

Supplementary Methods

Non-gradient spatial organisation at tumour–immune interfaces revealed by operator-based analysis

Anas Enoch, MD

Mohammed VI University of Health Sciences (UM6SS), Casablanca, Morocco

`anas_nour@um5.ac.ma`

Data, code, and ethics declarations are provided in the main manuscript.

Contents: Section S1 (mathematical framework and wedge algebra), S2 (GNN architecture and transport model), S3 (Lie-structured null model), S4 (NCG commutator grounding), S5 (Zeta spectral diagnostic), S6 (stochastic Hodge Bayes factors), S7 (KTS transition bias, GSE278936), S8 (baseline comparisons), S9 (remeshing and autocorrelation diagnostics); Supplementary Tables S1–S6; Supplementary Figures S1–S4.

Contents

1	S1 Mathematical Framework: Wedge Algebra and Hodge Decomposition	3
1.1	Oriented cell complex and differential forms	3
1.2	Incidence matrices and exterior derivatives	3
1.3	Wedge algebra and antisymmetry	3
2	S2 Governing Transport Model and GNN Architecture	4
3	S3 Lie-Structured Spatially Aware Null Model	4
4	S4 NCG Commutator Grounding	4
5	S5 Zeta Function Spectral Diagnostic	5
6	S6 Stochastic Hodge Decomposition and Bayesian Model Comparison	5
7	S7 Extended Prostate Cancer Validation: KTS Transition Bias and Graph-KS Instability	6
8	S8 Baseline Comparisons and Confound Exclusions	7
9	S9 Remeshing Controls and Autocorrelation Diagnostics	9
9.1	Remeshing sensitivity	9

9.2	Spatial autocorrelation of the coexact field	9
9.3	Bivector orientation analysis	9

1 S1 Mathematical Framework: Wedge Algebra and Hodge Decomposition

1.1 Oriented cell complex and differential forms

The tissue section is represented as an oriented cell complex $\mathcal{K} = (V, E, F)$: V are vertices (Visium spots or CosMx cells), E are oriented edges (kNN or Delaunay pairs), F are triangular faces. Discrete differential forms represent spatial quantities: 0-forms are scalar node functions (programme scores); 1-forms are antisymmetric edge functions (fluxes); 2-forms are face functions (circulations).

1.2 Incidence matrices and exterior derivatives

The signed node–edge incidence matrix $B_1 \in \{-1, 0, +1\}^{|E| \times |V|}$ represents $d_0 : C^0 \rightarrow C^1$: $(B_1)_{e,v} = +1$ if v is the head of e ; -1 if the tail; 0 otherwise. The edge–face matrix B_2 represents $d_1 : C^1 \rightarrow C^2$ with face-consistent orientation.

The k -Hodge Laplacian: $\Delta_k = d_{k-1}d_{k-1}^\top + d_k^\top d_k$. Explicitly: $\Delta_0 = B_1^\top B_1$; $\Delta_1 = B_1 B_1^\top + B_2^\top B_2$; $\Delta_2 = B_2 B_2^\top$. The Hodge decomposition theorem: $C^1 = \text{im}(B_1^\top) \oplus \text{im}(B_2) \oplus \text{ker}(\Delta_1)$ (exact \oplus coexact \oplus harmonic; orthogonal; unique).

1.3 Wedge algebra and antisymmetry

Construction. Given residualized programme fields $\tilde{a}, \tilde{b} \in C^0$, the wedge flux on oriented edge ($i \rightarrow j$) is:

$$\omega_{ij} = \frac{\tilde{a}_i \tilde{b}_j - \tilde{a}_j \tilde{b}_i}{\ell_{ij}}. \quad (1)$$

By construction, $\omega_{ij} = -\omega_{ji}$ (antisymmetric under edge reversal), so $\omega \in C^1$ with consistent orientation.

Uniqueness. Suppose $\phi : C^0 \times C^0 \rightarrow C^1$ is a local bilinear operator that is antisymmetric under programme exchange and vanishes when $\tilde{a} \propto \tilde{b}$ on a given edge. Then ϕ must be proportional to the wedge on each edge. Any symmetric bilinear operator yields an even-valued function on undirected edges mapping to a 0-form under B_1^\top ; its coexact projection is therefore zero. The wedge is the unique (up to scale) local construction producing a non-zero coexact projection.

Non-commutative geometry identity. For diagonal operators M_A, M_B on C^0 and any symmetric graph operator L :

$$[M_A, LM_B]_{ij} = L_{ij}(\tilde{a}_i \tilde{b}_j - \tilde{a}_j \tilde{b}_i). \quad (2)$$

The wedge flux is the commutator of programme operators weighted by graph connectivity [2]. The coexact component is the geometric image of operator non-commutativity; interface-specific localization follows because a construction artifact would distribute coexact content in proportion to field energy everywhere, whereas operator non-commutativity is maximal where programmes spatially oppose.

Decomposition algorithm. (1) Solve $\Delta_0 \alpha = B_1 f$ (conjugate gradient, tol 10^{-10}). (2) $f_{\text{exact}} = B_1^\top \alpha$. (3) Solve $\Delta_2 \beta = B_2 f$. (4) $f_{\text{coexact}} = B_2^\top \beta$. (5) $\gamma = f - f_{\text{exact}} - f_{\text{coexact}}$. Energy fractions: $\rho_k = \|f_k\|^2 / \|f\|^2$.

2 S2 Governing Transport Model and GNN Architecture

Reference transport model. Discrete advection–diffusion equation on 1-forms:

$$\frac{\partial \omega}{\partial t} = -\mathcal{L}_v \omega + D \Delta_1 \omega + R(\omega),$$

where $\mathcal{L}_v = \iota_v d + d \iota_v$ (discrete Cartan formula). This model is used as a structured null — the most permissive conservative dynamics class — not as a claim about tissue biology.

Learning objective.

$$\mathcal{L} = \mathcal{L}_{\text{data}}(\omega_\theta) + \lambda_{\text{PDE}} \mathcal{L}_{\text{ADR}}(\omega_\theta) + \lambda_{\text{Stokes}} \mathcal{L}_{\text{integral}}(\omega_\theta),$$

where $\mathcal{L}_{\text{integral}} = \sum_t \sum_s \left| \oint_{\partial \Omega_s} \omega_\theta(t) - \int_{\Omega_s} d\omega_\theta(t) \right|^2$ enforces Green–Stokes flux conservation over held-out subcomplexes. Conservation constraints enforce structural invariants, preventing the GNN from representing non-physical modes that fit observed data but violate flux balance on unseen subcomplexes.

Implementation: PyTorch with sparse DEC-consistent incidence matrices constructed from Delaunay triangulations with orientation-consistent indexing.

3 S3 Lie-Structured Spatially Aware Null Model

Edge weights.

$$w_e = [(\Delta_1 + \varepsilon I)^{-s}]_{ee}, \quad s(\Delta) = \log_2(1 + \Delta/\lambda_{\min}^+),$$

where ε excludes harmonic zero modes and Δ is the spectral gap of Δ_1 . Normalized: $\tilde{w}_e = w_e / \max_{e'} w_{e'} \in (0, 1]$.

Surrogate generation. For each surrogate k : $\theta_e^{(k)} \sim \mathcal{U}[-\pi \tilde{w}_e, \pi \tilde{w}_e]$; flux rotated by $R_e^{(k)} = \exp(-\frac{\theta_e^{(k)}}{2} B_e)$. Weights depend only on Δ_1 , not on the observed statistic; the null is operator-intrinsic without being diagnostically circular.

4 S4 NCG Commutator Grounding

Testable prediction. If the coexact component is the image of $[M_A, M_B]$, the per-node non-commutativity norm

$$\text{NC}_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} f_{ij}^2$$

should correlate with coexact energy $E_{\text{coexact},i}$ *specifically* at interface nodes ($\Delta\rho > 0$); a construction artifact would produce uniform correlation everywhere ($\Delta\rho \approx 0$).

Results across 18 valid TNBC sections. Spearman $\rho(\text{NC}, E_{\text{coexact}})$ within interface nodes: median 0.832, 18/18 ($p < 10^{-5}$). Interface-minus-other $\Delta\rho$: median +0.198, 18/18 ($p < 10^{-5}$). Individual section correlations: 0.530–0.923, all FDR < 0.001 .

Supplementary Table S1: **NCG commutator grounding: cohort-level results across 18 valid TNBC sections.** The per-node non-commutativity norm NC_i and node-level coexact energy $E_{\text{coexact},i}$ are highly correlated within interface-like nodes in every section, and this correlation is specifically elevated at interfaces relative to other regions. Both results are inconsistent with a uniform construction artifact and consistent with the NCG prediction that the coexact component is the geometric image of operator non-commutativity.

Metric	Median	Sections positive	Sign test p	Interpretation
$\rho(NC, E_{\text{coexact}})$ within interface nodes	0.832	18/18	$< 10^{-5}$	positive, consistent
$\Delta\rho$ (interface – other regions)	+0.198	18/18	$< 10^{-5}$	interface-specific elevation

5 S5 Zeta Function Spectral Diagnostic

Normalized Zeta functional.

$$Z(s) = \frac{\sum_k \alpha_k \lambda_k^{-s}}{\sum_k \alpha_k}, \quad \alpha_k = \langle E_{\text{coexact}}, \phi_k \rangle^2.$$

Removing energy dependence from the unnormalized form $Z_{\text{raw}} = Z \cdot \sum_k \alpha_k$ is essential: the unnormalized statistic inflates at high-energy (interface) nodes by construction, so apparent spectral enrichment before normalization reflects energy elevation, not independent spectral structure. $Z(s) > 1$ indicates bias toward low- λ spatially smooth modes. Null: 1,000 spatial label permutations. Energy-matched null: interface masks restricted to nodes with comparable coexact energy quantiles.

Results. Global coexact signal is non-flat in 17–19/19 sections (Supplementary Table S2). After energy normalization, interfaces are *not* more spectrally organized than equally energetic non-interface regions (2/43 TNBC, 2/23 prostate significant; sign test $p = 1.0$). Coexact enrichment at interfaces reflects elevated interaction intensity; interface-specific spectral geometry is not an additional independent property.

Supplementary Table S2: **Zeta function spectral diagnostic: cohort-level results across 19 TNBC sections.** Per-node coexact energy was projected onto eigenvectors of the node graph Laplacian. All four metrics indicate that the global coexact signal is not spectrally flat (consistent low-frequency concentration above a label-permutation null). $Z(s) > 1$ means coexact energy is biased toward spatially coherent large-scale modes. Note: these results use the label-permutation null; the energy-normalized null, applied in the GSE278936 cohort (Section S7), is required to test interface-specific spectral enrichment independently of energy effects. Sign test: one-sided binomial, H_0 : observed \leq null median.

Metric	Median	Sections $>$ null	Sign test p	Interpretation
$Z(s = 1)$	1.89	18/19	$< 10^{-4}$	low-freq concentrated
$Z(s = 2)$	2.13	18/19	$< 10^{-4}$	low-freq concentrated
Spectral Gini	0.663	17/19	$< 10^{-3}$	concentrated distribution
f_{low}	0.653	19/19	$< 10^{-5}$	low-freq concentrated

6 S6 Stochastic Hodge Decomposition and Bayesian Model Comparison

Model classes. M_0 (passive): zero cycle variance. M_{1a} (uniform active): $k = 30$ low-frequency cycle modes, uniformly distributed. M_{1b} (interface-localised): same modes, variance redistributed to interface-

adjacent edges (weight 3.0, trace-normalised for equal complexity with M_{1a}). Noise $\sigma = 0.5 \times \text{std}(P_c Y / \|Y\|)$ calibrated per section.

Log Bayes factors. $\log B_{1a/0}$: does the flux field require any non-gradient cycle structure? $\log B_{1b/1a}$: is that structure specifically interface-localised? Both computed analytically from Gaussian marginal likelihoods.

Results across 19 TNBC sections. $\log B_{1a/0}$: median +45.95 (19/19, $p < 10^{-5}$) — cycle structure is required. $\log B_{1b/1a}$: median +517.6 (19/19, $p < 10^{-5}$) — the cycle structure is specifically interface-localised. Posterior R under M_{1b} : 9.81 (19/19); under M_0 (passive null): 1.00 (11/19, $p = 0.32$).

Supplementary Table S3: **Stochastic Hodge decomposition: cohort-level Bayesian model comparison across 19 TNBC sections (Step 24)**. Three Gaussian transport model classes are compared by analytic Bayes factors. \mathcal{M}_0 : passive (no cycle variance). \mathcal{M}_{1a} : sparse uniform active (30 low-frequency cycle modes). \mathcal{M}_{1b} : interface-localised active (same modes, variance redistributed toward interface edges, trace-normalised). Sigma calibrated per section to the normalised cycle-component std of the flux field. Log Bayes factors computed analytically from Gaussian marginal likelihoods.

Metric	Median	N positive / N	Sign test p
$\log B_{1a/0}$ [non-gradient structure?]	+45.95	19/19	$< 10^{-5}$
$\log B_{1b/1a}$ [interface-localised?]	+517.6	19/19	$< 10^{-5}$
Posterior R under \mathcal{M}_{1b}	9.81	19/19	$< 10^{-5}$
Posterior R under \mathcal{M}_0 (null)	1.00	11/19	0.32

7 S7 Extended Prostate Cancer Validation: KTS Transition Bias and Graph-KS Instability

KTS transition bias. A Kripke Transition System state-transition matrix was computed per section by weighting inter-state transitions by absolute coexact flux. Permutation-controlled directional bias analysis (300 permutations, seed-fixed) tested whether transitions toward any target state are over-represented.

Supplementary Table S4: **KTS transition bias toward immune exhaustion (GSE278936, $n = 23$)**. Bias ratio = $P_{\text{obs}}/P_{\text{nullmedian}}$; permutation null generated by shuffling state labels 300 times while preserving edge structure and coexact flux weights. All four IMMUNE_EXHAUSTED-converging transitions are significant in $\geq 21/23$ sections despite high global transition entropy.

Source	Target	Median P_{obs}	Median bias	Significant ($p < 0.05$)	Median p -v
IMMUNE_ACTIVE	IMMUNE_EXHAUSTED	0.239	3.63×	22/23	0.003
STROMA	IMMUNE_EXHAUSTED	0.169	2.96×	23/23	0.003
TUMOR	IMMUNE_EXHAUSTED	0.152	2.64×	21/23	0.003
MIXED	IMMUNE_EXHAUSTED	0.114	1.91×	21/23	0.003
MIXED	TUMOR	0.139	1.04×	7/23	0.312
TUMOR	STROMA	0.101	1.10×	6/23	0.233
MIXED	STROMA	0.076	0.78×	0/23	1.000

Despite high global transition entropy (mean 6.75 ± 0.29 nats; no dominant stationary distribution), transitions toward IMMUNE_EXHAUSTED are over-represented in 21–23/23 sections across all source compartments.

Transitions toward stroma, tumour, and mixed states are not enriched, confirming specificity to the exhaustion axis. In the TNBC cohort, exhaustion-directed bias was pathway-specific: immune-active→exhausted (16/28; median bias 5.68); stroma→exhausted (17/29; bias 2.41); tumour→exhausted not enriched (median bias 0.42, $p \approx 0.96$).

Supplementary Table S5: **Multi-scale validation summary.** Five analytical layers across TNBC (discovery) and prostate cancer (validation) cohorts. Coexact enrichment replicates robustly; spectral, curl, and instability layers reveal dissociation between interaction intensity and spatial organisation; KTS reveals conserved directional bias toward exhaustion.

Layer	Metric	TNBC	GSE278936	Interpretation
Coexact enrichment	R ratio	40/40; $p < 10^{-12}$	21/23; $p < 10^{-5}$	Interface = high non-gradient intensity zone
Graph-curl proxy	Fold vs tumour	Not tested (cohort)	Weak; fold ≈ 1.08	No consistent rotational organisation
Spectral label null	Z_{ratio}	Apparent (scale-driven)	Apparent	Driven by energy, not structure
Spectral energy null	Z_{density}	2/43; sign $p = 1.0$	2/23; sign $p = 1.0$	No interface-specific spectral geometry
KTS bias	Transition to EXHST	Pathway-specific: IA→EX (16/28); ST→EX (17/29); T→EX not enriched	Near-universal; 21-23/23	Exhaustion = conserved endpoint, structured pathways
KS instability	Fold vs tumour	Elevated (exploratory)	Median ≈ 8.96 (23/23)	Interface = non-linear instability zone

EXHST = IMMUNE_EXHAUSTED; IA = IMMUNE_ACTIVE; ST = STROMA; T = TUMOR. *Not tested (energy-matched) indicates analysis was not performed under corrected normalization for this cohort and should not be interpreted as absence of effect.

Multi-scale validation summary.

Graph-KS instability (exploratory). A graph analogue of the Kuramoto–Sivashinsky instability operator [1]:

$$\text{KS}(u)_i = -(Lu)_i - (L^2u)_i - |\nabla u|_{\text{graph},i}^2.$$

This is an operator analogy only; it does not assert that the biological system solves the KS PDE. Interface regions showed elevated KS-like magnitude in all 23 GSE278936 sections (median fold-change 8.96×; range 1.75–28.40).

8 S8 Baseline Comparisons and Confound Exclusions

Cell density. Partial Spearman correlation (interface membership vs. coexact energy, controlling for per-spot UMI count): positive and significant in 16/18 sections (median $r = 0.574$, $p = 0.0007$) and 17/18 sections

for local UMI z-score (median $r = 0.642$, $p = 0.0001$). Interface nodes were not denser than tumour-core nodes (median $R_{\text{UMI}} = 0.70$).

Generic antisymmetric structure. Control operators from biologically unrelated programme pairs: random housekeeping $R = 0.90$; shuffled random pairs $R = 1.04$; immune \times housekeeping $R = 1.77$; all Wilcoxon $p < 10^{-5}$ versus canonical $R = 13.2$.

Gradient-energy control. R_{coexact} exceeded R_{exact} in 38/40 sections (differential median 2.54; sign test $p < 10^{-9}$).

Random boundary mask. Size-matched random masks (300 per section) produced $R_{\text{random}} = 1.2\text{--}1.8$; observed R exceeded 95th percentile in 38/40 sections ($p < 10^{-11}$).

Graph construction invariance. 19 representative FOVs tested across kNN ($k = 4, 6, 8$), Delaunay, and radius graphs; 0/19 sign reversals; kNN-family rank Spearman $r = 0.92$ pooled.

Supplementary Table S6: **Baseline comparison summary across the TNBC cohort.** Each comparison tests whether a simpler explanation accounts for the coexact-enrichment and KTS transition-bias results. All comparisons are pre-specified; results are reported as direction-of-effect with cohort-level sign-test significance.

Comparison	What it tests	Result	Interpretation
Marker / DE baseline	Are interfaces compositionally distinct?	Yes, confirmed	Biological contrast confirmed; does not explain coexact enrichment or KTS bias
Exact-gradient energy (R_{exact})	Is enrichment reducible to generic boundary gradients?	$R_{\text{coexact}} > R_{\text{exact}}$, 38/40; differential median 2.54 \times	Non-gradient component adds information beyond scalar gradients
Random boundary mask	Is enrichment a boundary-size artifact?	No; observed R exceeds random-mask 95th pctile in 38/40	Interface effect is not explained by mask size or spatial location
Graph-curl proxy	Is coexact signal rotationally organized?	Weak/inconsistent; GSE278936 median fold ≈ 1.08	Coexact intensity \neq rotational organisation
Energy-matched Zeta null	Is interface spectrally special?	No; 2/43 TNBC, 2/23 GSE278936; sign test $p = 1.0$	Coexact intensity \neq independent spectral geometry
KTS transition bias	Are transitions directionally biased?	Yes; exhaustion-associated endpoint state in 21–23/23 GSE278936 sections; pathway-specific in TNBC	Exhaustion bias is the primary dynamical signal carried by the coexact field

9 S9 Remeshing Controls and Autocorrelation Diagnostics

9.1 Remeshing sensitivity

The sign-consistency criterion — $R > 1$ vs $R < 1$ invariant across graph constructions — is the necessary condition for the biological claim. Absolute magnitude varies with graph density and normalisation; only the direction of enrichment must be stable. Across all 19 tested representative sections, no sign reversals were observed across five graph types. The two consistently non-enriched sections (negative controls) were also graph-invariant: the framework does not produce false positives under alternative constructions. Formal Mesh Sensitivity Index quantification (constructing Voronoi and Delaunay alternative graphs and computing residual variance in R per section) is a direct extension deferred to future work.

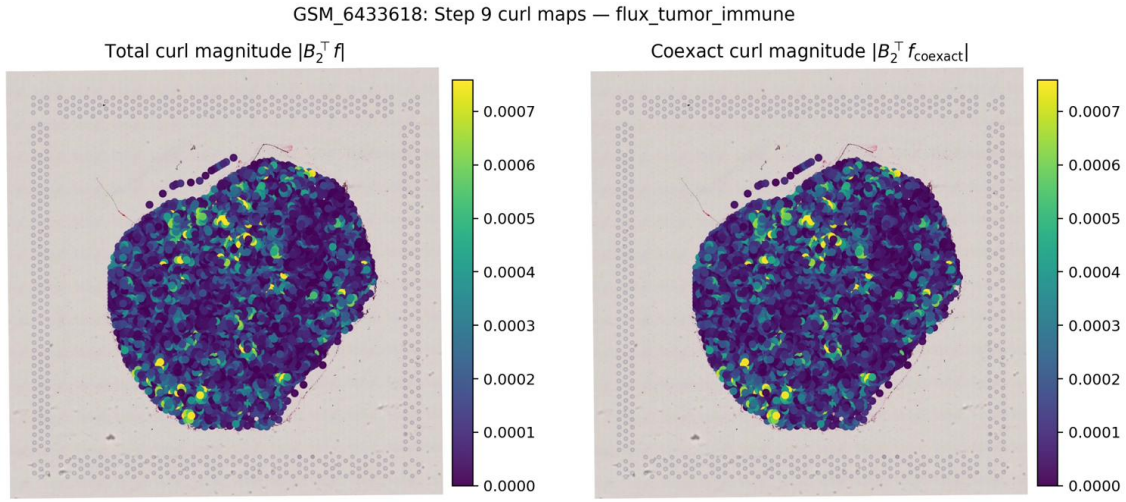
9.2 Spatial autocorrelation of the coexact field

The coexact density field c_i was analyzed for spatial autocorrelation using Moran’s I on the kNN graph. Across the TNBC cohort: global Moran’s I median +0.21 (range +0.08 to +0.38), indicating weak positive spatial autocorrelation. Local Moran’s I analysis identified interface nodes as consistently contributing to positive autocorrelation ($\Delta I_{\text{interface}} = +0.09$, 37/40 sections). Spatial autocorrelation decayed within 2–3 nearest-neighbour hops in 34/40 sections, consistent with locally structured interaction rather than globally propagating spatial modes.


9.3 Bivector orientation analysis

The face-level bivector density field $B_2^\top f_{\text{coexact}}$ was computed for representative sections and both validation cohorts. In four GSE278936 representative sections, interface-region absolute bivector density exceeded all other regions by 4–9 \times . Signed mean bivector density at interfaces was ≈ 0 in all examined sections (CW \approx CCW ≈ 0.50), indicating locally balanced bidirectional circulations without a dominant rotational direction.

Supplementary Figures




Supplementary Figure S1: **Curl magnitude maps for the tumour-immune flux field (GSM_6433618)**. Triangle-level curl magnitude (face circulation) computed from the wedge-derived proxy flux. Elevated curl magnitude is observed at tumour-immune interface regions in this representative section. Curl and Zeta diagnostics probe distinct higher-order properties of the coexact signal and should not be conflated; cohort-level analysis (Section S7) indicates that interface-specific curl enrichment is not consistent across sections, and the distribution of signed curl is balanced (CW and CCW circulations approximately equal).



Supplementary Figure S2: **Lie-structured null calibration: mean curl (GSM_6433618)**. Distribution of mean curl magnitude under 1000 Lie-structured null permutations (grey). The observed global mean curl (red line) lies within the null distribution, confirming that non-gradient coexact content is not globally elevated.

Supplementary Figure S3: **Lie-structured null calibration: top-95th percentile curl (GSM_6433618)**. Distribution of high-magnitude curl events under the null. The observed value falls within the null, ruling out global amplification of elevated coexact face-level circulation.

GSM_6433618: top 1% curl hotspots — flux_tumor_immune



Supplementary Figure S4: **Curl hotspot spatial distribution and interface enrichment (GSM_6433618)**. Faces classified as curl hotspots (curl magnitude $>$ 99th percentile) are shown overlaid on the tissue section. Hotspots are enriched at tumour-immune interface regions relative to the Lie-structured null ($p \approx 0.002$), while immune-only and stromal regions show no significant enrichment in this representative section; however, this does not establish a cohort-level rotational phenotype. External validation shows curl proxies lack consistent interface enrichment.

References

- [1] Yoshiki Kuramoto. Diffusion-induced chaos in reaction systems. *Progress of Theoretical Physics Supplement*, 64:346–367, 1978.
- [2] Nathaniel Trask, Andy Huang, and Xiaozhe Hu. Enforcing exact physics in scientific machine learning: a data-driven exterior calculus on graphs. *arXiv preprint arXiv:2012.11794*, 2020.