f$ native-grid block by its mean velocity and reduces the grid resolution accordingly, at the frozen factors $\{1, 2, 4, 8, 16\}$.

Derived fields. From the (coarse-grained) velocity field the adapter computes, via boundary-safe finite differences (no FFT derivatives, no periodicity assumed): divergence, vorticity, enstrophy, strain-squared $\|S\|^2$, the Q-criterion, and helicity.

Boundary policy. The 256^3 cutout is not treated as periodic. A fixed physical margin of 16 native-grid cells is excluded from object detection, and any object touching the retained interior boundary is discarded as truncated.

Object segmentation. Objects are connected components of the rotation-dominated region

$$Q \;\geq\; 1.5 \times \mathrm{RMS}(Q)$$

on the boundary-safe interior, using 26-connectivity, with a minimum native-volume criterion of 64 voxels, subject to a floor of 2 voxels on the coarse grid. These are two separate frozen constraints, not a single equivalence: at coarsening factor 16 a coarse voxel represents $16^3 = 4096$ native voxels, so the coarse-grid floor of 2 voxels is the binding constraint at that scale, not a restatement of the native-grid threshold. The multiplier 1.5 and the connectivity/size thresholds are frozen in the adapter configuration and are not retuned between pilots or scales. Figure 3 shows an example: slice intersections of segmented 3D objects on one native-resolution slice of the actual Pilot B velocity field, computed with this exact rule.

Candidate relations and eligibility. Only positive spatial overlap creates a candidate relation. Across scale, the coarser object's label field is expanded exactly onto the finer block grid before overlap is counted; across time, same-scale label fields at adjacent stored frames are compared directly. For each candidate pair, overlap_src = |intersection|/|source|, overlap_dst = |intersection|/|destination|, and IoU = |intersection|/|union|; the relation `confidence` used as the routing edge weight is overlap_src. A candidate is accepted as a relation when *any* of overlap_src $\geq$ 0.10, overlap_dst $\geq$ 0.10, or IoU $\geq$ 0.05 holds. The routing graph built from the accepted relations is defined next (Section 2.5).

Object physical weight. Each object is assigned integrated enstrophy $w = \int(|\omega|^2/2)\, dV$, a positive, extensive object-level quantity, not asserted to be conserved along accepted relations. It does *not* enter the construction of the confidence-weighted $S_{s,t}/T_{s,t}$ routing distributions defined in Section 2.5; those are built entirely from the edge-level `confidence`

(= overlap_src) introduced above. Integrated/mean enstrophy is instead retained downstream as a physical-intensity descriptor and tested for association with $D_\square$ in STITCH-MECH (Section 2.8).

2.5 STRUC-ROUTE-I: routing model and stitching defect

STRUC-ROUTE-I's primitive object is a directed, layered graph $\mathcal{G} = (V, E_s, E_t)$ over the adapter's structural-object population, where E_s contains adjacent-scale relations and E_t contains adjacent-time relations accepted under the eligibility rule above. Edges are restricted to strict adjacency: a scale edge requires Δscale_idx $= +1$, Δtime_idx $= 0$, and a time edge requires Δtime_idx $= +1$, Δscale_idx $= 0$; an edge violating adjacency is a fatal configuration error, not a silently dropped relation.

For a source object at scale-time cell (s, t), let $S_{s,t}$ and $T_{s,t}$ denote its (confidence-weighted, normalized) outgoing routing distributions along the scale and time axes respectively. The two-step routing distributions compared for the stitching defect are

$$P_{\mathrm{ST}} = S_{s,t}\, T_{s+1,t} \qquad \text{versus} \qquad P_{\mathrm{TS}} = T_{s,t}\, S_{s,t+1}\,:$$

scale-then-time advances scale first and then time from the resulting cell; time-then-scale advances time first and then scale from the resulting cell. $D_\square$ (Eq. 1) is computed for every source object for which both two-step distributions exist, i.e. for which the object has eligible outgoing relations on both paths. Figure 2 shows this square of routes schematically.

The null model ("N0") performs directed, layer-degree-preserving endpoint swaps: an ensemble of graphs is generated that preserves each object's axis, transition layer, edge count, and source out-/destination in-degree, while rewiring *which* objects are connected via repeated directed endpoint swaps. Scale persistence, time persistence, and $D_\square$ are each tested one-sided for favorability against this ensemble (z_{fav}, p_{fav}, $\alpha = 0.05$, minimum 20 nulls, 100 null graphs at 5.0 swap attempts per edge).

Criterion 2.2 (Verdict rule)**.** CONSTRAINED_ROUTING is assigned when ≥ 2 of the three inferential tests (scale persistence, time persistence, $D_\square$) are favorable and 0 are opposing, under a null ensemble passing quality gates (mean mobility fraction ≥ 0.01, unique-graph fraction ≥ 0.1, real-graph match fraction ≤ 0.9).

2.6 STRUC-I: perturbative admissibility

STRUC-I tests whether a one-column ladder L remains admissible under uniform perturbation $\delta_i \sim \mathrm{Uniform}[-\varepsilon, \varepsilon]$ at scale $\varepsilon = \kappa \cdot \mathrm{median(gaps)}$, via

$$\mathrm{inv}(P_\varepsilon;\, L) \;\leq\; \nu(V_\varepsilon(L))\,, \tag{2}$$

evaluated at 40 κ steps over $\kappa \in [0.01, 1]$. The reported admissibility coefficient $A_\kappa \in [0, 1]$ is averaged ($\langle A_\kappa \rangle$) and minimized ($\min A_\kappa$) across steps. Regime boundaries (frozen, shared with the companion charge-routing chamber program): GEOMETRIC PERSISTENCE for $\langle A_\kappa \rangle \gtrsim 0.93$, else STRUCTURAL BOUNDARY; below 0.60, STRUCTURAL BOUNDARY degrades further to STRUCTURAL INSTABILITY. States within GEOMETRIC PERSISTENCE are STABLE STRUCTURE ($\langle A_\kappa \rangle = \min A_\kappa = 1.000$), BOUNDARY-STABILIZED ($\langle A_\kappa \rangle \gtrsim 0.975$), WEAK PERSISTENCE ($0.93 \lesssim \langle A_\kappa \rangle < 0.975$); within STRUCTURAL BOUNDARY, TRANSITIONAL STRUCTURE; within STRUCTURAL INSTABILITY, RANDOM STRUCTURE.

Criterion 2.3 (Precision rerun rule)**.** If $\langle A_\kappa \rangle$ lies within 0.01 of a regime boundary, the ladder is rerun at 10,000 Monte Carlo perturbations with all other settings unchanged.

2.7 STRUC-PERC-I: percolation

STRUC-PERC-I constructs a graph over ladder rows, connecting values within a κ-scaled tolerance, and reports the giant-component ratio (GR), isolated-node count, and connectivity threshold κ_{conn}. FULL_PERCOLATION requires GR = 1.000 and isolated = 0; otherwise HARD_FRAGMENTATION.

2.8 STITCH-MECH: harder null hierarchy

STITCH-MECH independently reproduces the real $D_\square$ mean from first principles (geometry recomputation) and evaluates it against two harder, nested null families beyond the STRUC-ROUTE-I N0:

- **N1**: distance-stratified, layer-preserving degree-swap null (100 replicates; block factors and stratification bins frozen).
- **N2**: distance-*and*-feature-similarity-stratified, layer-preserving degree-swap null (100 replicates; the principal mechanism-survival target).

Each null ensemble is quality-gated on mean mobility fraction, unique-graph fraction, and real-graph match fraction, identically to N0. Every frozen quantity in this section — distance-bin count, feature-bin count, swap count, seed, and permutation count — is recorded exactly as configured, not retuned between pilots:

- **N1** additionally preserves each edge's within-transition normalized-distance quantile stratum (5 distance bins; 100 null graphs; seed `20260905`; 5.0 swap attempts per edge).
- **N2** additionally preserves each edge's within-transition feature-similarity stratum on top of N1's distance stratum (4 distance bins $\times$ 4 feature bins; 100 null graphs; seed `20261905`; 5.0 swap attempts per edge). This is the harder local-geometry-*and*-state control and the principal mechanism-survival target (R6).

Both null families use the same directed, degree/layer-preserving swap kernel as N0 (Section 2.5), with the additional stratum constraint on which edges may be swapped; pair-specific distance and feature-similarity values are recomputed after each accepted swap.

STITCH-MECH additionally computes object-level associations between $D_\square$ and two physical-intensity proxies: mean enstrophy, and a dissipation proxy $2\nu \times$ (mean enstrophy $- 2 \times$ mean Q-criterion) $= 2\nu \times$ mean strain-squared $\|S\|^2$, where the frozen kinematic viscosity $\nu = 0.000185$ is carried over from the Pilot A adapter configuration (it does not alter ranks or any STRUC-ROUTE-I metric). This is algebraically the conventional strain-based dissipation form $\epsilon = 2\nu S_{ij} S_{ij}$; we retain the term "proxy" because it is reconstructed from the adapter's boundary-safe, finite-difference, coarse-grained velocity-gradient field (Section 2.4) rather than measured on a native-resolution DNS dissipation field — the reconstruction, not the formula, is what is approximate. The mean strain-squared $\|S\|^2$ appearing here is the same quantity used in adapter segmentation. Both proxies are tested raw (Spearman) and partial-rank (conditioned on scale, time, outgoing overlap geometry, IoU, normalized distance, feature similarity, and branch/merge state), and upper-decile $D_\square$ enrichment in each proxy is tested via stratified permutation (5000 permutations, seed `20260905`).

2.9 Preregistered criteria R1–R8

Eight criteria were preregistered, before Pilot B acquisition, for criterion-by-criterion (not weighted-composite) evaluation:

R1 — Route organization.
STRUC-ROUTE-I returns CONSTRAINED_ROUTING under a usable primary null ensemble.

R2 — Primary N0 stitching contrast.
Real mean $D_\square$ is lower than N0 with $p_{\text{fav}} < 0.05$.

R3 — STRUC-I.
`D_STITCH` reaches GEOMETRIC PERSISTENCE or stronger under canonical settings (with the precision-rerun rule applied where boundary-adjacent).

R4 — STRUC-PERC-I.
`D_STITCH` reaches giant ratio ≥ 0.95 with ≥ 100 gap vertices after deduplication.

R5 — N1.
N1 is usable and real $D_\square <$ N1 mean with $p_{\text{fav}} < 0.05$.

R6 — N2 (principal mechanism target).
N2 is usable and real $D_\square <$ N2 mean with $p_{\text{fav}} < 0.05$.

R7 — Physical localization.
The top-decile $D_\square$ tail is enriched (enrichment > 1, stratified permutation $p < 0.05$) in top-decile enstrophy and in top-decile dissipation proxy.

R8 — Restructuring localization.
At least one preregistered branch/merge class shows significantly higher $D_\square$ ($p < 0.05$).

Interpretation labels are fixed in advance: **Core Replication** (R1–R6 all pass); **Structural Partial Replication** (R2 passes and at least one of R5/R6 passes, but one or more cross-chamber conditions fail); **Route-Only Replication** (R1, R2 pass but both N1/N2 mechanism-survival criteria fail or are underresolved); **No Replication** (R2 fails under a usable N0 ensemble). `Underresolved` is reserved for unusable null ensembles and is never converted to a pass or fail after the fact.

Stopping rule. *Pilot B is one preregistered physical replication sample. New cutouts are not selected under the Pilot B label until one reproduces Pilot A; if Pilot B diverges, the result is frozen and interpreted, and any further sample is separately preregistered as Pilot C or later.*

3 Pilot-B Replication Results

3.1 Adapter output and corpus scale

The physical adapter produced 10,515 structural objects and 14,169 relation candidates (14,118 eligible) from the Pilot B cutout, organized into 5 scale layers and 10 time layers, with 4,830 scale-eligible source nodes and 8,806 time-eligible source nodes. Of these, 3,962 objects are stitching-eligible (have both a well-formed ST and TS route).

3.2 Primary STRUC-ROUTE-I result

The primary run (Pilot B, all five scale factors) returns **CONSTRAINED_ROUTING** (tags: `PERSISTENT_TIME`, `LOW_STITCH_DEFECT`), with 100/100 nulls accepted, mean null mobility 0.998859, 100 unique null graphs, and zero real-graph null matches. Two of three inferential tests are favorable; zero are opposing (Table 2).

Scale persistence is not significant ($p_{\text{fav}} = 0.7426$); time persistence and the stitching defect are both strongly favorable ($p_{\text{fav}} = 0.0099$ for each). The real stitching defect (0.014172) sits dramatically below its null mean (0.996336), $z_{\text{fav}} = 868.9$.

Per-scale-index breakdown (Table 3) shows the eligible population collapsing sharply above the $\times 4$ scale gap (2,609 $\rightarrow$ 1,234 $\rightarrow$ 114 $\rightarrow$ 5 eligible objects across the four populated scale indices), with the sparsest layer ($\times 8 \rightarrow \times 16$) contributing negligibly. Per-time-index breakdown across the 9 consecutive frame transitions (Table 4) shows the low stitching defect distributed evenly across the 10-frame window rather than concentrated at one time boundary (range 0.0085–0.0206).

3.3 Scale sensitivity [1,2,4,8]

A preregistered robustness run excluding the sparsest scale factor (16) returns CONSTRAINED_ROUTING again, with materially unchanged qualitative reads across all three inferential tests (Table 5). This robustness result is *not* R2 itself — R2 is the primary N0 contrast — but supports that the primary verdict is not driven by the sparsest scale layer.

3.4 Ladder extraction

Deterministic extraction from the frozen STRUC-ROUTE-I output yields `D_STITCH` ($n =$ 3,962), `P_TIME` ($n =$ 8,806), and `P_SCALE` ($n =$ 4,830), matching the adapter's eligible-source counts exactly, with no deduplication, jitter, smoothing, normalization, or rescaling.

3.5 STRUC-I result

Table 6 reports the primary (2,000 Monte Carlo) and precision (10,000 Monte Carlo) STRUC-I results. `D_STITCH` reaches $\langle A_\kappa \rangle = 0.851863$, classified STRUCTURAL BOUNDARY/TRANSITIONAL STRUCTURE — 0.078 below the GEOMETRIC PERSISTENCE boundary (not boundary-adjacent; no rerun triggered). `P_SCALE` and `P_TIME` both fall to STRUCTURAL INSTABILITY/RANDOM STRUCTURE. `P_TIME` sits only 0.0028 below the 0.60 STRUCTURAL INSTABILITY/STRUCTURAL BOUNDARY line, triggering the frozen precision rule; the 10,000-MC rerun changes the mean by < 0.001 and leaves the classification unchanged (0.598235, 0.0018 below the line).

3.6 STRUC-PERC-I result

All three ladders reach FULL_PERCOLATION (Table 7). `D_STITCH` clears the preregistered R4 minimum decisively: giant ratio 1.000, isolated 0, and 336 gap vertices after deduplication (minimum required: 100).

3.7 STITCH-MECH result

STITCH-MECH independently reproduces the real $D_\square$ mean exactly (0.014171635400793336; geometry-recomputation maximum errors $\leq 10^{-16}$ for both distance-norm and feature-similarity checks) and evaluates the harder N1/N2 null hierarchy (full detail in Section 5). Both nulls pass all quality gates with no flags; the mechanistic verdict is **SURVIVES_LOCAL_GEOMETRY_CONTROL**.

4 Pilot A ↔ Pilot B: Direct Comparison

Table 8 gives the frozen R1–R8 record. Six of eight criteria pass: R1, R2, R4, R5, R6, R8. Two fail: R3, R7. The frozen classification is **Structural Partial Replication** — R2 passes and both R5 and R6 pass (the mechanism-survival criteria), while the cross-chamber R3 condition fails, so Pilot B does not qualify as **Core Replication** (which requires R1–R6 all pass).

Table 9 maps each criterion onto both pilots directly. R1, R2, R4, R5, and R6 replicate cleanly. R7 fails outright in Pilot B (Section 6). R8 replicates in a narrower form: Pilot A showed significant elevated $D_{\square}$ for scale branching, time branching, time merging, and any restructuring; Pilot B replicates only the two branching contrasts.

Finding 4.1. All five preregistered route-mechanism criteria — R1, R2, R4, R5, and R6 — pass cleanly across two preregistered, non-overlapping samples of the same forced-isotropic DNS realization. This is the strongest possible outcome for the mechanism itself within a two-pilot design: it is not merely non-contradicted, it is positively reproduced under a preregistered, criterion-by-criterion protocol.

5 Mechanism Result: The N0→N1→N2→Real Hierarchy

5.1 The collapse

The defining quantitative fact of both pilots is the same four-point hierarchy (Figure 4, Table 10): the degree-preserving null (N0) predicts $D_{\square}$ near its maximum (≈ 0.996 in both pilots); the distance-stratified null (N1) predicts a large but partial reduction (≈ 0.33 in both); the distance-*and*-feature-stratified null (N2) predicts a further reduction (≈ 0.21–0.22); and the real population sits an order of magnitude below N2 in both samples.

$$\underbrace{0.9963}_{\text{N0}} \longrightarrow \underbrace{0.3262}_{\text{N1}} \longrightarrow \underbrace{0.2115}_{\text{N2}} \longrightarrow \underbrace{0.0142}_{\text{REAL}} \qquad (\textit{Pilot B}) \tag{3}$$

Local graph geometry (distance stratification in N1; distance + feature similarity in N2) therefore explains part of the naive N0 contrast, but a large residual — the real system is 93.3% below N2 in Pilot B, and 88.1% below N2 in Pilot A — remains unexplained by the tested local controls.

5.2 The separation is not weaker in the replication

Table 10 gives the direct numeric comparison, shown visually in Figure 6. The null hierarchy itself is nearly sample-invariant (N0 differs by 0.045% between pilots; N1 by 0.88%; N2 by 5.45%), while the real defect is 46.8% lower in Pilot B than in Pilot A. Consequently the real/null separation ratios are *larger* in Pilot B: N1/real = 23.02× (Pilot A: 12.36×) and N2/real = 14.92× (Pilot A: 8.40×).

Finding 5.1. The route mechanism does not merely survive replication; it reproduces with a *stronger* real/null separation in the independent sample. Because the null hierarchy itself stays close across pilots while the real defect differs, the reproducible object is not a fixed scalar value of $D_{\square}$ but the relationship between the real defect and its local-geometry-controlled null — the routing architecture, not a magnitude.

6 Divergences: R3 and R7

6.1 R3: cross-chamber admissibility does not reach Geometric Persistence

Pilot A's `D_STITCH` ladder reached Geometric Persistence/Weak Persistence ($\langle A_\kappa \rangle = 0.971708$); Pilot B's reaches only Structural Boundary/Transitional Structure ($\langle A_\kappa \rangle = 0.851863$). This is a clean, substantive divergence, not a boundary-adjacent near-miss: Pilot B's value sits 0.078 below the 0.93 threshold, well outside the 0.01 precision-rerun band. We record this as a genuine failure of the strict scalar-admissibility criterion, without qualification.

6.2 R7: the frozen selector is degenerate, not merely non-significant

Pilot B's $D_\square$ distribution has an unusually large atom at exactly zero: the frozen 90th-percentile threshold evaluates to 0.0 because more than 90% of the 3,962 eligible objects lie at $D_\square = 0$. The nominal "top-decile" tail is therefore the entire eligible population (3,962/3,962), producing enrichment exactly 1.0× and $p = 1.0$ for both enstrophy and the dissipation proxy (Table 13). This is recorded as **FAIL**, not `UNDERRESOLVED` — the frozen protocol reserves the latter label for unusable N1/N2 tests, not for a degenerate quantile selector on an otherwise usable population. No post-hoc tail redefinition is applied to rescue the criterion.

Descriptively, however, alternative tail definitions — never substituted for the preregistered test, but computed for interpretive context — recover a modest, consistent positive enrichment (Table 14): this holds both for fixed-fraction tails (top-5%, top-1%) and for fixed-$D_\square$ threshold tails ($D_\square \geq 0.1$, $D_\square \geq 0.5$), 1.13–1.32× across all four, roughly half Pilot A's effect size (2.50× enstrophy, 1.87× dissipation). Ordinary (unconditioned) rank association also remains positive ($\rho(D_\square, \text{enstrophy}) = 0.1021$; $\rho(D_\square, \text{dissipation}) = 0.1266$), as does the geometry-conditioned partial-rank association ($\rho_{\text{partial}} = 0.1027$ and 0.0494 respectively), weaker than Pilot A's 0.179 and 0.120 but not absent.

Finding 6.1. R7 fails on a selector artifact interacting with a genuinely different $D_\square$ distribution shape, not on a null result for every tail definition tested. The frozen preregistered localization test does not replicate in Pilot B, and that failure stands as recorded. Descriptive physical-intensity associations under alternative, non-preregistered tail definitions remain positive but are weaker and selector-sensitive; they are reported as context, not as a substitute verdict. This is exactly the kind of sample-dependence a preregistered replication is designed to expose rather than paper over.

6.3 R8: restructuring localization narrows

Table 15 gives the full branch/merge contrast. Scale branching and time branching both replicate at $p = 0.00020$; scale merging, time merging, and any-restructuring — all significant in Pilot A — do not replicate in Pilot B ($p = 1.0$ for each). R8 passes on the frozen rule (at least one preregistered class significant), but the replicated statement narrows: elevated $D_\square$ localizes specifically around branching events in both pilots, not around merging or aggregate restructuring. Figure 8 shows this localization directly as marginal empirical CDFs rather than only a median contrast: across most of the $D_\square$ range, the non-branching population's empirical CDF lies above the branching population's in both pilots, with the curves converging in the far tail as both approach 1 — i.e. branching objects are systematically shifted toward larger $D_\square$ across the distribution, not only at the median.

7 Interpretation

7.1 A transformation grammar, not a catalogue of shapes

The very low $D_{\square}$ says something more specific than "turbulent structures persist." It says that, for most structural objects, the route obtained by transforming scale first and then advancing time is highly compatible with the route obtained by advancing time first and then transforming scale. Scale evolution and temporal evolution are, for most of the population, nearly interchangeable in order. What reproduces across two independent samples is this relational property — not a universal vortex shape, not a fixed catalogue of morphological species, but a *transformation grammar* governing how structural evolution across scale and structural evolution across time relate to one another.

Principle 7.1 (Scale–time commutation)**.** Turbulent structural evolution occupies a largely commuting scale–time route regime: for most structural objects, transforming scale and advancing time are mutually compatible operations under the routing representation, and this compatibility — not a fixed scalar admissibility value — is the reproducible object.

7.2 Why scale–time commutation matters physically

Coarse-graining (changing scale) and temporal evolution (advancing time) are, in general, operations that need not commute in a multiscale dynamical system: applying them in different orders can in principle produce different structural descriptions, and the size of that disagreement is itself informative about how strongly scales are coupled. A large, generic order-sensitivity would say that a structural object's identity depends on *when* in the scale/time processing pipeline it is examined — a discouraging property for any attempt to track structures across a cascade. What both pilots show instead is that, under the routing representation constructed here, order-sensitivity is the exception rather than the rule: for most structural objects, coarse-graining before advancing time yields nearly the same structural routing as advancing time before coarse-graining. That is a nontrivial empirical property of a multiscale turbulent system, independent of and prior to any derivation of it from the Navier–Stokes equations. It constrains how a future first-principles account would need to behave: it would have to explain why route disagreement is small and bulk-dominated rather than explain away a large disagreement.

7.3 Six concrete properties of the structure

The two-pilot record supports six specific, falsifiable statements about the discovered object:

1. **Primarily relational, not merely geometric.** No universal vortex shape is claimed. What reproduces is a relationship between two structural transformations (scale-then-time vs. time-then-scale), not a morphological template. Pilot A and Pilot B have materially different scalar STRUC-I strengths for `D_STITCH` (Geometric Persistence vs. Structural Boundary), yet the routing mechanism itself reproduces.
2. **Largely commuting.** Most of the turbulent structural evolution sits in a regime where changing scale and advancing time are compatible operations, under the routing representation constructed here. This is the two-pilot formulation of Principle 7.1.
3. **Failures of commutation are localized, not uniform.** Elevated $D_{\square}$ repeatedly appears around branching events in both pilots. Pilot B does not reproduce every merging association seen in Pilot A, but scale branching and time branching replicate significantly

in both.

4. **Highly connected under the value-space criterion.** The `D_STITCH` ladder forms a substantive giant component under the STRUC-PERC-I value-space connectivity criterion in both pilots (Pilot A: giant ratio 0.996933, 652 gap vertices; Pilot B: giant ratio 1.000, 336 gap vertices). This is a statement about connectivity among numerically proximate ladder values, not a claim about spatial percolation of turbulent objects in the flow: the low-defect population is not a collection of isolated coincidences in the ladder representation, but no claim is made here about a percolating spatial vortex network.

5. **Reproducible while its scalar strength is variable.** Pilot A's `D_STITCH` ladder reaches Geometric Persistence; Pilot B's reaches only Structural Boundary. The robust object is therefore not a fixed scalar admissibility value — it is the routing architecture itself, which reproduces even as its perturbative admissibility strength differs between samples.

6. **Not reducible to the tested local geometry.** N1 and N2 constrain distance and feature similarity, yet the real system remains dramatically more commuting than either control, in both pilots (Section 5). This is arguably the single most important mechanistic fact obtained by this study.

7.4 Scientific Gain, Structural Interpretation, and Relation to Turbulence Theory

The preceding sections establish *that* the route mechanism replicates and *where* it does not. This subsection makes explicit what that buys scientifically: the two-pilot experiment reveals several layers of organization that were invisible after Pilot A alone, and a single-sentence summary ("Pilot B replicated Pilot A") would understate what the comparison actually shows.

7.4.1 What we have actually gained

The first gain is a new empirical object. The reproducible quantity is not a particular vortex geometry, a particular value of $D_\square$, or even a particular STRUC-I regime. What survives independent sampling is a *relationship between transformations*: scale-then-time and time-then-scale structural routing are far more compatible in the real flow representation than in progressively constrained null graphs. Pilot B gives

$$D_\square^{\mathrm{N0}} = 0.9963 \;\rightarrow\; D_\square^{\mathrm{N1}} = 0.3262 \;\rightarrow\; D_\square^{\mathrm{N2}} = 0.2115 \;\rightarrow\; D_\square^{\mathrm{real}} = 0.0142,$$

and Pilot A exhibits the same ordering. Crucially, the null hierarchy stays close across pilots while the real $D_\square$ changes substantially: Pilot B's real defect is 46.8% below Pilot A's (Table 10). The study therefore shows that the reproducible feature is an *architecture of separation from admissible controls*, not a universal numerical value. Before Pilot B, a low $D_\square$ could have been a peculiarity of one cutout; after Pilot B, the evidence supports a reproducible scale–time route relation within the frozen representation — already a conceptual advance over what a single pilot could show.

The second gain is that the phenomenon has been partially decomposed. N0 shows that degree structure alone cannot account for it. N1 shows that adding local spatial-distance structure explains a substantial fraction of the naive N0 contrast. N2 shows that adding feature similarity explains more still. Yet the real population remains 93.3% below N2 in Pilot B. The result is therefore not simply REAL $\neq$ RANDOM; it is closer to

$$\text{degree} \;\rightarrow\; \text{local geometry} \;\rightarrow\; \text{feature compatibility} \;\rightarrow\; \boxed{\text{remaining route organization}}.$$

That residual is presently unexplained. Identifying it as a residual — rather than absorbing it into a single significance statement — is one of the principal scientific products of this study.

7.4.2 A revelation from the failed criteria

R3 and R7 did more than lower the replication label: they show what does *not* constitute the stable object.

R3 shows that the routing architecture and scalar perturbative admissibility are separable. Pilot A's `D_STITCH` ladder reaches Geometric Persistence, whereas Pilot B's reaches only Structural Boundary; nevertheless, the N0/N1/N2 routing mechanism replicates strongly in both. This means

$$\boxed{\text{route organization} \;\neq\; \text{fixed scalar robustness}}.$$

That is an actual discovery generated by replication: Pilot A alone gave no way to know that these two properties could separate.

R7 yields an equally interesting lesson. More than 90% of Pilot B's stitch-eligible objects have representation-level $D_{\square} = 0$, which is what collapses the preregistered 90th-percentile selector. The failed test must remain a failure, but the distribution (Figure 7) supports a useful empirical decomposition,

$$\text{population} \;=\; \underbrace{\text{commuting bulk}}_{\text{very large}} + \underbrace{\text{noncommuting defect set}}_{\text{sparse}}.$$

We are deliberately careful with the word "exact" here: $D_{\square} = 0$ means exact agreement *under our routing representation*, not mathematical equality of filtered Navier–Stokes operators. The zero-atom is nonetheless a striking property of that representation.

R8 tells us something about the defect set itself: its departures are not spread uniformly through the population. Scale branching and time branching carry significantly elevated $D_{\square}$ in both pilots, while Pilot B does not reproduce Pilot A's merging contrasts. This gives a considerably more specific structural picture than a generic claim that "turbulence contains coherent structures":

$$\boxed{\text{dominant route compatibility} \;+\; \text{sparse branching-associated failures}}.$$

7.4.3 What the discovery actually is

We state the discovery carefully:

The discovery, stated precisely. *Within a frozen multiscale structural representation of forced-isotropic DNS, structural routing across scale and time exhibits a reproducible near-commutation architecture: the real scale–time stitching defect lies far below progressively geometry-constrained graph nulls in two non-overlapping samples, while the comparatively rare failures of route compatibility preferentially localize around branching events.*

That claim is fully supported by the present experiment. We would *not* yet write "scale evolution and time evolution commute in the Navier–Stokes equations" — that has not been demonstrated. The difference matters: the first is an empirical discovery about a specific, frozen structural representation; the second would be a mathematical or physical statement about the governing equations themselves, requiring a different level of evidence than a two-sample instrument replication can provide.

7.4.4 Why this matters conceptually

Classical turbulence analysis typically asks how amplitudes, moments, energy fluxes, or correlations change with scale [2, 3]. Our question is different: *does the order in which a turbulent structure is viewed across scale and evolved through time matter to its structural ancestry?* That turns turbulence into an operator-order problem as well as a scaling problem. The answer from these two samples is, on its face, unexpected: under the chosen structural representation, order matters remarkably little for most of the population. There was no *a priori* requirement that scale-then-time and time-then-scale routing should produce similar structural distributions — the null models show precisely how easily that agreement disappears once the observed relations are reorganized (Section 5). The new object is therefore not merely structure; it is *compatibility between structural transformations.*

7.4.5 Relation to established turbulence theories

The relation to existing theory is one of complementarity, not competition (Table 12).

The connection to Germano's filtering approach [7] is worth developing specifically. Germano formulated turbulence in terms of algebraic relations among filtered representations at different resolutions, rather than only mean/fluctuation decomposition. Our work is conceptually adjacent but distinct: we do not derive filtered Navier–Stokes identities; we ask empirically whether a structural-routing representation remains compatible when scale transformation and temporal evolution are performed in opposite orders. This is a legitimate operatorial tradition to sit alongside, without any claim that $D_{\square}$ is a Germano identity.

Relation to cascade theory. Kolmogorov-type theories describe statistical scale similarity and energy transfer [2, 3]. She–Leveque and related intermittency approaches [4, 5] place particular emphasis on hierarchies of coarse-grained dissipation and rare singular structures. Our result asks an orthogonal question: instead of measuring only how a quantity changes as $\ell \to \ell/2 \to \ell/4$, we ask whether a scale change remains structurally compatible with temporal evolution, i.e. whether $S(T(X))$ and $T(S(X))$ agree, for a structural object X. That introduces an additional dimension into cascade organization: not only what happens across scale, but whether the scale trajectory remains compatible with the temporal trajectory. If future work shows that high-$D_{\square}$ branching events coincide with anomalous structure-function contributions or bursts of coarse-grained dissipation, the sparse noncommuting population could become a structural interpretation of intermittency. At present that connection is a testable hypothesis (Section 7.5), not a result.

Relation to stochastic cascade models. There is a suggestive connection to the Friedrich–Peinke description of turbulence as a stochastic process in scale [11], which uses conditional probability distributions and Fokker–Planck dynamics to represent the cascade as transitions among velocity increments across scales. Our framework may eventually be viewed as complementary: a Markov-in-scale cascade characterizes transition statistics $P(X_{\ell_2} \mid X_{\ell_1})$ through scale, while route commutation compares $P(\mathrm{ST})$ against $P(\mathrm{TS})$ — whether scale transitions are compatible with temporal evolution, not only how they transition. This could eventually connect the UNNS notion of a "cascade grammar" to an established stochastic-cascade literature, but that comparison has not been carried out here.

Relation to coherent-structure and dynamical-systems theories. Coherent-structure approaches have long argued that turbulent flows are not featureless randomness. Exact-coherent-structure (ECS) research [12] treats unstable invariant solutions as a scaffold organizing turbulent trajectories in state space; Lagrangian coherent-structure (LCS) research [13] seeks robust organizing skeletons behind particle trajectories. Our result is distinct from both: we have not found an invariant solution, and we have not found a universal vortex morphology. We have instead found evidence for what might be called a *transformation-space skeleton*: many different objects, shapes, and intensities nevertheless obey a similar relation between what happens when scale is changed and what happens when time advances. In short: ECS organize *states*; LCS organize *transport*; this result organizes *transformations*.

7.4.6 The most important implication

If the phenomenon survives stronger tests, the implication is larger than this first paper. High-dimensional turbulent evolution might possess constraints that are easier to detect in relations *between* transformations than in instantaneous fields themselves. In other words, a highly unpredictable velocity field does not necessarily imply equally unconstrained structural-transformation relations. That was close to the hypothesis with which this research branch began; this study does not yet establish the full "nested admissible families" hypothesis, but it gives a first empirical piece of evidence for a more modest version of it:

The real intellectual gain. *There appears to be a reproducible rule governing how scale and time transformations relate, within the tested representation.*

7.4.7 A possible deeper revelation: a hierarchy of structural invariance

The Pilot A/Pilot B divergence pattern suggests a hierarchy of structural properties by how strongly they replicate:

route architecture $>$ connectivity $>$ scalar admissibility strength $>$ localization detail.

The first two replicate strongly; the third changes materially; the fourth is sample-dependent. This suggests that turbulent structure may possess *levels* of invariance, with relational properties more stable across samples than scalar or localization properties. This was not something the experiment was designed specifically to discover: it emerged because the replication partially failed. We regard it as one of the manuscript's strongest interpretive points — replication did not merely confirm or reject Pilot A; it resolved the original signature into a stable core and a variable periphery.

Stable core	**Variable periphery**
CONSTRAINED_ROUTING verdict	STRUC-I regime strength (GEOMETRIC PERSISTENCE vs. STRUCTURAL BOUNDARY)
Low $D_\square$ vs. N0	Extreme-tail intensity enrichment (R7)
N1/N2 survival (R5, R6)	Merging localization (R8)
Branching localization (R8)	—

7.4.8 What it could eventually imply

Contingent on the phenomenon surviving different Reynolds numbers, flow classes, field-level controls, and alternative representations (Section 8.1), several further consequences become testable, though none is demonstrated by the present study:

1. **Prediction.** If noncommutation rises *before* branching, $D_\square$ could become an early structural indicator of reorganization. We have shown association with branching, not temporal prediction.
2. **Intermittency.** If the sparse noncommuting set carries disproportionate dissipation, enstrophy, or anomalous high-order structure-function weight, it could provide a structural mechanism behind intermittent corrections to otherwise regular cascade behavior.
3. **Model reduction.** If the dominant turbulent population occupies a strongly constrained scale–time route architecture, one might model allowed structural transitions without predicting every velocity degree of freedom.
4. **Multiscale closure.** Germano-style coarse-graining already requires consistency across representations at different scales. A demonstrated physical scale–time compatibility could eventually become a diagnostic for whether a reduced model preserves structural evolution across resolution, not only spectra and stresses. This is a future implication, not something demonstrated by the present study.
5. **Universality.** The critical universality question may not be whether $D_\square$ itself has one universal value — Pilot B already argues against that — but whether the ordering $D_\square^{\mathrm{real}} \ll D_\square^{\mathrm{N2}} < D_\square^{\mathrm{N1}} \ll D_\square^{\mathrm{N0}}$ and the bulk/branching-defect organization persist across flows. That is a substantially more sophisticated target for high-Reynolds-number and cross-regime tests than a fixed constant would be.

7.5 The zero-atom as a clue, not only a nuisance

The full $D_\square$ population has a pronounced two-component, zero-inflated structure: a large atom sits at exactly (or numerically indistinguishably from) zero, while a comparatively small population occupies the remaining $(0, 1]$ range. We avoid calling this "bimodal," a term that properly describes two separated modes of a continuous density; what is observed is a discrete atom at zero plus a separate continuous tail, which is why we describe it as zero-inflated and two-component instead. Recomputed directly from the frozen per-object ladder (`D_STITCH.csv`, both pilots), this atom accounts for 91.0% of Pilot B's 3,962 eligible objects and 85.6% of Pilot A's 4,957; the exact figure for Pilot B is precisely what breaks the R7 quantile selector (Section 6). Figure 7 shows the full split for both pilots: the zero-atom fraction, and the shape of the nonzero tail alongside it.

Read scientifically rather than only as a measurement problem, this exposes something about the structure itself. The population is not usefully described as "mostly commuting with some noise"; it is better described as *two components*: a dominant, essentially exactly commuting bulk, and a sparse, defect-bearing exceptional set whose typical $D_\square$ is not small. Under this routing representation, exact or effectively exact route compatibility is not rare — it dominates the population in both pilots — while noncommutation is concentrated in a comparatively small exceptional set that is itself preferentially associated with branching (Section 6, R8). This two-component picture is more concrete than "largely commuting": it identifies a specific population split to be explained, rather than a single continuous distribution to be summarized by its mean.

A candidate link to intermittency (hypothesis only). Turbulent intermittency is conventionally diagnosed through the anomalous scaling of velocity structure functions and the spatial rarity of intense dissipation events [5, 4]. The branching-localized noncommutation reported here is suggestively similar in shape — a sparse exceptional population embedded in an otherwise regular bulk — and it is tempting to read branching-localized $D_\square$ defects as a structural-routing marker for the same intermittent reorganization events that conventional diagnostics detect through dissipation statistics. We flag this only as a hypothesis. No equivalence with conventional intermittency diagnostics has been tested in this study, and doing so would require computing both on the same object population, which is future work.

7.6 A translation to conventional turbulence language

This study is expressed in the UNNS routing/admissibility vocabulary because that is the vocabulary the frozen instruments compute in. Table 11 gives a plain-language correspondence to more familiar turbulence quantities, to make the claims legible to a reader who has not used these instruments before (Table 11).

Empirical result vs. physical hypothesis. ***Empirical result:*** *within the frozen structural-routing representation described in Section 2, scale and time transformations exhibit reproducibly low route disagreement relative to the N0/N1/N2 controls, in two operationally independent samples.* ***Physical hypothesis:*** *this commutation structure reflects a genuine constraint on multiscale turbulent organization, rather than merely a property of the particular routing representation and object definition used to measure it. This manuscript establishes the first statement; it does not establish the second. Section 8 states what would be required to test the second.*

7.7 What we have not yet shown

At the outset of this research branch, the working hypothesis was that turbulence might be organized as a dynamically sustained routing among nested families of admissible coherent structures, with the expectation that distinct morphological families would eventually be identified in a strong sense (for example, a fixed number of universal turbulent structural species). This study does not establish that stronger, more concrete claim. What it establishes is one level more abstract: turbulence, in the tested corpus, appears to possess a reproducible transformation grammar in which structural evolution across scale and structural evolution across time usually commute, and departures from that compatibility are sparse and preferentially associated with branching or reorganization events. That is a more specific and more informative statement than "persistent structures exist," even though it stops short of a morphological classification of those structures.

8 Limitations

1. **Two samples, one Reynolds number, one turbulence class.** Both pilots are 256^3 cutouts of the same `isotropic1024coarse` forced-isotropic DNS. No claim is made of Reynolds-number generality or of generality beyond forced isotropic turbulence.
2. **Graph-local, not field-level, null controls.** N1 and N2 constrain distance and feature-similarity strata of the routing graph. The frozen STRUC-ROUTE-I export does not retain original voxel masks for arbitrary rewired pairs, so N1/N2 are graph-local geometry controls, not exact field-level voxel-overlap- or IoU-preserving surrogates.
3. **R3 is a genuine, unexplained divergence.** This manuscript does not explain *why* `D_STITCH`'s scalar admissibility differs between pilots; it reports that it does, at a magnitude well outside measurement noise.
4. **R7's failure is selector-specific, not a proof of absent localization.** Section 6 shows that alternative, non-preregistered tail definitions recover a positive but weaker signal. This is reported descriptively; it is explicitly not used to convert the frozen R7 result to a pass.
5. **R8's replicated statement is narrower than Pilot A's.** Only branching-type restructuring replicates; merging and aggregate-restructuring associations seen in Pilot A do not reproduce in Pilot B.
6. **Categorical thresholds are not physical constants.** The 0.93/0.60 STRUC-I regime boundaries and the R1–R8 significance threshold ($\alpha = 0.05$) are analysis conventions fixed before observation, not derived physical constants.
7. **No claim of full cross-chamber replication.** The preregistered classification is **Structural Partial Replication**, not **Core Replication**. This manuscript reports that result at full strength and does not soften it because other criteria are strongly positive.
8. **Two-pilot corpus, not an exhaustive survey.** All claims are scoped to the tested two-sample corpus. A third sample (Pilot C, Section 9) is proposed but not yet run; the present claims do not depend on its outcome.

8.1 Representation dependence and falsifiability

Every claim in this manuscript is established *within* the frozen structural-routing representation: a specific object definition, a specific ladder extraction, and null families (N0, N1, N2) that are graph-local rather than field-level surrogates. This is a scope statement, not an apology, and it converts directly into a falsifiability program rather than a standing disclaimer. The physical hypothesis in Section 7.6 — that scale–time commutation reflects a genuine constraint on turbulent organization rather than an artifact of this representation — predicts that the same qualitative signature (a dominant near-zero $D_\square$ population, sparse branching-localized defects, survival against N1/N2) should reappear under: (i) alternative object definitions not tied to the specific segmentation used here; (ii) field-level, rather than graph-local, surrogate nulls; (iii) different Reynolds numbers within the same JHTDB forced-isotropic family; and (iv) other turbulence classes (e.g. channel flow, decaying turbulence). Each of these is a concrete, falsifying test the representation could fail. None has yet been run; Pilot C (Section 9) addresses only the narrower R3/R7 sample-dependence question, not representation dependence in general.

9 Future Test: A Preregistered Pilot C

The R3/R7 divergence is scientifically warranted grounds for a third independent sample, but it is not a precondition for the claims of this manuscript, which are frozen and complete for the two-pilot record as reported. Pilot C is proposed as the next preregistered confirmatory study, specifically aimed at adjudicating whether Pilot B's Structural Boundary-level `D_STITCH` admissibility (R3) and degenerate-selector tail behavior (R7) are sample-specific or recur.

Two design decisions should be fixed in the Pilot C preregistration itself, before acquisition, precisely because they are informed by what Pilot B revealed:

1. **A selector specified blind to Pilot C's data.** R7's frozen 90th-percentile selector is degenerate whenever $D_\square$ has a large zero-atom, as in Pilot B. A preregistered alternative (for example, a fixed absolute threshold rather than a sample-conditional quantile) should be committed to in the Pilot C protocol text itself, not retrofitted after Pilot C's outcome is seen.
2. **A graded R3 criterion.** R3's Geometric-Persistence threshold is binary, but the underlying quantity is continuous. A secondary continuous read — distance-to-boundary, in the same spirit as the boundary-adjacency rule already used for the precision-rerun decision — would prevent a near-miss from being flattened into a bare pass/fail in future replications.

If Pilot C again shows a strong route mechanism with Pilot B-like softness in R3/R7, that pattern becomes part of the finding itself: a stable route mechanism with variable secondary expression. If Pilot C instead reproduces Pilot A's full signature, Pilot B is identified as the outlier and the mechanism claim can be stated without the R3/R7 qualification. Either outcome is informative under the stopping rule already in force (Section 2.9): Pilot B is not replaced under its own label, and any further sample is preregistered as Pilot C or later.

10 Data and Reproducibility

10.1 Archive

The raw source cutouts and their provenance/checksum companion records for both pilots are archived at Zenodo [15] as the source-data archive for this replication:

DOI: 10.5281/zenodo.22650769

This archive does not itself contain the full derived-result tree (per-instrument outputs, ladder extractions, and frozen chamber records); those are identified by content hash below and held in the project repository alongside the frozen `R1_R8`/`SYNTHESIS` records.

10.2 Source identity and hashes

	Pilot A	**Pilot B**
Source bytes	—	2,015,302,728 (1.8769 GiB)
Source SHA-256	`e32c9225...e02108a46`	`977e6ab3...ca0b969f`
Selection-lock SHA-256	—	`385cca0a...a40875e`
Acquisition	Hugging Face mirror	authenticated JHTDB `getCutout`
HDF5 validation	—	PASS

10.3 Frozen chamber records

Instrument	**Frozen version**
Physical adapter	`JHTDB_ROUTE_ADAPTER_v0_1_0`
STRUC-ROUTE-I	`STRUC_ROUTE_I_v0_1_2`
Ladder extraction	`ROUTE_LADDER_EXTRACT_v0_1_1` (native); `ROUTE_EXTRACT_PB_v013` (Pilot-B I/O wrapper)
STRUC-I	`STRUC_I_v1_0_4`
STRUC-PERC-I	`STRUC_PERC_I_v2_5_0`
STITCH-MECH	`STITCH_MECH_v0_1_1`

10.4 Canonical result records

The frozen record status is `FINAL_FROZEN_RESULT`, created 2026-09-07T22:11:59Z, comprising: `R1_R8.json` / `R1_R8.md` (the criterion record; Table 8 reproduces it in full); `SYNTHESIS.json` / `SYNTHESIS.md` (the Pilot A↔Pilot B comparison synthesis); the primary STRUC-ROUTE-I run (`jhtdb_pilot_b_20260907_145640`); the sensitivity run (`jhtdb_pilot_b_20260907_154708`); and the STITCH-MECH result bundle (`STITCH_MECH_RESULT.zip`, `STITCH_MECH_REPORT.json`). No frozen analysis setting was altered after Pilot B's data were inspected.

10.5 Attribution

JHTDB data are attributed to the Johns Hopkins Turbulence Database project [1]. Public dataset and cutout-service documentation: `https://turbulence.pha.jhu.edu/Forced_isotropic_turbulence.aspx` and `https://turbulence.pha.jhu.edu/newcutout.aspx`.

11 Conclusion

A preregistered, independent-sample replication of a scale–time commutation signature in forced isotropic turbulence is complete and frozen. Its classification is **Structural Partial Replication**: R1, R2, R4, R5, R6, and R8 pass; R3 and R7 fail.

The route mechanism — constrained routing, an exceptionally low real scale–time stitching defect, and survival of that low defect under both a distance-controlled (N1) and a distance-plus-feature-controlled (N2) graph-local null — reproduces in a spatially and temporally disjoint sample of the same forced-isotropic dataset, and does so with a *larger* real/null separation than in the original sample. It does not reproduce the complete cross-chamber signature: the scalar STRUC-I admissibility of `D_STITCH` falls short of Geometric Persistence in the replication, and the frozen upper-decile physical-intensity localization test degenerates against a $D_{\square}$ distribution with an unusually large zero-atom.

The two-pilot record separates a question that one pilot alone could not: route-mechanism reproducibility is strongly supported; full cross-chamber and localization reproducibility is not established. We read this as evidence for a reproducible transformation grammar — a largely commuting scale–time route regime whose sparse noncommuting exceptions localize preferentially around branching — rather than for a fixed, universal morphological classification of turbulent structures. Pilot B is now frozen as observed. A third preregistered sample (Pilot C) is proposed to adjudicate the R3/R7 divergence, but is future work rather than a precondition for the claims reported here.

Principle 11.1 (Final)**.** The route mechanism reproduces; the full cross-chamber signature does not. The final principle of this replication is:

> Within the tested routing representation, scale transformation and temporal evolution largely commute for turbulent structural objects, reproducibly across two operationally independent samples; the strength of that commutation and the localization of its rare exceptions are sample-dependent.

Tables

Table 1: Independent-sample audit (Section 2.2).

Property	Pilot A	Pilot B
Spatial cube (1-based)	$x, y, z \in [1{:}256]$	$x, y, z \in [513{:}768]$
Spatial cube (0-based)	$x, y, z \in [0{:}255]$	$x, y, z \in [512{:}767]$
Stored frames (1-based)	[1:10]	[501:510]
Stored Δt	0.002	0.002
Physical time window	$[0.000, 0.018]$	$[1.000, 1.018]$
Spatial grid-point overlap	0	
Time-frame overlap	0	
Physical-time gap (A end → B start)	0.982	

Table 2: Primary STRUC-ROUTE-I evidence, Pilot B (run `jhtdb_pilot_b_20260907_145640`). $n = 100$ nulls throughout.

Test	Real	Null mean	z_{fav}	p_{fav}	Read
scale persistence	0.346636	0.346954	-0.590	0.7426	not significant
time persistence	0.500681	0.488256	12.520	0.0099	supported
$D_{\square}$ stitch defect	0.014172	0.996336	868.885	0.0099	strongly supported

Table 3: Stitch defect by scale index (`STITCH_BY_SCALE.csv`, Pilot B primary run).

Scale index	Eligible	mean $D_{\square}$	mean enstrophy	mean diss. proxy
0 ($\times1 \to \times2$)	2,609	0.014390	1585.76	336.50
1 ($\times2 \to \times4$)	1,234	0.015058	950.92	170.91
2 ($\times4 \to \times8$)	114	0.000195	488.31	86.84
3 ($\times8 \to \times16$)	5	0.0	208.01	58.87

Table 4: Stitch defect by time index (`STITCH_BY_TIME.csv`, Pilot B primary run; 9 consecutive frame transitions across the 10-frame stored window).

Transition	Eligible	mean $D_{\square}$
t01→t02	455	0.011261
t02→t03	442	0.008479
t03→t04	438	0.013874
t04→t05	438	0.014857
t05→t06	430	0.010284
t06→t07	444	0.017507
t07→t08	439	0.020633
t08→t09	435	0.013132
t09→t10	441	0.017521

Table 5: Scale-sensitivity comparison, factors $\{1, 2, 4, 8\}$ vs. primary $\{1, 2, 4, 8, 16\}$ (Pilot B; run `jhtdb_pilot_b_20260907_154708`).

Metric	**Primary**	**Sensitivity**
verdict	CONSTRAINED_ROUTING	CONSTRAINED_ROUTING
objects	10,515	10,487
relations supplied / eligible	14,169 / 14,118	14,139 / 14,088
scale persistence (p_{fav})	0.7426	0.7426
time persistence (p_{fav})	0.0099	0.0099
$D_{\square}$ (p_{fav})	0.0099	0.0099
stitch-eligible sources	3,962	3,957
null mean mobility	0.998859	0.998869

Table 6: STRUC-I results, Pilot B (2,000-MC primary; 10,000-MC precision rerun for `P_TIME`, triggered by boundary-adjacency).

Ladder	**Regime / State**	$\langle A_\kappa \rangle$	min A_κ
`D_STITCH` (2,000 MC)	Structural Boundary/Transitional Structure	0.851863	0.7925
`P_SCALE` (2,000 MC)	Structural Instability/Random Structure	0.564837	0.5395
`P_TIME` (2,000 MC)	Structural Instability/Random Structure	0.597188	0.5620
`P_TIME` (10,000 MC, precision)	Structural Instability/Random Structure	0.598235	0.5784

Table 7: STRUC-PERC-I results, Pilot B.

Ladder	**Verdict**	**Giant ratio**	**Isolated**	κ_{conn}	**Gap** n
`D_STITCH`	FULL_PERCOLATION	1.000000	0	10	336
`P_SCALE`	FULL_PERCOLATION	1.000000	0	0.010	3
`P_TIME`	FULL_PERCOLATION	1.000000	0	1	9

Table 8: Frozen R1–R8 criterion record, Pilot B. Record status `FINAL_FROZEN_RESULT`, created 2026-09-07T22:11:59Z.

ID	Frozen test	Pilot-B result	Status
R1	Route organization	constrained_routing; 100/100 nulls; mobility 0.998859; unique 1.0; match 0	**PASS**
R2	Primary N0 stitching contrast	real 0.014172 $\ll$ N0 0.996336; $z = 868.9$; $p = 0.0099$	**PASS**
R3	STRUC I `D_STITC` $\rightarrow$ Geome or stronger	$\langle A_\kappa \rangle = 0.851863$; Structural Boundary	**FAIL**
R4	STRUC PERC-I connectivity	full_percolation; giant 1.0; gaps 336; isolated 0	**PASS**
R5	N1 survival	mean 0.326204; real/N1 0.0434; $z = 98.9$; $p = 0.0099$	**PASS**
R6	N2 survival (principal target)	mean 0.211452; real/N2 0.0670; $z = 36.0$; $p = 0.0099$	**PASS**
R7	Physica intensit localization	enstrophy 1.0×, $p = 1.0$; dissipation 1.0×, $p = 1.0$	**FAIL**
R8	Restruc localization	scale branching $p = 0.00020$; time branching $p = 0.00020$	**PASS**

Classification: ***Structural Partial Replication.*** *Additional preregistered robustness (not itself an R-criterion): [1,2,4,8] scale sensitivity also returns constrained_routing.*

Table 9: Direct Pilot A ↔ Pilot B criterion comparison. The Pilot-A column maps the frozen baseline onto the later R1–R8 definitions for direct comparison; this is not a claim that Pilot A was itself prospectively preregistered under these labels.

ID	Pilot A	Pilot B	Comparison
R1	PASS — CONSTRAINED_ROUTING	PASS — CONSTRAINED_ROUTING	route verdict replicates
R2	PASS — real 0.026628 vs. N0 0.996786, $p = 0.0099$	PASS — real 0.014172 vs. N0 0.996336, $p = 0.0099$	primary contrast replicates
R3	PASS — $A_\kappa = 0.971708$, GEOMETRIC PERSISTENCE	FAIL — $A_\kappa = 0.851863$, STRUCTURAL BOUNDARY	cross-chamber admissibility does not replicate
R4	PASS — giant 0.996933, gaps 652	PASS — giant 1.0, gaps 336	connectivity replicates
R5	PASS — N1 0.329091, $p = 0.0099$	PASS — N1 0.326204, $p = 0.0099$	N1 survival replicates
R6	PASS — N2 0.223649, $p = 0.0099$	PASS — N2 0.211452, $p = 0.0099$	N2 survival replicates (principal target)
R7	PASS — enstrophy 2.50×, diss. 1.87×	FAIL — both 1.0×, $p = 1.0$	tail localization does not replicate
R8	PASS — branching + time merging + any restructure	PASS — branching only	replicates in narrower form

Table 10: Quantitative stitching-defect hierarchy, Pilot A vs. Pilot B.

Quantity	Pilot A	Pilot B	B relative to A
REAL $D_\square$	0.026628	0.014172	46.8% lower
N0 mean	0.996786	0.996336	0.045% lower
N1 mean	0.329091	0.326204	0.88% lower
N2 mean	0.223649	0.211452	5.45% lower
N1 / REAL	12.36×	23.02×	larger separation in B
N2 / REAL	8.40×	14.92×	larger separation in B

Table 11: Translation from UNNS instrument quantities to conventional turbulence language. The right column is descriptive, not a claim of formal equivalence.

UNNS quantity	**Conventional turbulence reading**
$D_\square$ (scale–time stitching defect)	Scale/time route incompatibility: how much a structural object's ancestry depends on the order of coarse-graining and temporal advance
Branching (scale or time)	Structural splitting / topological reorganization of a tracked object
Mean enstrophy (physical-intensity proxy)	Local vorticity intensity
Dissipation proxy (2ν[enstrophy $-$ $2Q] = 2\nu\|S\|^2$) [9]	Algebraically the conventional strain-based rate $\epsilon = 2\nu S_{ij}S_{ij}$, reconstructed from the adapter's coarse-grained finite-difference fields rather than measured on a native-resolution DNS dissipation field; a proxy in that reconstructive sense, not a different formula
P_{TIME}	Persistence of a structural object's identity across temporal evolution
P_{SCALE}	Persistence of a structural object's identity across coarse-graining
A_κ (perturbative admissibility)	Robustness of the low-$D_\square$ population to small controlled perturbations of the routing graph
FULL_PERCOLATION / HARD_FRAGMENTATION (percolation state)	Whether low-$D_\square$ ladder values form one connected value-space population or fragment into isolated clusters
N0/N1/N2 (null families)	Degree-preserving, then distance-stratified, then distance-and-feature-stratified randomizations of the same routing graph

Table 12: Relation to established turbulence frameworks (Section 7.4.5). The relation is one of complementarity: no framework listed here is contradicted by the present result, and none of them is subsumed by it.

Existing framework	**What it asks**	**Relation to our result**
Kolmogorov / inertial-range scaling [2, 3]	How velocity statistics scale with separation	Our result concerns the compatibility of structural transformations, not scaling exponents
Intermittency / multifractals / She–Leveque [4, 5]	How rare intense events produce anomalous scaling	Sparse noncommuting branching events are a candidate structural correlate, but this has not yet been demonstrated
Coarse-graining / LES / Germano [7]	How filtered representations at different resolutions are related	Closest operatorial analogue; our test compares structural routes across resolution and time rather than filtered stresses

Existing framework	What it asks	Relation to our result
Scale-local cascade theories	How energy transfer couples neighboring scales	Potentially compatible with route locality, but $D_\square$ is not an energy-flux locality measure
Markov cascade / stochastic cascade models [11]	Transition probabilities of velocity increments through scale	Our route grammar may provide a structural counterpart involving both scale and time
Coherent-structure theories	Which organized structures inhabit turbulent fields	We do not identify a universal shape; we identify a reproducible relation among transformations
Exact coherent structures / dynamical systems [12]	Whether recurrent/invariant Navier–Stokes states organize turbulent phase space	Potentially complementary: ECS give a state-space skeleton; our result suggests a transformation-space skeleton
Lagrangian coherent structures [13]	Which robust transport skeletons organize particle trajectories	Distinct target (transport, not routing); both seek organization beneath apparent randomness
Renormalization / coarse-graining ideas	What survives repeated change of scale	Suggestive analogy only; we have not established a fixed point, universality class, or RG flow

Table 13: R7 frozen 90th-percentile selector, Pilot B (degenerate: $p_{90}(D_\square) = 0$).

Measure	Enrichment	Permutation p	topD n / eligible n
enstrophy	1.0×	1.0	3,962 / 3,962
dissipation proxy	1.0×	1.0	3,962 / 3,962

Table 14: R7 alternative tail definitions, Pilot B (descriptive only; not substituted for the preregistered test).

Tail def.	n	frac.	$D_\square$ median	enstrophy ratio	diss. ratio
top10pct (frozen)	3,962	100%	0.0	—	—
top5pct	199	5.02%	0.2286	1.187×	1.290×
top1pct	41	1.03%	0.5887	1.129×	1.272×
$D_\square \geq 0.1$	144	3.63%	0.3341	1.173×	1.275×
$D_\square \geq 0.5$	33	0.83%	0.6516	1.151×	1.321×

Table 15: R8 branch/merge contrasts, Pilot B: $D_{\square}$ median difference, flagged vs. unflagged.

Flag	**Flagged** n	**Unflagged** n	**Median** Δ	p	**Replicated?**
scale branching	237	3,725	+0.0002181	0.00020	**YES**
time branching	81	3,881	+0.0000112	0.00020	**YES**
scale merging	178	3,784	0	1.0	no
time merging	60	3,902	0	1.0	no
any restructuring	428	3,534	0	1.0	no

Figures

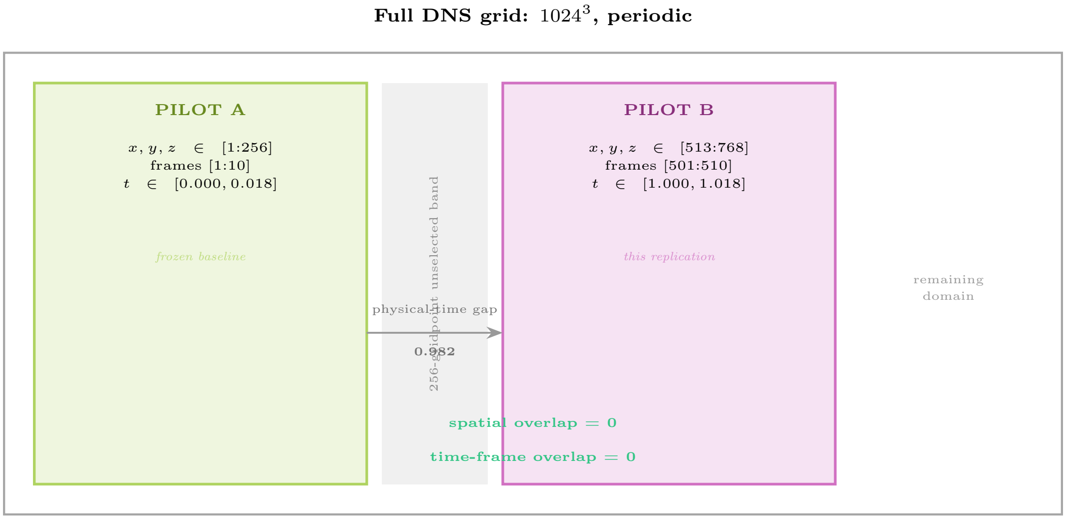

Figure 1: Independent-sample geometry. Pilot A (frozen baseline) and Pilot B (this replication) are disjoint 256^3 cutouts of the same `isotropic1024coarse` JHTDB simulation, related by a half-domain translation in x, y, and z: zero shared spatial grid points, zero shared stored frames, and a 0.982 physical-time gap between the end of Pilot A's window and the start of Pilot B's.

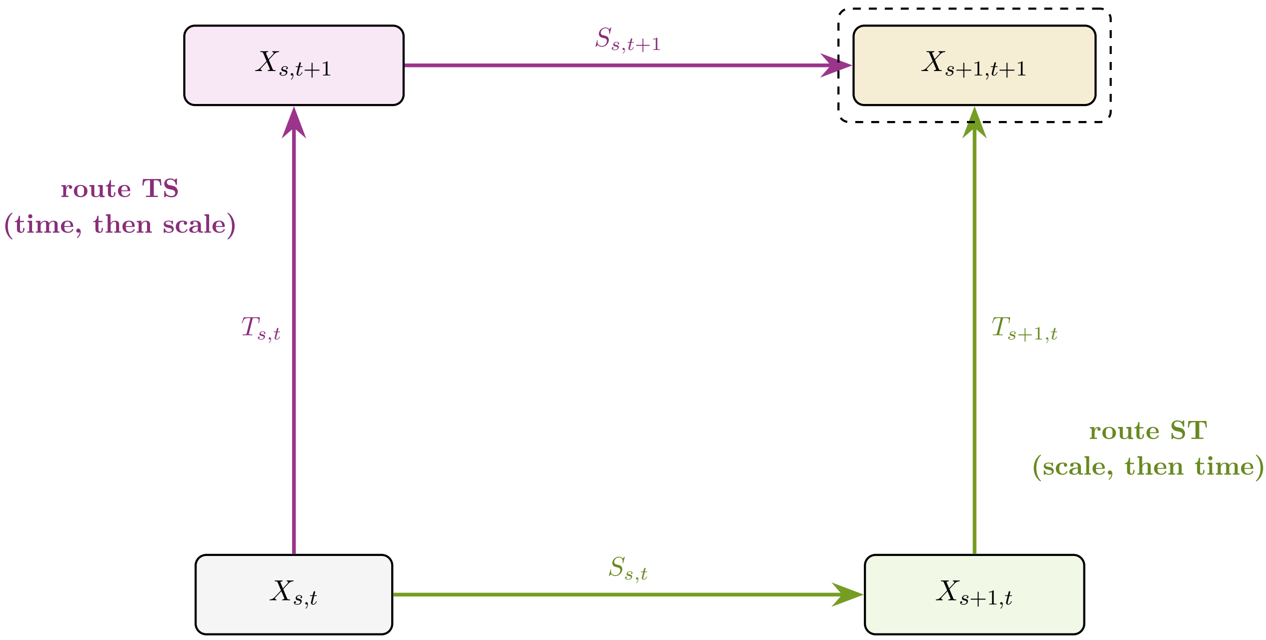

Figure 2: The scale–time commutation square. Each arrow is a confidence-weighted outgoing routing distribution over candidate destination objects (Section 2.5), not a deterministic single-object trajectory. Route ST advances scale first, then time; route TS advances time first, then scale. $D_\square$ (Eq. 1) is the normalized Jensen–Shannon divergence between the two resulting two-step destination distributions P_{ST} and P_{TS}, evaluated only for source objects with eligible outgoing relations on both routes. Low $D_\square$ means the two routes are nearly interchangeable in the distributions they induce; it is not a claim that any single object literally traverses both paths.

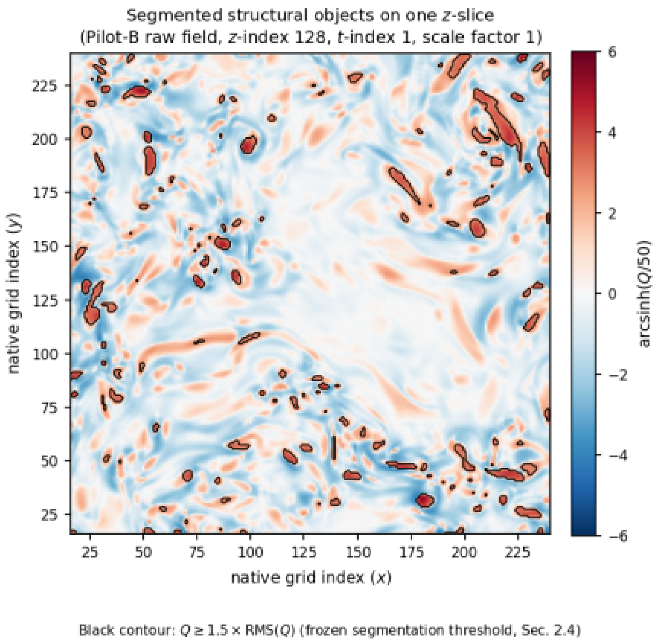

Figure 3: Slice intersections of actual segmented 3D structural objects. One native-resolution (256^3, block factor 1) z-slice through the real Pilot-B raw velocity field, z-index 128, t-index 1 (first stored frame), chosen by a fixed rule (mid-plane, first frame, finest scale) rather than selected for appearance. Objects are defined as 3D 26-connected components (Section 2.4); what is shown here is the 2D cross-section of those 3D objects induced by this single slice, not the objects in their full volumetric extent. Background: $\mathrm{arcsinh}(Q/50)$, the Q-criterion computed by the same boundary-safe finite-difference scheme used by the adapter. Black contours: the frozen segmentation boundary $Q \geq 1.5 \times \mathrm{RMS}(Q)$ on the boundary-safe interior — exactly the rule that defines a structural object throughout this manuscript. This translates the abstract routing-graph language of Sections 2.4–2.5 back into the velocity field it is built from.

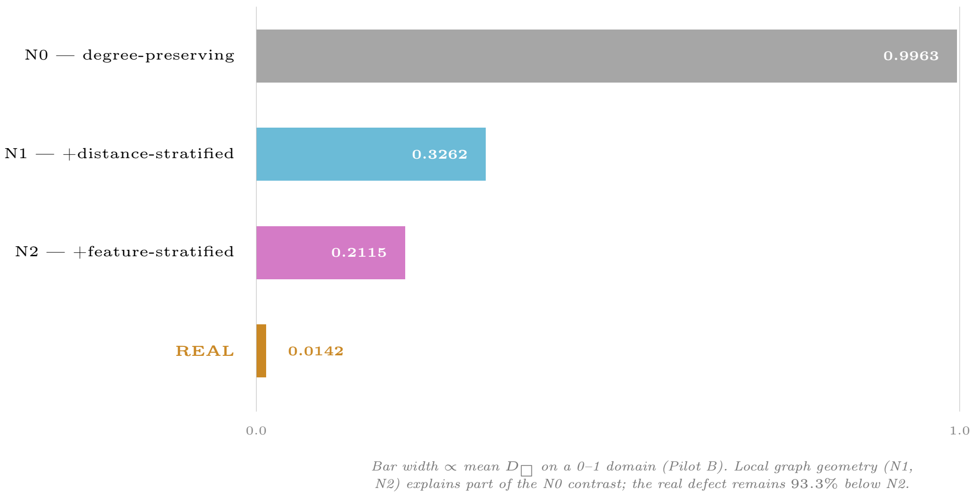

Figure 4: The N0→N1→N2→real collapse (Pilot B). Each nested null constrains additional local geometry (N1: distance; N2: distance + feature similarity), yet the real population mean $D_\square$ remains an order of magnitude below even the hardest tested null. The same qualitative collapse occurs in Pilot A ($0.9968 \to 0.3291 \to 0.2236 \to 0.0266$; Table 10).

ID	Pilot A	Pilot B	Comparison
R1	PASS	PASS	*route replicates*
R2	PASS	PASS	*contrast replicates*
R3	PASS	FAIL	*admissibility diverges*
R4	PASS	PASS	*connectivity replicates*
R5	PASS	PASS	*N1 survival replicates*
R6	PASS	PASS	*N2 survival replicates*
R7	PASS	FAIL	*localization diverges*
R8	PASS	PASS	*narrower replication*

6/8 pass in Pilot B · Classification: Structural Partial Replication

Figure 5: Direct R1–R8 criterion comparison, Pilot A vs. Pilot B. Six of eight criteria replicate; R3 (cross-chamber scalar admissibility) and R7 (physical-intensity tail localization) fail in Pilot B. R8 replicates in a narrower form (branching only; Table 15).

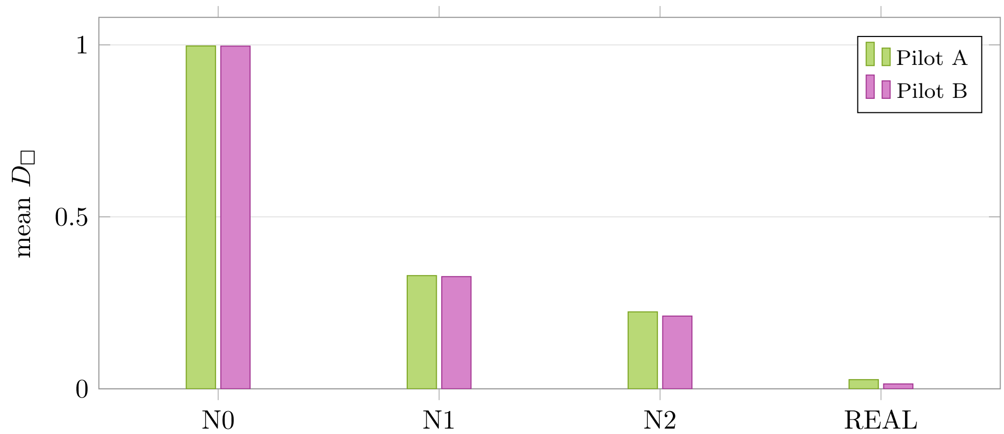

Figure 6: Two-pilot null hierarchy comparison. Mean $D_\square$ under each nested null (N0, N1, N2) and for the real population, Pilot A vs. Pilot B (Table 10). The null hierarchy is nearly unchanged between pilots while the real value drops substantially further in Pilot B: the real/null separation is the reproducing quantity, not a fixed real value.

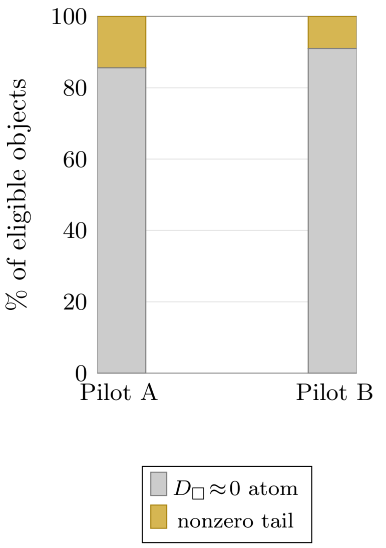

(a) Zero-atom vs. nonzero-tail split.

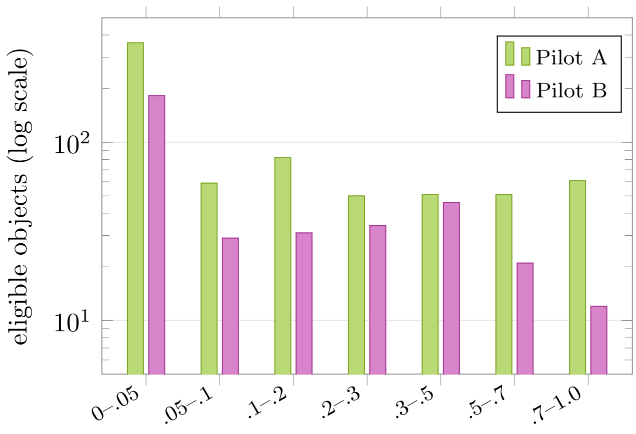

(b) Nonzero tail only (zero-atom excluded, shown in (a) instead).

Figure 7: $D_\square$ population structure, Pilot A vs. Pilot B. (a) The full population splits into a dominant zero (or numerically indistinguishable from zero) atom and a comparatively small nonzero population, computed directly from the frozen per-object ladder (`D_STITCH.csv`): Pilot A, 85.6% of 4,957 eligible objects at $D_\square \approx 0$ (716 nonzero); Pilot B, 91.0% of 3,962 eligible objects at $D_\square \approx 0$ (356 nonzero). **(b)** The nonzero population alone (panel (a)'s excluded atom), by bin: both pilots show the same qualitative shape, a long, thin tail extending to $D_\square \approx 1$. Together the two panels make explicit what Section 7.5 calls the two-component, zero-inflated structure of the full population.

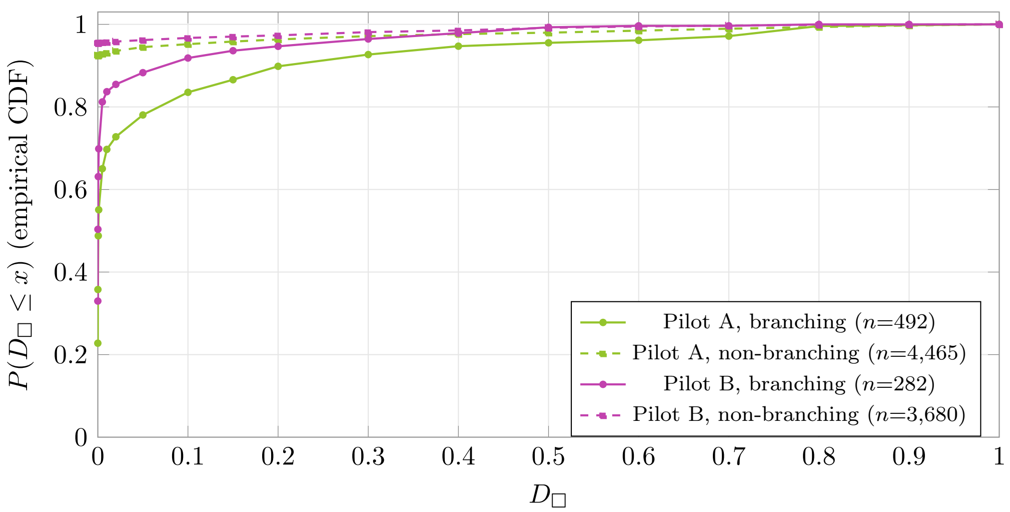

Figure 8: Branching localizes elevated $D_\square$: marginal empirical CDFs, both pilots. Computed directly from `STITCH_MECH_OBJECTS.csv` (frozen per-object records), splitting each pilot's stitch-eligible population by whether the object carries a scale-branch or time-branch flag. Across most of the $D_\square$ range the non-branching curve (dashed) lies above the branching curve (solid) in both pilots, with the curves converging in the far tail as both approach 1: the non-branching population's mass is concentrated at or near $D_\square = 0$, while the branching population's $D_\square$ is systematically shifted toward larger values across the distribution, not only at the median. This is the same R8 contrast reported in Table 15 (Section 6.3), shown here as the full marginal distribution rather than a single median statistic; these are unconditioned (raw) distributions, not the geometry-partial-rank association reported in Appendix B.

A R1–R8 Preregistered Criteria (Full Text)

Reproduced from the frozen preregistration for reference; see Section 2.9 for discussion.

R1 STRUC-ROUTE-I returns CONSTRAINED_ROUTING with a usable primary null ensemble.

R2 Real mean $D_\square$ is lower than N0 with $p_{\text{fav}} < 0.05$.

R3 `D_STITCH` is GEOMETRIC PERSISTENCE or stronger under canonical STRUC-I settings; if $\langle A_\kappa \rangle$ lies within 0.01 of a regime boundary, rerun at 10,000 MC with all other settings unchanged.

R4 `D_STITCH` has giant-component ratio ≥ 0.95 and ≥ 100 gap vertices after deduplication.

R5 N1 is usable and $D_\square^{\text{real}} < D_\square^{\text{N1}}$ with $p_{\text{fav}} < 0.05$.

R6 N2 is usable and $D_\square^{\text{real}} < D_\square^{\text{N2}}$ with $p_{\text{fav}} < 0.05$ (principal mechanism replication target).

R7 Top-decile $D_\square$ tail is enriched (enrichment > 1, stratified permutation $p < 0.05$) in top-decile enstrophy and top-decile dissipation proxy.

R8 At least one preregistered restructuring class (scale branching, time branching, scale merging, time merging, any restructuring) has higher $D_\square$ with $p < 0.05$.

Interpretation labels. CORE REPLICATION: R1–R6 all pass. STRUCTURAL PARTIAL REPLICATION: R2 passes and at least one of R5/R6 passes, but one or more cross-chamber criteria fail. ROUTE-ONLY REPLICATION: R1, R2 pass but both N1/N2 mechanism-survival criteria fail or are underresolved. NO REPLICATION: R2 fails under a usable N0 ensemble. `Underresolved` is reserved for unusable null ensembles.

B Object-Level Mechanism Detail, Pilot B

B.1 Correlations

Ordinary (unconditioned) rank association between $D_\square$ and physical intensity, over the 3,962 stitch-eligible objects: $\rho(D_\square, \text{enstrophy}) = 0.10207$; $\rho(D_\square, \text{dissipation proxy}) = 0.12663$. Geometry-conditioned partial-rank association (conditioned on scale, time, outgoing overlap geometry, IoU, normalized distance, feature similarity, and branch/merge state): $\rho_{\text{partial}}(D_\square, \text{enstrophy}) = 0.10273$; $\rho_{\text{partial}}(D_\square, \text{dissipation proxy}) = 0.04936$. For comparison, Pilot A's corresponding partial-rank values are 0.17902 and 0.12005.

B.2 STITCH-MECH null quality gates

Null	Mean $D_\square$	real/null	mean mobility	unique / match
N1 (distance-stratified)	0.326204	0.04344	0.044413	100/100, 0
N2 (+feature-stratified)	0.211452	0.06702	0.014257	100/100, 0

All quality-gate thresholds (min mean-mobility 0.01; min unique-graph fraction 0.1; max real-match fraction 0.9) are passed with no flags in both N1 and N2.

References

[1] Perlman, E., Burns, R., Li, Y., Meneveau, C. (2007). *Data Exploration of Turbulence Simulations Using a Database Cluster.* Proceedings of the ACM/IEEE Conference on Supercomputing (SC07). Dataset and cutout-service documentation: `https://turbulence.pha.jhu.edu`.

[2] Kolmogorov, A. N. (1941). *The Local Structure of Turbulence in Incompressible Viscous*

Fluid for Very Large Reynolds Numbers. Doklady Akademii Nauk SSSR, 30, 301–305.

[3] Frisch, U. (1995). *Turbulence: The Legacy of A. N. Kolmogorov.* Cambridge University Press.

[4] She, Z.-S., Leveque, E. (1994). *Universal Scaling Laws in Fully Developed Turbulence.* Physical Review Letters, 72(3), 336–339.

[5] Meneveau, C., Sreenivasan, K. R. (1991). *The Multifractal Nature of Turbulent Energy Dissipation.* Journal of Fluid Mechanics, 224, 429–484.

[6] Farge, M. (1992). *Wavelet Transforms and Their Applications to Turbulence.* Annual Review of Fluid Mechanics, 24, 395–457.

[7] Germano, M. (1992). *Turbulence: The Filtering Approach.* Journal of Fluid Mechanics, 238, 325–336. DOI: `https://doi.org/10.1017/S0022112092001733`.

[8] Ghosal, S., Moin, P. (1995). *The Basic Equations for the Large Eddy Simulation of Turbulent Flows in Complex Geometry.* Journal of Computational Physics, 118(1), 24–37.

[9] Hunt, J. C. R., Wray, A. A., Moin, P. (1988). *Eddies, Streams, and Convergence Zones in Turbulent Flows.* Center for Turbulence Research, Proceedings of the Summer Program, CTR-S88.

[10] Lin, J. (1991). *Divergence Measures Based on the Shannon Entropy.* IEEE Transactions on Information Theory, 37(1), 145–151.

[11] Friedrich, R., Peinke, J. (1997). *Description of a Turbulent Cascade by a Fokker–Planck Equation.* Physical Review Letters, 78(5), 863–866.

[12] Waleffe, F. (1998). *Three-Dimensional Coherent States in Plane Shear Flows.* Physical Review Letters, 81(19), 4140–4143.

[13] Haller, G. (2015). *Lagrangian Coherent Structures.* Annual Review of Fluid Mechanics, 47, 137–162.

[14] UNNS Substrate Research Program. *STRUC-ROUTE-I, STRUC-I, STRUC-PERC-I, and STITCH-MECH: Frozen Instrument Specifications.* Internal technical reports, versions v0.1.2, v1.0.4, v2.5.0, and v0.1.1 respectively, 2026.

[15] UNNS Collective. *UNNS Turbulence — JHTDB Analysis Cutouts and High-Re Export v0.1.* Zenodo, 2026. DOI: 10.5281/zenodo.22650769.