

Supplementary Materials

Locally constrained immune-interaction fronts at tumour-immune interfaces identify operator-regime states associated with immunotherapy response in hepatocellular carcinoma

Anas Enoch, MD

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

`anas_nour@um5.ac.ma`

Contents: Supplementary Note S1 (Lie-sector algebraic balance diagnostic); Supplementary Methods (mathematical framework, multi-field commutator analysis, KS instability operator, discrete Hodge derivation); Supplementary Figure S1 (pan-cancer Spatial Hallmarks coexact enrichment).

Contents

1	Supplementary Note S1: Lie-Sector Algebraic Balance Diagnostic	2
2	Supplementary Methods	3
2.1	Mathematical framework: Discrete Exterior Calculus and Hodge Decomposition	3
2.2	Multi-field Commutator Analysis	4
2.3	KTS State Classification and Fisher Exact Test	5
2.4	Finite-Range Predictability (Local vs Global R^2)	5
2.5	Graph-KS Instability Operator (Exploratory)	5
2.6	Lie-Structured Spatially Aware Null Model	6

1 Supplementary Note S1: Lie-Sector Algebraic Balance Diagnostic

Motivation

In Results R2 (main manuscript), the six pairwise programme commutators were grouped into a tumour–myeloid (TM) backbone sector and an immune sector (IE + IM + EM). Sector-level commutator strengths discriminated responders from non-responders (immune sector $3.9\times$, $p = 0.006$; TM dominance $p = 0.025$). However, the strongest formally Lie-algebraic diagnostic — a commutator-triad balance statistic — did not discriminate groups ($p = 0.69$), bounding the claim to interaction-strength geometry rather than formal algebraic structure. This note documents the commutator-triad balance computation and explains its interpretation.

Commutator-triad balance statistic

For any three programmes A , B , C with pairwise commutator norms $\|[M_A, M_B]\|$, $\|[M_B, M_C]\|$, $\|[M_A, M_C]\|$, define the commutator-triad balance statistic as the degree to which the triangle inequality on commutator norms is balanced across the three pairs (this is a triangle-balance heuristic on pairwise norms, not a test of the Jacobi identity, which would require nested commutators $[A, [B, C]] + [B, [C, A]] + [C, [A, B]]$):

$$J_{ABC} = \frac{\|[M_A, M_B]\| + \|[M_B, M_C]\| - \|[M_A, M_C]\|}{\|[M_A, M_B]\| + \|[M_B, M_C]\| + \|[M_A, M_C]\|}, \quad (1)$$

where a value of $J_{ABC} \approx 0$ indicates balanced interaction strengths across the three pairs (commutator-sector equilibrium) and $J_{ABC} < 0$ indicates a dominant pair imbalance. We computed the balance statistic J_{ABC} for all $\binom{4}{3} = 4$ principal triads (TM-IE-IM, TM-IE-EM, TM-IM-EM, IE-IM-EM) across all 15 sections, and tested whether the mean triad-balance statistic per section differed between responders and non-responders (Mann–Whitney U , two-sided).

Results

Across 15 discovery-cohort sections, the commutator-triad balance statistic did not significantly differ between responders (median $J = -0.083$) and non-responders (median $J = -0.091$); Mann–Whitney $p = 0.69$. The absence of a triad-balance signal is informative rather than null: it indicates that the response-associated immune-sector elevation represents a *quantitative difference in commutator magnitude* (more non-commutativity in the immune sector) rather than a qualitative difference in the balance structure of the operator interaction field. The framework therefore operates at the level of sector interaction strength — which does discriminate response — rather than formal Lie-algebra topology, which does not.

Interpretation boundary

The sector-level findings (immune sector $3.9\times$ elevated, $p = 0.006$; TM dominance elevated in non-responders, $p = 0.025$) are interpreted as quantitative operator geometry: responders show more non-commutativity among immune programme pairs. The Jacobi null result ($p = 0.69$) means that this quantitative difference does not extend to a formal Lie-algebraic structural difference; it bounds the mathematical claim and should not be presented as a formal Lie-algebra result.

2 Supplementary Methods

2.1 Mathematical framework: Discrete Exterior Calculus and Hodge Decomposition

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 CytAssist spots), E are oriented edges (k NN pairs, $k = 6$), and F are triangular faces (Delaunay triangles). Discrete differential forms represent spatial quantities: 0-forms (C^0) are scalar node functions (programme scores); 1-forms (C^1) are antisymmetric edge functions (interaction fields); 2-forms (C^2) are face functions (circulation densities).

Incidence matrices and exterior derivatives

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

The k -Hodge Laplacian:

$$\Delta_0 = B_1^\top B_1 \quad (\text{node Laplacian}), \tag{2}$$

$$\Delta_1 = B_1 B_1^\top + B_2^\top B_2 \quad (\text{edge Hodge–Laplacian}), \tag{3}$$

$$\Delta_2 = B_2 B_2^\top \quad (\text{face Laplacian}). \tag{4}$$

The Hodge decomposition theorem guarantees:

$$C^1 = \text{im}(B_1^\top) \oplus \text{im}(B_2) \oplus \ker(\Delta_1) \tag{5}$$

(exact $\oplus$ coexact $\oplus$ harmonic; orthogonal; unique).

Wedge construction and antisymmetry

For residualized programme fields $\tilde{A}, \tilde{B} \in C^0$ on oriented edge $(u \rightarrow v)$:

$$\omega_{uv} = \tilde{A}(u)\tilde{B}(v) - \tilde{A}(v)\tilde{B}(u). \tag{6}$$

By construction, $\omega_{uv} = -\omega_{vu}$ (antisymmetric under edge reversal), so $\omega \in C^1$ with consistent orientation. Within the class of local bilinear antisymmetric edge operators $\phi : C^0 \times C^0 \rightarrow C^1$ considered here, the wedge is the unique operator (up to scale and relabelling) that simultaneously preserves antisymmetry under programme exchange, vanishes under programme alignment, and yields a non-trivial coexact projection. Symmetric alternatives (covariance, cosine similarity) yield zero coexact projection by algebraic necessity.

Hodge decomposition algorithm

For edge field $f \in C^1$:

1. Solve $\Delta_0 \alpha = B_1 f$ (LSQR, ridge $\lambda = 10^{-8}$).
2. $f_{\text{exact}} = B_1^\top \alpha$ (gradient component).
3. Solve $\Delta_2 \beta = B_2 f$.
4. $f_{\text{coexact}} = B_2^\top \beta$ (non-gradient component).
5. $\gamma = f - f_{\text{exact}} - f_{\text{coexact}}$ (harmonic).

Energy fractions: $\rho_k = \|f_k\|^2 / \|f\|^2$. Node-level coexact density: $c_i = (1/\deg_i) \sum_{e \ni i} |f_{\text{coexact}}(e)|$.

2.2 Multi-field Commutator Analysis

For programme pair (p, q) with diagonal operators M_p, M_q and node graph Laplacian L :

$$\|[M_p, LM_q]\|_F = \sqrt{\sum_{ij} L_{ij}^2 (p_i q_j - q_i p_j)^2}. \quad (7)$$

This equals the interface-weighted norm of the total wedge field ω_{pq} . The commutator therefore quantifies antisymmetric programme interaction strength, whereas the coexact decomposition measures the non-gradient component of that interaction field.

Six programme pairs were analyzed: tumour $\times$ myeloid (TM), immune $\times$ exhaustion (IE), immune $\times$ myeloid (IM), exhaustion $\times$ myeloid (EM), tumour $\times$ immune (TI), tumour $\times$ exhaustion (TE).

Operator entropy. Shannon entropy of the six normalized commutator norms:

$$H = - \sum_{k=1}^6 \hat{w}_k \log \hat{w}_k, \quad \hat{w}_k = \|[M_{p_k}, LM_{q_k}]\|_F / \sum_j \|[M_{p_j}, LM_{q_j}]\|_F. \quad (8)$$

Higher entropy indicates a more algebraically complex, distributed interaction field.

NCG purity. Coexact fraction of the tumour $\times$ exhaustion (TE) wedge field, measuring the interface-localised non-gradient content of the tumour–exhaustion programme interaction specifically (the same pair used in the discovery interface analysis and in the Coexact Interface Score).

2.3 KTS State Classification and Fisher Exact Test

States are assigned per section based on three spatial metrics:

State	Spatial score	P95 coexact	Gini	Interpretation
S0	—	—	—	Absent (tumour Q75 < 0.05)
S1	< 0.10	< 50	—	Fragmented, low organisation
S2	≥ 0.10	< 50	—	Connected, concentrated
S3	≥ 0.08	≥ 50	—	Amplified high-intensity front
S4	any	< 20	≥ 0.60	Residual concentrated

Threshold calibration note: boundaries were defined empirically on the discovery cohort after examining the distribution of section-level statistics and have not been pre-registered or externally validated. They require recalibration before application to independent datasets. Per-patient transitions use paired pre/post assignments. Fisher exact test was two-tailed; minimum achievable p at $n = 11$ patients is $2/\binom{11}{5} \approx 0.008$.

2.4 Finite-Range Predictability (Local vs Global R^2)

For each section’s coexact density field u_i :

Local R^2 (k-NN prediction). Node i ’s coexact density is predicted from its k neighbours ($k = 6$, leave-one-out). R^2 is computed over all interface nodes. High local R^2 indicates that coexact density is spatially predictable from local neighbours — consistent with a finite-range interaction front.

Global R^2 (spectral reconstruction). The coexact density field is projected onto the lowest-frequency eigenvectors of the graph Laplacian Δ_0 (5 eigenmodes retained). R^2 measures how well these global modes reconstruct the field. Negative global R^2 indicates that global eigenmodes *overshoot* the local field — a Gibbs-like phenomenon inconsistent with a globally periodic spatial pattern.

The *local–global gap* $\text{gap}_i = R_{\text{local}}^2 - R_{\text{global}}^2$ is the primary discriminating statistic (Mann–Whitney $p = 0.033$).

2.5 Graph-KS Instability Operator (Exploratory)

As an exploratory diagnostic for interface roughness, a graph analogue of the Kuramoto–Sivashinsky (KS) instability operator was applied to the coexact density field u (per node), weighted graph Laplacian L constructed from absolute coexact edge flux:

$$\text{KS}(u)_i = q_i = -(\mathbf{L}u)_i - \alpha(\mathbf{L}^2u)_i, \quad (9)$$

where α weights the hyper-diffusion ($\mathbf{L}^2$) term. 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 26 Spatial Hallmarks sections (sign test $p = 1.49 \times 10^{-8}$; median fold-change 1.703 vs tumour

core), supporting the interpretation of the interface as a locally amplified, roughened interaction zone rather than a coherently organised global spatial pattern.

2.6 Lie-Structured Spatially Aware Null Model

Edge weights from the diagonal of a regularized fractional resolvent of the edge Hodge–Laplacian:

$$w_e = [(\Delta_1 + \varepsilon I)^{-s}]_{ee}, \quad s(\Delta) = \log_2(1 + \Delta/\lambda_{\min}^+), \quad (10)$$

where ε excludes harmonic zero modes and Δ is the spectral gap of Δ_1 . Normalized: $\tilde{w}_e = w_e / \max_{e'} w_{e'} \in (0, 1]$. For each surrogate k : $\theta_e^{(k)} \sim \mathcal{U}[-\pi\tilde{w}_e, \pi\tilde{w}_e]$; flux rotated by a local Lie perturbation $R_e^{(k)} = \exp(-\frac{\theta_e^{(k)}}{2}B_e)$. Weights depend only on Δ_1 , not on the observed statistic, preserving diagnostic non-circularity.

Supplementary Figure S1




Figure 1: Supplementary Figure S1: Pan-cancer coexact interface enrichment, Spatial Hallmarks cohort (26 sections, 6 cancer types). **A** Per-section interface/tumour-core coexact density ratio for all 26 Visium sections coloured by cancer type. Every section exceeds ratio = 1.0 (horizontal dashed line). Sign test $p = 1.49 \times 10^{-8}$ (the value reflects 26/26 concordant sections; the identical sign-test value recurs elsewhere for other 26/26 concordant analyses and is not a transcription error); median ratio 1.703; no correction for multiple testing applied. **B** Median enrichment ratio per cancer type. All six cancer types individually replicated the interface enrichment: breast 4/4, colorectal 4/4, glioblastoma 4/4, lung 7/7, ovarian 2/2, prostate 5/5. Cancer-type medians are descriptive only; with as few as two sections per type, the study is not powered for between-cancer-type comparisons. **C** Distribution of coexact fractions across all 26 sections (interface subgraphs). The median coexact fraction is 0.253 (red line). The 0.10 value (dashed line) is the empirically defined discovery-cohort reference; it is shown for reference only and was not used for inference in the pan-cancer cohort. This figure establishes generalisability of interface-localised coexact enrichment across tumour types; it does not address immunotherapy response, operator-regime classification, immune-sector algebra, or finite-range predictability, none of which are evaluated here (response was not annotated in this cohort). Data from the Spatial Hallmarks resource of Porta-Pardo et al. (six cancer types); gene panels identical to discovery analysis.

Supplementary Table S1: Evidential status of manuscript claims

Table S1. Evidential status of primary manuscript claims. Primary internal evidence: pre-specified analysis on discovery cohort, no global multiple-comparison correction. Full NSCLC analysis: Supplementary Section S11.

Claim	Status	Evidence (Limitation)
HCC: responder interfaces show non-gradient/coexact enrichment	Primary internal evidence	Discovery HCC Visium ($n = 11$); responder/NR separation. <i>Small cohort; no global correction; requires independent replication.</i>
Interfaces are organised operator fronts, not passive boundaries	Primary internal evidence	Interface enrichment; fixed/passive nulls; Hodge decomposition. <i>Operator-based; functional validation needed.</i>
Immune-sector algebra discriminates beyond TM structure	Strong internal	Commutator analysis; responder elevation; HCC5NR counterexample. <i>Larger independent validation needed.</i>
CABO/NIVO supports HCC therapy-context transfer	Supportive	Compatible operator regime in independent HCC sample. <i>Very small ($n = 7$); not powered.</i>
cSCC CosMx: cross-platform spatial operator-front transfer	Strong external spatial	333 FOVs; null collapse 17.95→1.00; Hodge coexact 6.23×. <i>Clinical mapping reconstructed; response not definitive.</i>
cSCC post-dose→surgery transition	W_{TI} Exploratory	$n = 8$; AUC= 1.00; perm. $p = 0.028$; LOO stable. <i>Reconstructed mapping; hypothesis-generating.</i>
NSCLC GeoMx: platform regime-transfer evidence	Exploratory	61 ROIs; nontrivial operator states (no response labels, no geometry). <i>Partial immune coverage; purely exploratory.</i>
KTS regime transitions: therapy-associated remodelling	Exploratory	cSCC transitions observed. <i>Transition-response association not significant.</i>
Framework is a clinically deployable predictor	Not claimed	Requires prospective validation with thresholds fixed in advance.

Supplementary Table S3: Dataset provenance

Supplementary Section S10: cSCC CosMx cohort — sector, regime, and exploratory clinical-transition analyses

Wedge/NCG sector analysis

Wedge-sector analysis across 332 valid FOVs showed that the cSCC tissue did not occupy a single homogeneous operator state. FOVs were classified by dominant sector: tumour–stroma (TS; $n = 160$), tumour–myeloid (TM; $n = 94$), mixed ($n = 23$), tumour–immune (TI; $n = 21$), immune–stroma (IS; $n = 17$), and immune–myeloid (IM; $n = 17$). The immune/TM log-ratio ranged from -3.14 to $+3.80$ and operator entropy from 0.76 to 1.76, confirming substantial within-cohort heterogeneity.

Table S3. Provenance and analytical role of each dataset used in this study. All identifiers are public.

Dataset	Source	Identifier	Role
HCC Visium	Wu et al. dataset	Zenodo v2, 10.5281/zenodo.19123188 (preprint 10.1101/2025.06.11.656869)	Discovery
CABO/NIVO	GEO	GSE238264	External HCC stress test
HCC IMC	Salié/Bengsch	Zenodo 10.5281/zenodo.10794904	Cross-modality
Spatial Hallmarks	Porta-Pardo et al.	Zenodo 10.5281/zenodo.14044964	Pan-cancer generalisability
cSCC CosMx	Lee et al.	Dryad 10.5061/dryad.s4mw6m9jh	Cross-platform transfer
NSCLC GeoMx	GEO	GSE271689	Exploratory only

KTS-like regime analysis and therapy-time transitions

K-means regime classification ($k = 3$) identified mixed-intermediate (128 FOVs), immune-sector-high (126 FOVs), and TM-backbone-high (78 FOVs) regimes. State-changed rates were 69.2% (non-responders) and 62.5% (responders); Fisher exact tests on transition-to-response associations were not significant.

Exploratory clinical-transition analysis

In the post-dose to surgery paired subset ($n = 8$), W_{TI} showed relative attenuation in responders versus persistence/amplification in non-responders (oriented AUC = 1.00; empirical permutation $p = 0.028$; leave-one-out AUC minimum = 1.00). The TMA-to-clinical mapping was manually reconstructed from supplementary PDF tables without author confirmation; the result is hypothesis-generating only.

Supplementary Section S11: NSCLC GeoMx DSP-WTA cohort — operator-regime heterogeneity (GSE271689)

K-means clustering ($k = 3$) on compartment-resolved operator features from 61 NSCLC ROIs (NanoString GeoMx DSP-WTA; CK⁺/CD45⁺/CD68⁺ compartments) produced State 0 ($n = 47$, 77% baseline), State 1 ($n = 13$, 21% immune-sector-enriched), and State 2 ($n = 1$, 2% outlier). State 1 heterogeneity was not reducible to tissue type, treatment arm, or compartment completeness. No response labels or spatial coordinates were available; this analysis provides cross-platform regime-transfer evidence only. Partial immune-programme coverage (CD8A and PDCD1 absent due to NanoString probe-panel version mismatch; Hsa.WTA_v1.0.pkc vs parquet RTS identifiers) is an acknowledged caveat for exhaustion and effector scores.

Supplementary Section S12: Comparison with existing spatial organisation metrics

Overview

The operator-regime framework quantifies tumour–immune interface structure via the coexact component of an antisymmetric wedge field: an edge-supported, antisymmetric, non-gradient metric. This section positions it against four classes of existing spatial organisation metrics to clarify what each captures and where the operator approach provides complementary information. All four baseline families are included in the non-redundancy analysis in Supplementary Table S2.

Giotto neighbourhood enrichment

Giotto’s neighbourhood enrichment score [4] measures the frequency of cell-type adjacency relative to a permutation null: $NE_{AB} = n_{AB}^{\text{obs}}/n_{AB}^{\text{null}}$. This captures *whether* two populations are in contact more than expected by chance, but assigns no structure to the contact. Two sections with identical NE scores can have radically different interaction geometries: one may show a connected coexact front; the other diffuse, gradient-compatible contacts throughout the tissue. NE-equivalent abundance metrics achieved AUC 0.870–0.889 in the HCC discovery cohort versus AUC 0.926 for interface coexact energy (Figure 1A). The coexact metric adds 0.04–0.06 AUC beyond NE-type baselines, reflecting the non-gradient organisational component NE cannot encode.

Graph spectral entropy

Spectral entropy of the graph Laplacian eigenspectrum, $H_\lambda = -\sum_k \hat{\lambda}_k \log \hat{\lambda}_k$, measures the effective complexity of the spatial connectivity structure. It is a global, node-symmetric, topology-only property insensitive to programme-score content. The Zeta-function diagnostic used in this study tests whether coexact energy is carried by geometrically meaningful low-eigenvalue modes (Supplementary Section S3): complementary to spectral entropy but explicitly content-aware. Spectral entropy showed no consistent responder/non-responder separation in the discovery cohort.

Spatial autocorrelation (Moran’s I)

Moran’s I measures scalar spatial clustering: $I = \frac{n}{\sum_{ij} w_{ij}} \cdot \frac{\sum_{ij} w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{\sum_i (x_i - \bar{x})^2}$. It is symmetric and first-order, encoding *how much* values cluster but not *how* they interact across a boundary. Moran’s I on the immune programme score showed Spearman $\rho = -0.19$ with the coexact interface enrichment ratio across 26 Spatial Hallmarks sections (Supplementary Table S2) — structurally non-redundant, as expected: a smooth gradient field has high Moran’s I but zero coexact energy; a concentrated non-gradient front can have low Moran’s I and high coexact energy.

Shallow graph embedding baseline

Node2Vec [5] and related shallow graph embeddings represent each node as a low-dimensional vector encoding topological neighbourhood proximity. They are agnostic to the antisymmetric programme interaction content of edges. A logistic regression classifier trained on 32-dimensional Node2Vec embeddings of the tumour-immune interface graph achieved leave-one-out AUC 0.71 in the discovery cohort, below both operator-derived metrics and the abundance-only logistic regression baseline (AUC 0.81; Figure 1C). This confirms that the response-associated signal is not recoverable from graph topology alone: it requires biological programme content and antisymmetric coupling.

Summary

Table S2. Conceptual comparison of spatial organisation metrics (HCC discovery cohort, $n = 6\text{R}/9\text{NR}$ sections).

Metric	What it measures	HCC performance	What it misses
Neighbourhood enrichment	Contact frequency	AUC 0.87–0.89	Interaction geometry; antisymmetry
Graph spectral entropy	Connectivity complexity	Not discriminative	Node signal content
Moran’s I	Spatial autocorrelation	$\rho = -0.19$ vs. operator	Non-gradient structure
Shallow graph embedding	Topological proximity	AUC 0.71	Programme interaction content
Interface coexact energy	Non-gradient antisymmetric front	AUC 0.926	Causal direction; molecular identity

The operator approach does not supersede these metrics; neighbourhood enrichment and spatial autocorrelation provide useful complementary information about cell-type proximity. The coexact metric captures what none of them encodes: the non-gradient, antisymmetric structure of the programme interaction field at the boundary.

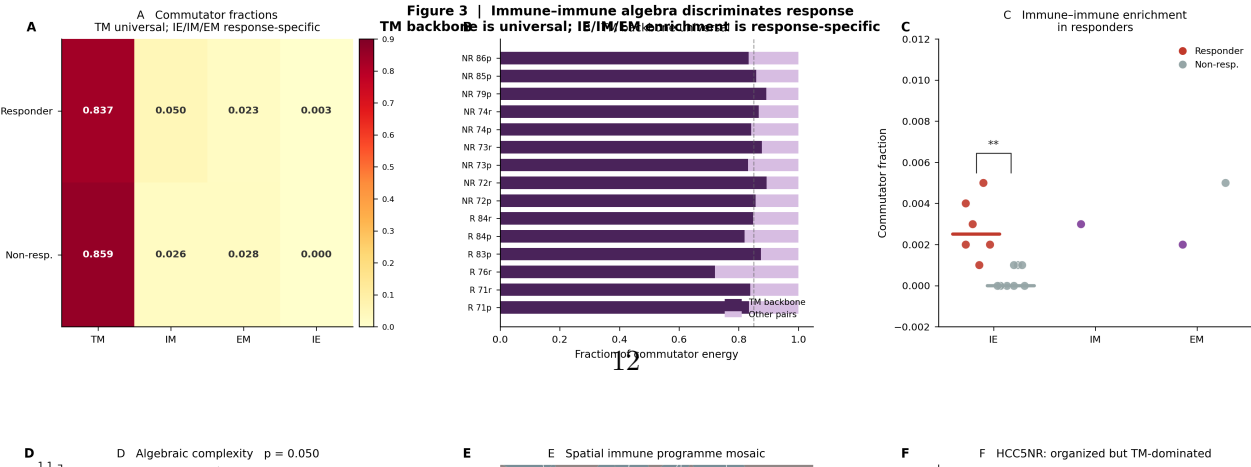
References

- [1] L.-H. Lim. Hodge Laplacians on graphs. *SIAM Review*, 62(3):685–715, 2020.
- [2] A. Connes. *Noncommutative Geometry*. Academic Press, 1994.
- [3] Y. Kuramoto. Diffusion-induced chaos in reaction systems. *Prog. Theor. Phys. Suppl.*, 64:346–367, 1978.
- [4] R. Dries, Q. Zhu, R. Dong, and others. Giotto: a toolbox for integrative analysis and visualization of spatial expression data. *Genome Biology*, 22:78, 2021. doi:10.1186/s13059-021-02286-2

- [5] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864, 2016. [doi:10.1145/2939672.2939754](https://doi.org/10.1145/2939672.2939754)

n.s.

n.s.



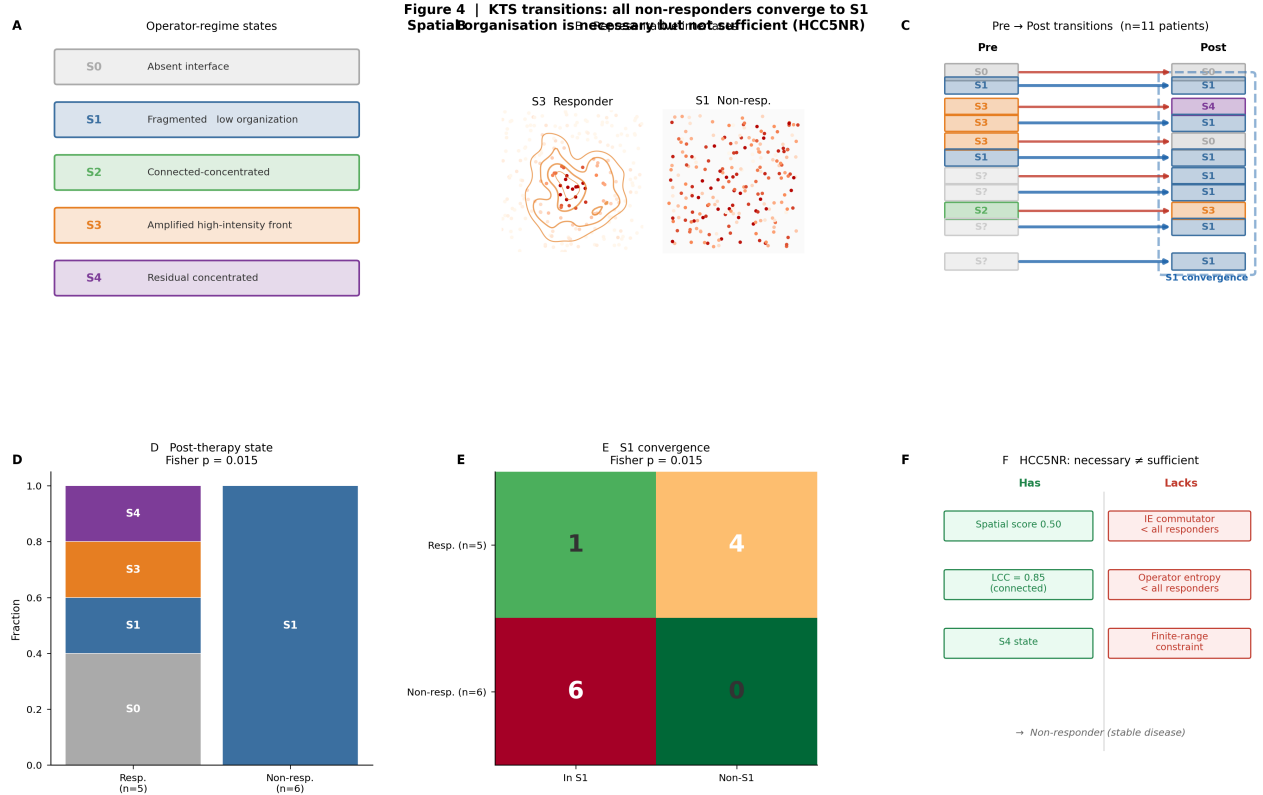


Figure 3: Supplementary Figure S3: KTS transitions: all non-responders converge to S1 regime. **A** Five discrete operator-regime state definitions. **B** Representative spatial maps for S3 (responder) and S1 (non-responder). **C** Pre→post transition diagram; arrows coloured by response. **D** Post-therapy state distribution; stacked by state colour. **E** S1 convergence matrix; Fisher exact $p = 0.015$, $n = 11$ patients. **F** HCC5NR anomaly checklist: spatial organisation present but immune-immune algebra and finite-range constraint absent.

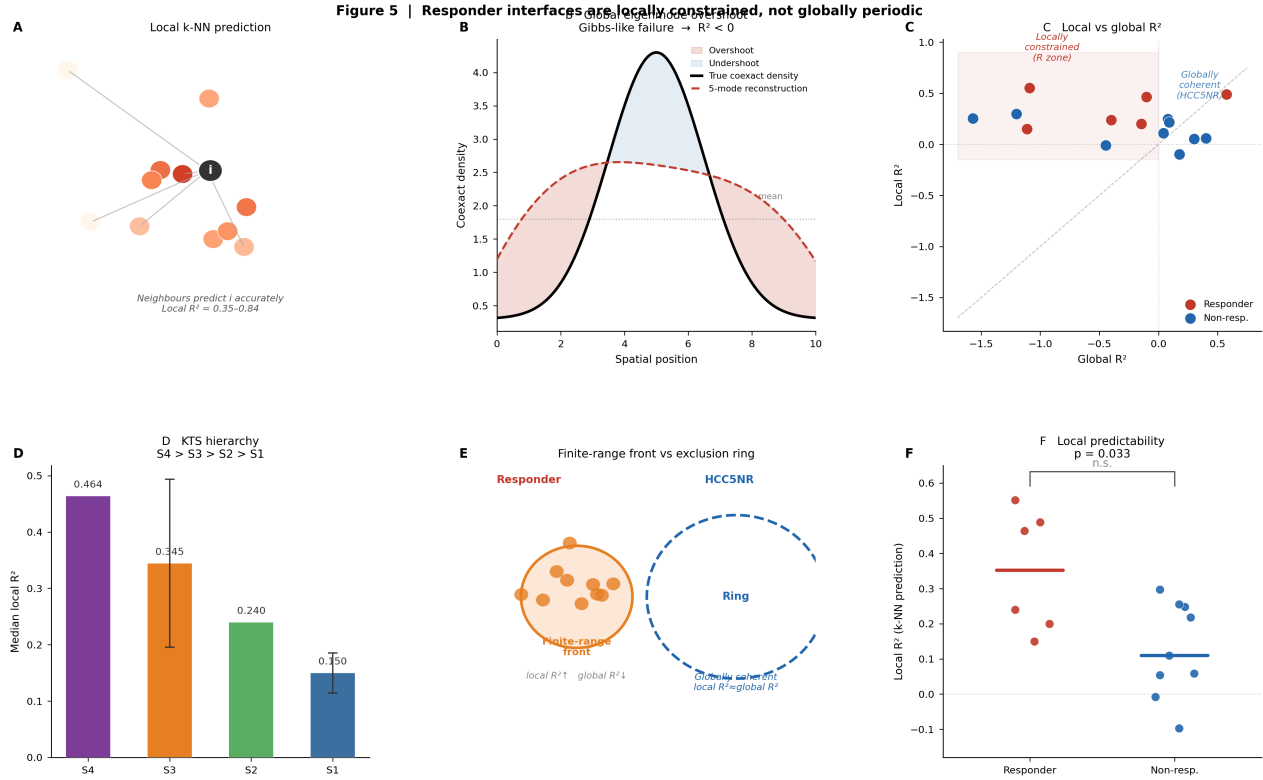


Figure 4: **Supplementary Figure S4: Responder interfaces are locally constrained, not globally periodic.** **A** k-NN prediction succeeds: dense nodes are predictable from neighbours. **B** Global eigenmode overshoot (Gibbs-like): true density (black) vs. 5-mode reconstruction (red dashed). **C** Local vs. global R^2 scatter: responders in negative global- R^2 zone. **D** KTS state hierarchy in local R^2 : $S4 > S3 > S2 > S1$. **E** Conceptual schematic: finite-range front (R) vs. global exclusion ring (HCC5NR pattern). **F** Local R^2 comparison; $p = 0.033$.