# HODGEDEPTH: SPATIAL REPRESENTATION LEARNING THROUGH GLOBAL GEOMETRY

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

A central challenge in spatial representation learning is to connect local feature patterns with global geometric patterns. For example, understanding layered tissues requires connecting local molecular variation to global spatial organization. We introduce *HodgeDepth*, a global-geometry-aware framework that learns an interpretable depth coordinate describing progression across layers and corresponding positions along curved tissue contours. The method aligns noisy local directions of molecular change and integrates them into a coherent coordinate, making global geometry an explicit representation-learning target. The learned depth provides a spatial axis for inspecting tissue contours and gene-expression profiles, with a complementary construction for recurrent layers. Our theory characterizes geometric bias and contour recovery, establishes conditions for correct global ordering from local estimates, and bounds recovery of a canonical depth scale based on cumulative spatial mass. We evaluate HodgeDepth on 23 human cortical Visium sections across two cohorts. On the 12-section DLPFC benchmark, mean absolute depth–layer Spearman correlation increases from 0.690 for GASTON to 0.885. Downstream clustering achieves mean ARI 0.618 versus 0.533 for SEDR, the strongest of nine tested baselines, a 16% relative improvement. Controlled synthetic examples further illustrate recovery of layered geometry. Together, these results support global spatial depth as an interpretable representation of tissue organization with measurable benefits for downstream analysis.

## 1 INTRODUCTION

Spatial representation learning should recover the organization of a tissue as a whole. In layered systems, this requires both an ordering across layers and comparable coordinates along their curved contours. Distant locations can occupy corresponding positions in this organization, while nearby locations can straddle a molecular boundary. Neighborhood-based methods combine expression similarity and spatial proximity for denoising and domain discovery (Long et al., 2023), but these objectives do not uniquely determine a global depth. We study *spatial transcriptomic depth*: an interpretable scalar coordinate learned from molecular measurements and location to describe the tissue's dominant layered geometry.

Layered and continuous spatial organization is already a target of Belayer (Ma et al., 2022), GASTON (Chitra et al., 2025a), and GASTON-Mix (Chitra et al., 2025b). GASTON learns a global isodepth through an integrable scalar neural field, while GASTON-Mix learns domain-specific coordinates. SLOPER estimates gene-specific spatial score fields and derives isodepth from gradient outer products (Wang et al., 2025). We address the geometric estimation problem connecting local differential evidence to a global coordinate: *how can noisy, unoriented local axes be assembled into a coherent tissue-wide depth, and what controls the fidelity of its geometry?*

We introduce HODGEDEPTH, a *global-geometry-aware* framework that reconstructs an ordered tissue foliation from local molecular directions. Under a regular single-index model, feature gradients share a contour-normal axis with no preferred sign. We represent these axes by sign-invariant projectors, fuse them using directional confidence and spatial support, synchronize their orientations, and fit a global scalar potential by discrete Hodge projection. This construction makes compatibility across graph cycles an explicit constraint. Unit-normal increments use spatial displacement to set geometric spacing, separating depth parameterization from the magnitude of molecular change. A source-eikonal extension represents recurrent layers by unsigned distance from an inferred internal sheet.



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

All reported HODGEDEPTH experiments estimate local Jacobians by weighted ridge regression. The general framework also accommodates Jacobians obtained by differentiating neural feature fields, including a shared nonlinear network, within the same fusion, synchronization, and integration pipeline. The recovery analysis applies to either estimator under its stated assumptions on derivative accuracy, geometry, and the graph; neural estimation remains an unevaluated extension of the implemented ridge specialization.

Interpretability is central to the representation. As with biological pseudotime (Trapnell et al., 2014), a scalar axis organizes molecular variation; here the axis describes spatial tissue organization. Depth contours support comparisons across curved layers, while expression profiles along depth support investigation of spatial markers and transitions. This benefit is shared with isodepth methods such as GASTON. HODGEDEPTH also exposes the local directions, their confidence, and integration residuals, allowing users to inspect how molecular evidence determines the coordinate and where it disagrees with a single scalar geometry.

Our contributions are:

- **Global geometry from local molecular axes.** We formulate spatial depth as reconstruction of an ordered tissue foliation and develop a confidence-weighted estimator that couples projector fusion, orientation synchronization, and global Hodge integration.
- **Theory linking local evidence to global depth.** We establish geometric bias and contour recovery, full-pipeline depth-order recovery, and canonical mass-coordinate recovery for a specified orientation and spatial measure. The analysis separates geometric model bias, directional estimation error, and sampling discrepancy. A population extension identifies a positive integration speed under a prescribed mass law, enabling contour recovery beyond transnormal geometry.
- **Interpretable depth and empirical validation.** Two human cortical cohorts comprising 23 Visium sections evaluate downstream domain recovery. The main DLPFC study also assesses continuous anatomical ordering and label-free spatial coherence, connecting the learned coordinate's geometric interpretation to measurable fidelity. Synthetic examples illustrate recovery under controlled geometry.

## 2 RELATED WORK

**Neighborhood-based spatial representations.** BayesSpace uses a spatial prior for clustering (Zhao et al., 2021); SpaGCN aggregates expression and histology through graph convolution (Hu et al., 2021); STAGATE learns spatial attention within an autoencoder (Dong & Zhang, 2022); and GraphST combines graph encoding with local-context contrastive learning (Long et al., 2023). SpatialPCA places a spatial covariance kernel on latent factors (Shang & Zhou, 2022). These methods support denoising and domain identification through spatially informed representations. HODGEDEPTH regularizes the normal direction of a latent tissue foliation, specifying which molecular variation the global coordinate should retain.

**Continuous spatial coordinates.** Pseudotime organizes expression-defined state trajectories (Trapnell et al., 2014), while SpaceFlow and stLearn incorporate spatial information into related representations (Ren et al., 2022; Pham et al., 2023). Spatial depth targets the arrangement of molecular variation in physical tissue. GASTON models a tissue-wide isodepth and piecewise linear expression along it (Chitra et al., 2025a); GASTON-Mix learns domain-specific coordinates through a spatial mixture of experts (Chitra et al., 2025b). These approaches share our single-index motivation and the interpretability of an explicit depth axis. Their scalar neural fields are integrable by construction. HODGEDEPTH estimates a local direction field and fits its global potential through a geometric objective, exposing directional reliability, orientation consistency, and integration residuals. SLOPER learns gene-specific spatial score fields under an inhomogeneous Poisson point-process model and derives isodepth from gradient outer products (Wang et al., 2025). This provides a generative approach to estimating molecular gradients. Our contribution centers on assembling dominant projective directions through reliability-weighted fusion, synchronization, and global integration.

**Synchronization and Hodge foundations.** Vector diffusion maps use parallel transport and a connection Laplacian to organize vector-valued data (Singer & Wu, 2012). Discrete Hodge theory

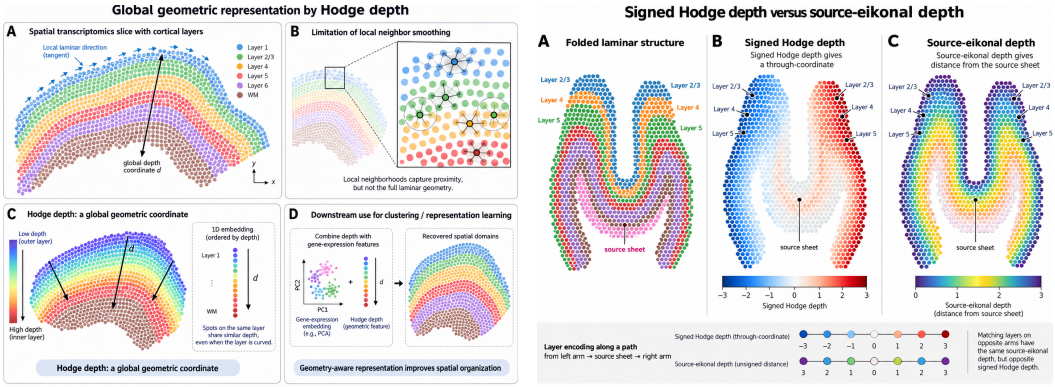


Figure 1: *Left:* HODGEDEPTH reconstructs global layered geometry. *Right:* Signed Hodge depth orders tissue positions; source-eikonal depth measures unsigned graph distance from a candidate internal sheet in recurrent layers.

separates edge flows into gradient and cyclic components (Jiang et al., 2011). HODGEDEPTH combines these foundations in an estimator that respects the sign ambiguity of molecular axes and separates directional evidence from depth parameterization. Projector fusion shares local evidence without choosing signs, synchronization establishes coherent orientations, and Hodge projection fits globally compatible increments. Our analysis quantifies geometric bias, propagates local errors to depth ordering, and establishes canonical mass-coordinate recovery. Appendix F discusses broader connections and why spatial smoothness alone does not identify depth.

## 3 LEARNING INTERPRETABLE GLOBAL SPATIAL DEPTH

HODGEDEPTH estimates local molecular axes, fuses their projectors, synchronizes orientations, and integrates the resulting field into *Hodge Depth* (Figure 1). A complementary *Source–Eikonal* coordinate measures distance from an internal sheet in recurrent layers. Appendix B gives implementation details.

### 3.1 SPATIAL GRADIENT AND THE DEPTH FOLIATION

Given locations $s_i \in \Omega \subset \mathbb{R}^2$ and molecular features $z_i \in \mathbb{R}^p$, we estimate spatial Jacobians by differentiating fitted multivariate fields:

$$f_{\theta_i} : \Omega \to \mathbb{R}^p, \quad \widehat{J}_i = Df_{\widehat{\theta}_i}(s_i). \tag{1}$$

Parameters may be shared across locations or fitted locally. Feature reconstruction and derivative regularization use no anatomical labels or scalar depth bottleneck.

Under the single-index model

$$z(s) = m(d_0(s)) + \varepsilon(s), \tag{2}$$

the unknown feature profile $m$ absorbs monotone relabelings of latent depth $d_0$, leaving its ordering and level sets as the target. The noiseless Jacobian is

$$J_z(s) = m'(d_0(s))\nabla d_0(s)^\top. \tag{3}$$

When both factors are nonzero, it has rank one: its leading right-singular vector identifies the contour-normal axis, represented without sign ambiguity by

$$P_0(s) = \frac{\nabla d_0(s)\nabla d_0(s)^\top}{\|\nabla d_0(s)\|^2}. \tag{4}$$

The regular level sets $\mathcal{L}_t = \{s : d_0(s) = t\}$ form a tissue foliation; the Jacobian nullspace gives their tangent lines.

All reported experiments use local affine fields $f_{\theta_i}(s) = z_i + A_i(s - s_i)$ fitted independently by weighted ridge regression, with fixed intercept $z_i$ and spatially varying Jacobians $\widehat{J}_i = \widehat{A}_i$. Shared nonlinear neural fields provide a generalization not evaluated here. Appendix B.7 gives the common objective and its affine specialization.

### 3.2 PROJECTIVE LINE FIELDS AND EIGEN-CONTRAST

A normal axis lies in $\mathbb{RP}^1$, which identifies a unit vector with its negative. Let $\tilde{n}_i$ be the leading right-singular vector of $\widehat{J}_i$, with singular values $\widehat{\sigma}_{i1} \geq \widehat{\sigma}_{i2}$. The *eigen-contrast*

$$\chi_i = \frac{\widehat{\sigma}_{i1}^2 - \widehat{\sigma}_{i2}^2}{\widehat{\sigma}_{i1}^2 + \widehat{\sigma}_{i2}^2}, \tag{5}$$

when defined, measures concentration along one spatial direction. It quantifies anisotropy, not absolute signal strength: high contrast alone cannot exclude anisotropic noise, and neighborhoods without feature change do not identify an axis.

On a connected spatial graph $G = (V, E)$, confidence $c_i$ combines eigen-contrast with spatial support, downweighting one-sided neighborhoods. We fuse sign-invariant projectors $\widetilde{P}_i = \tilde{n}_i \tilde{n}_i^\top$ through

$$\min_{Q_i = Q_i^\top} \sum_i c_i \left\| Q_i - \widetilde{P}_i \right\|_F^2 + \lambda_P \sum_{(i,j)\in E} w_{ij} \left\| Q_i - Q_j \right\|_F^2, \tag{6}$$

where $w_{ij}$ weights neighboring locations. The leading unit eigenvector $u_i$ of each fitted $Q_i$ defines a refined axis. Fusion therefore regularizes directions before constructing depth, without cancellation from arbitrary signs. Stabilization and support-harmonic refinement are specified in Appendix B.

### 3.3 ORIENTATION SYNCHRONIZATION AND GLOBAL HODGE INTEGRATION

To orient the refined axes coherently, set $g_{ij} = \mathrm{sign}(u_i^\top u_j)$ and use positive weights $a_{ij}$ combining proximity, confidence, and axis agreement. A spectral relaxation of

$$\max_{\sigma_i \in \{-1,+1\}} \sum_{(i,j)\in E} a_{ij} g_{ij} \sigma_i \sigma_j. \tag{7}$$

gives oriented normals $n_i = \sigma_i u_i$, up to an overall sign (Appendix B.4). Their proposed edge increments are

$$\omega_{ij} = \tfrac{1}{2}(n_i + n_j)^\top (s_j - s_i), \;\; \omega_{ji} = -\omega_{ij}. \tag{8}$$

Unit normals set spacing through spatial displacement: stronger molecular change does not automatically create larger depth gaps. Confidence controls the increments' influence. The population analysis in Section 4 establishes exact contour preservation under transnormal geometry and quantifies bias beyond that setting.

Synchronization resolves orientation but does not ensure integrability. Noisy increments can disagree around cycles; Hodge Depth fits a globally compatible scalar potential:

$$\widehat{d} = \operatorname*{arg\,min}_{\mathbf{1}^\top d = 0} \left\| W^{1/2}(Bd - \omega) \right\|_2^2. \tag{9}$$

Here $B$ is the oriented graph-incidence matrix, $(Bd)_{ij} = d_j - d_i$, and $W = \mathrm{diag}(a_{ij})$. The Laplacian $L = B^\top W B$ yields a sparse solve, with centering fixing the additive ambiguity. Fitted increments $B\widehat{d}$ sum to zero around every cycle; the residual lies in $\ker(B^\top W)$. Its normalized magnitude $R_H$ measures incompatibility with a scalar potential, not biological accuracy. We normalize $\widehat{d}$ to $\phi \in [0, 1]$; Appendix B.5 specifies the solve, residual, and normalization.

### 3.4 SOURCE–EIKONAL DEPTH AND DOWNSTREAM USE

When layers recur on both sides of an internal sheet, unsigned distance can give corresponding layers comparable coordinates. We locate a candidate sheet by comparing expression profiles at matched offsets on opposite sides of a Hodge-depth level. For a narrow source band $\mathcal{S}$, define

$$q_i = \min_{v \in \mathcal{S}} \mathrm{dist}_G(v, i), \;\; \psi_i = \frac{q_i}{\max_j q_j}. \tag{10}$$

Graph distance uses physical edge lengths. A label-free gate selects $\psi$ only when molecular reflection is supported across feature scales and it improves both piecewise expression reconstruction and cross-validated expression prediction; otherwise it retains $\phi$. The gate selects between existing coordinates without refitting the normal field or Hodge potential (Appendix B.6).

**Output and downstream refinement.** The output comprises depth, directional confidence, and orientation and integration diagnostics. Depth supports gene-expression profiling and can complement expression features in downstream representations or ordered domain decoding. Local expression and connectivity can refine assignments without redefining depth; cohort-specific appendices give the decoder configurations.

## 4 From local molecular directions to global geometric depth

Under the regular single-index model $z = m \circ d_0$, nonvanishing derivatives and connected level sets identify the depth foliation up to monotone relabeling (Theorem 4). We establish three recovery results linking local directions to geometry, ordering, and a canonical scale. Supporting statements and proofs are in Appendix A.

### 4.1 Geometric bias and contour recovery

Let $r = \|\nabla d_0\|$ and $n = \nabla d_0 / r$. The continuum Hodge potential $u_H$ minimizes $\mathcal{J}(u) = \int_\Omega w \|\nabla u - n\|^2$ over weighted-mean-zero $H^1$ functions. A competing coordinate $u = h(d_0)$ preserves the true contours. On $\Sigma_t = d_0^{-1}(t)$, define $A_j(t) = \int_{\Sigma_t} w r^{j-1} \, d\mathcal{H}^1$ and $d\mu_t = w \, d\mathcal{H}^1/(A_0 r)$. Here $w$ is a fixed positive weight and $\|f\|_w^2 = \int_\Omega w \|f\|^2$.

**Theorem 1** (Optimal contour-preserving depth and Hodge bias)**.** *Suppose $\Omega$ is bounded, connected, and Lipschitz, $d_0$ extends smoothly to a neighborhood of its closure with connected regular level sets, $0 < r_- \le r \le r_+$, and $w$ is bounded above and away from zero (Appendix A.12). The optimal contour-preserving coordinate is $u_F = h_F(d_0)$, where $1/r_+ \le h_F' = A_1/A_2 \le 1/r_-$. With $I = d_0(\Omega)$, its residual and its discrepancy from Hodge integration satisfy*

$$\mathcal{D}_F := \mathcal{J}(u_F) = \int_I \left( A_0 - \frac{A_1^2}{A_2} \right) dt = \int_I A_0 \frac{\mathrm{Var}_{\mu_t}(r)}{\mathbb{E}_{\mu_t}[r^2]} \, dt, \tag{11}$$

$$\|\nabla(u_H - u_F)\|_w^2 = \mathcal{D}_F - \mathcal{J}(u_H). \tag{12}$$

*The defect vanishes exactly when $r = q(d_0)$ for a positive function $q$, the transnormal condition.*

Within-contour variation of gradient magnitude obstructs exact unit-normal integration. A variable-width example (Appendix A.12) shows contour distortion even with perfect local axes. For directional error $\|\widehat{n} - n\|_w \le \delta$, Corollary 10 bounds the gradient error by $\delta + \sqrt{\mathcal{D}_F - \mathcal{J}(u_H)}$; a leafwise Poincaré inequality controls average within-contour dispersion. The appendix separately bounds finite-graph discretization error.

### 4.2 From local Jacobians to global depth ordering

On a connected spatial graph, orient $e = (i, j)$ and set $x_e = s_j - s_i$, $\ell_e = \|x_e\|$, $(Bd)_e = d_j - d_i$, and $(\mathsf{A}v)_e = (v_i + v_j)^\top x_e / 2$. For positive integration weights $W$, write $L = B^\top W B$. For a vertex pair, let $h_{uv} = e_v - e_u$, $j^{uv} = WBL^\dagger h_{uv}$, and $\zeta^{uv} = \mathsf{A}^\top j^{uv}$. The block $\zeta_i^{uv}$ measures how node $i$'s normal affects that pair's recovered depth difference.

The projector smoother is $H = (C_P + \lambda_P L_g + \epsilon_P I)^{-1} C_P$, where $L_g$ is the spatial Laplacian, $C_P = \mathrm{diag}(c_i^P)$, $c_i^P \ge 0$, $\sum_i c_i^P > 0$, and $\lambda_P, \epsilon_P > 0$. Its positive row sums $q_i = \sum_j H_{ij}$ normalize the fused projectors. Appendix A.13 specifies the spectral synchronization and Hodge estimators.

**Theorem 2** (Depth-order recovery from local Jacobian bounds)**.** *Let $d_\star$ be $C^3$ on a neighborhood of all edge segments, with $n_\star = \nabla d_\star$, $\|n_\star\| = 1$, and third-derivative norm bounded by $M$. Write $n_i^\star = n_\star(s_i)$ and $P_i^\star = n_i^\star (n_i^\star)^\top$. At positive-fidelity nodes, suppose $J_i^\star = b_i (n_i^\star)^\top$, $\kappa_i = \|b_i\| > 0$, and $\left\| \widehat{J}_i - J_i^\star \right\|_2 \le \eta_i \le \kappa_i / 2$. Set $p_i = 2\eta_i/\kappa_i$ there and $p_i = 1$ elsewhere. Define*

$$t_i = \frac{\sum_j H_{ij} (p_j + \left\| P_j^\star - P_i^\star \right\|_2)}{q_i} < \frac{1}{2}, \qquad \varepsilon_i = \arcsin \frac{t_i}{1 - t_i}, \quad \delta_i = 2 \sin(\varepsilon_i / 2).$$

*If* $\angle(n_i^\star, n_j^\star) + \varepsilon_i + \varepsilon_j < \pi/2$ *on every synchronization edge, there is one* $\sigma \in \{-1, 1\}$ *such that* $\|\sigma\widehat{n}_i - n_i^\star\| \le \delta_i$ *for every node. The centered Hodge estimator satisfies, simultaneously for every distinct pair,*

$$\left|\sigma(\widehat{d}_v - \widehat{d}_u) - [d_\star(s_v) - d_\star(s_u)]\right| \le \underbrace{\sum_i \|\zeta_i^{uv}\|\,\delta_i + \frac{M}{12}\sum_e |j_e^{uv}|\ell_e^3}_{E_{uv}}. \tag{13}$$

*Every pair with true depth gap greater than* $E_{uv}$ *is correctly ordered, up to this common reversal.*

The radii $\delta_i$ propagate local Jacobian error and the bias from fusing differently oriented projectors. The bound retains their shared effects across incident edges and separates them from quadrature error. It holds for realized data-dependent weights. Choosing $d_\star$ as the transnormal representative gives the ordering of $d_0$. The assumptions require derivative accuracy and geometric margins; they do not establish neural-training convergence or correctness of source inference.

### 4.3 Recovery of a canonical mass coordinate

Fix an orientation and spatial probability measure $\mu$, and write $F_d(t) = \mu\{s : d(s) \le t\}$. The cumulative-mass coordinate $r_0 = F_{d_0} \circ d_0$ is invariant to increasing depth relabeling (Proposition 11). For a fixed oracle $d_\star$, let $x_i = d_\star(s_i)$, $y_i = \widehat{d}_i$, and $\|v\|_{2,N}^2 = N^{-1}\sum_i v_i^2$. Define the centered depth error $E_N = \inf_c \|y + c\mathbf{1} - x\|_{2,N}$ and the oracle sampling discrepancy $\Delta_N = \sup_t |N^{-1}\sum_i \mathbf{1}\{x_i \le t\} - F_{d_\star}(t)|$.

**Theorem 3** (Canonical mass-coordinate recovery)**.** *If* $F_{d_\star}$ *is* $M$*-Lipschitz, let* $R_i$ *be the rank of* $y_i$ *among the* $N$ *fitted depths, using ranks* $1, \dots, N$ *and average ranks for ties. Then*

$$\|(R_i/N)_i - (F_{d_\star}(x_i))_i\|_{2,N} \le 2ME_N + \Delta_N. \tag{14}$$

*The bound is deterministic and allows ties and fitting and ranking on the same observations.*

Optimal one-dimensional matching converts depth error into spatially assigned rank error. Appendix A.15 connects $E_N$ to Hodge error and graph connectivity, and controls $\Delta_N$ under independent sampling. When $d_\star$ is an increasing relabeling of $d_0$, the target is $r_0$. Mass calibration fixes a coordinate convention under $\mu$; it is optional and is not applied in the reported experiments.

**Supporting results and extension.** The appendix establishes foliation identifiability (Theorem 4), approximate geometry recovery (Corollary 10), projector perturbation (Theorem 5), exact orientation (Proposition 6), transnormal integration (Corollary 7), Hodge stability and direction-to-depth error (Theorem 8 and Corollary 9), canonical-coordinate uniqueness and population rank stability (Proposition 11 and Lemma 12), and joint spatial-depth transport recovery (Corollary 13). Finally, Proposition 14 establishes exact population recovery with a learned positive integration speed and prescribed mass law, without transnormality. This extension has no noisy-field stability or computational guarantee yet.

## 5 Numerical studies

Our main evidence comes from two human cortical cohorts. The 12-section DLPFC study evaluates both downstream domain recovery and the continuous depth coordinate; an independent 11-section Batiuk cohort assesses downstream utility in a second dataset. Synthetic examples illustrate recovery under controlled geometry (Appendix D).

### 5.1 Main study: 12 human DLPFC sections

We use the Visium sections 151507–151510 and 151669–151676 from Maynard et al. (2021). Comparisons include GraphST (Long et al., 2023), SEDR (Xu et al., 2024), STAGATE (Dong & Zhang, 2022), BASS (Li & Zhou, 2022), SpatialPCA (Shang & Zhou, 2022), BayesSpace (Zhao et al., 2021), SpaGCN (Hu et al., 2021), GASTON (Chitra et al., 2025a), and SpaceFlow (Ren et al., 2022). HodgeDepth fitting and configuration selection use no reference-label ARI. A downstream decoder

Table 1: Main DLPFC benchmark: equal-weight means over 12 sections. Both metric groups assess the same partitions. Bold marks the best mean, including displayed ties. Full results: Appendix C.

| | Annotation agreement ↑ | | | Label-free coherence ↓ | | |
|---|---|---|---|---|---|---|
| Method | ARI | NMI | Acc. | CHAOS | PAS | LISI |
| HodgeDepth | **0.618** | **0.695** | **0.737** | **0.0582** | **0.0110** | **1.225** |
| GraphST | 0.523 | 0.657 | 0.678 | **0.0582** | 0.0125 | **1.225** |
| SEDR | 0.533 | 0.663 | 0.683 | 0.0583 | 0.0197 | 1.274 |
| STAGATE (PyG) | 0.509 | 0.642 | 0.679 | 0.0589 | 0.0337 | 1.298 |
| BASS | 0.515 | 0.650 | 0.671 | 0.0594 | 0.0381 | 1.310 |
| SpatialPCA | 0.503 | 0.640 | 0.664 | 0.0583 | 0.0205 | 1.263 |
| BayesSpace | 0.469 | 0.608 | 0.617 | 0.0589 | 0.0334 | 1.259 |
| SpaGCN (no HE) | 0.434 | 0.567 | 0.603 | 0.0595 | 0.0718 | 1.348 |
| GASTON | 0.420 | 0.592 | 0.575 | 0.0584 | 0.0172 | 1.236 |
| SpaceFlow | 0.305 | 0.458 | 0.463 | 0.0588 | 0.0266 | 1.262 |

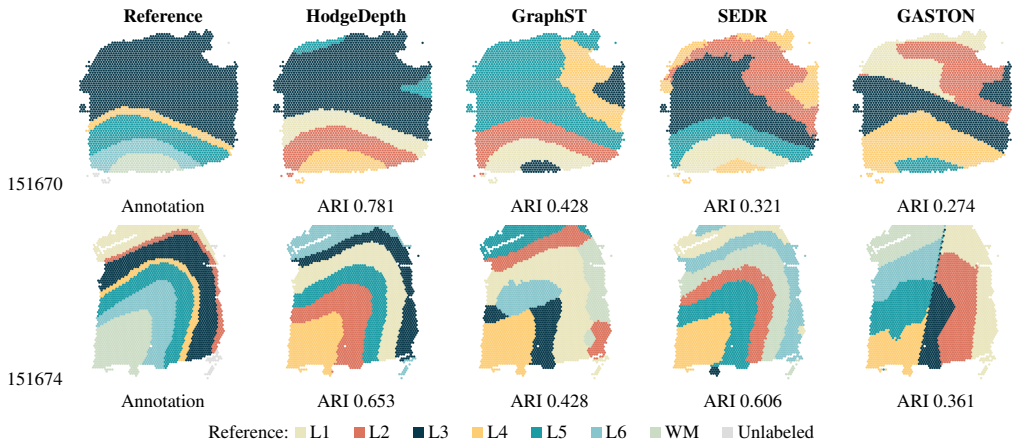


Figure 2: Domain recovery in 151670 (Br5595) and 151674 (Br8100). Predicted cluster colors need not match anatomical colors. All 12 sections: Appendix C.

combines depth and expression, followed by spatial label voting; these steps do not update the depth. Software versions and configurations are in Appendix C. Cohort summaries weight sections equally. We do not include SLOPER as a separate depth-ordering comparison because its authors report that the SLOPER index is strongly correlated with GASTON isodepth on the same tissues (Wang et al., 2025), so GASTON serves as the representative of this family.

**(a) Downstream clustering: agreement and spatial coherence.** We assess annotation agreement using ARI, NMI, and one-to-one matched accuracy. Complementary label-free measures evaluate spatial coherence: CHAOS, abnormal-spot fraction (PAS), and local inverse Simpson index (LISI), with lower values preferred under the benchmark convention. These measures evaluate predicted partitions without reference layers; coherence alone does not establish anatomical correctness.

HodgeDepth attains mean ARI 0.618, versus 0.533 for SEDR and 0.523 for GraphST, and leads in NMI and accuracy (Table 1). It ranks first in six sections and exceeds GraphST and GASTON in 11 each. Its PAS is lowest (0.0110); CHAOS (0.0582) and LISI (1.225) tie GraphST at displayed precision. Thus anatomical agreement is accompanied by coherent spatial partitions. Figure 2 illustrates layer continuity and fragmentation in two sections from different donors; the quantitative comparison uses all 12 sections.

**(b) Continuous depth: anatomical ordering.** To evaluate the representation independently of a clustering decoder, we compare depth with the ordinal reference $\mathrm{L1} < \mathrm{L2} < \cdots < \mathrm{L6} < \mathrm{WM}$ on annotated spots. Each coordinate receives one global sign alignment per section for evaluation only. Table 2 reports rank agreement, adjacent-layer separation, and isotonic fit.

Table 2: Depth ordering across 12 DLPFC sections: section means, except the pooled ordered-pair count; higher is better. Sign alignment uses reference layers for evaluation only. Definitions: Appendix C.2.

| Coordinate | Spearman $\lvert\rho\rvert$ | Kendall $\lvert\tau_b\rvert$ | Adjacent-pair AUC | | Pairs ordered | Monotone $R^2$ |
|---|---|---|---|---|---|---|
| | | | mean | min | | |
| HodgeDepth | **0.885** | **0.784** | **0.901** | **0.757** | **61/64** | **0.883** |
| GASTON isodepth | 0.690 | 0.605 | 0.834 | 0.661 | 60/64 | 0.785 |
| Spatial-graph Fiedler vector | 0.661 | 0.521 | 0.668 | 0.561 | **61/64** | 0.530 |

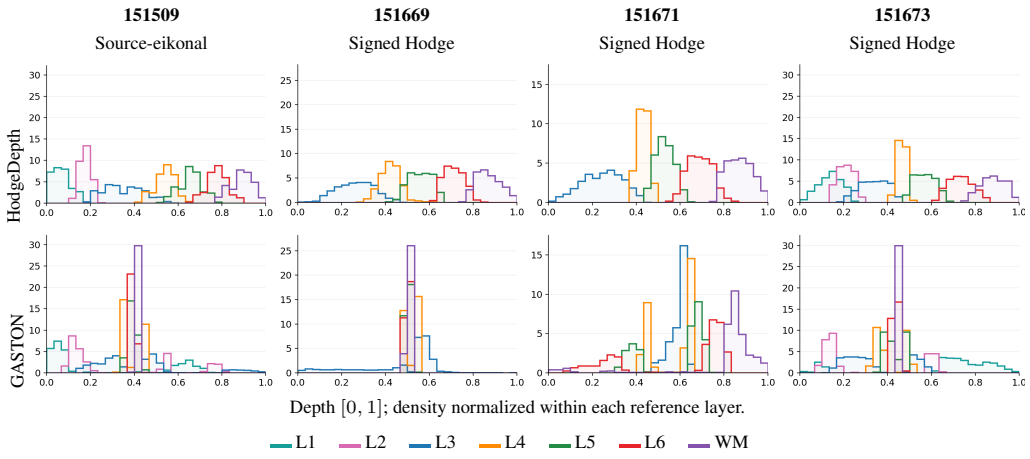


Figure 3: Layer-conditioned depth distributions: gate-selected HodgeDepth (top; source-eikonal for 151509, signed otherwise) and GASTON (bottom). Densities use 30 bins and section-specific scales. Reference maps, unselected candidates, and display details: Appendix B.9.

HodgeDepth has the highest mean absolute Spearman correlation (0.885), Kendall correlation (0.784), and monotone $R^2$ (0.883). Its mean adjacent-pair AUC is 0.901, compared with 0.834 for GASTON and 0.668 for the spatial-graph Fiedler vector. Both HodgeDepth and the Fiedler vector order 61/64 adjacent pairs correctly, but their AUCs distinguish the extent of layer overlap. The four-section distributions in Figure 3 illustrate this separation using the gate-selected coordinates. These results support anatomical ordering, without establishing physical depth calibration or cross-section correspondence.

Focused ablations on the 12 DLPFC sections support HodgeDepth's advantage over the tested geometric coordinates, the robustness of global sign synchronization, and a modest benefit from unit-direction increments for mean ordering fidelity (Appendix C.5).

### 5.2 Independent Cortical Cohort: Batiuk

We assess downstream domain recovery on 11 Batiuk DLPFC sections (five control, six schizophrenia) (Batiuk et al., 2022). HodgeDepth, GraphST, and GASTON are scored on common annotated spots with requested $K = 7$; summaries weight sections equally. HodgeDepth achieves mean ARI 0.5663, versus 0.4651 and 0.4499, and ranks first in nine sections. GraphST leads on MB14 and GASTON on MB50. This second cohort supports downstream utility beyond the primary benchmark; Table 20 and Appendix E give aggregate and section-level results.

**Synthetic illustration.** Controlled layered examples illustrate how the geometry translates into recovered depth. Across six primary settings, HodgeDepth's mean affine-calibrated depth NRMSE is 2.01%, versus 7.31% for GASTON; shared-decoder ARI is 0.915 versus 0.621. These illustrative single-realization results do not replace the real-tissue benchmarks. Appendix D contains the protocol, complete results, recovery maps, controls, and documented failure.

## 6 Discussion and conclusion

HodgeDepth makes global coordinate fidelity an explicit objective of spatial representation learning. Accurate local directions alone need not yield a faithful global depth: their orientation, integration, and parameterization jointly determine the recovered geometry. The learned coordinate provides an inspectable spatial axis for relating molecular profiles to tissue organization. Directional confidence and integration residuals expose agreement with a scalar geometric description, while the ordering evaluation assesses the coordinate separately from downstream clustering. This representation supports scientific exploration of spatial expression profiles and candidate markers.

The model is most appropriate when one dominant direction organizes local molecular variation. Independent programs, branching, and sharp transitions can challenge this representation or its derivative estimates. Unit-normal integration preserves the latent foliation under transnormal geometry; beyond this setting, contour distortion can persist even with perfect axes. The theory separates this geometric bias from estimation error and gives sufficient conditions for global ordering. These guarantees concern the stated models and do not certify anatomical validity. Experiments evaluate the local affine estimator; shared neural fields remain an unevaluated extension.

Learning integration speed jointly with depth is a natural next step. Under exact normals and a prescribed mass law, the population extension identifies a unique contour-preserving coordinate without transnormality. Stability under estimated normals and a practical optimization method are needed to turn this result into an estimator. Further evaluation should examine how sampling and geometric misspecification affect this construction. HodgeDepth provides an explicit framework for connecting local molecular evidence to global tissue organization and for assessing the resulting coordinate alongside its downstream uses.

## Broader impact

An interpretable tissue coordinate can support visualization and downstream analysis of spatial molecular organization. A one-dimensional representation can also suppress branching, local niches, or orthogonal biological programs. Reliability, synchronization, and Hodge-residual maps should therefore accompany the scalar depth to support assessment of the representation's suitability. These diagnostics describe compatibility with the model and do not certify anatomical correctness. Layer annotations and tissue provenance can reflect human or sampling biases; agreement with such annotations provides evidence of anatomical alignment within the evaluated data.

### AI use statement

Generative AI tools were used to assist with implementing the proposed method, debugging experimental code, and generating the synthetic datasets described in Appendix D. They also assisted with drafting and editing portions of the manuscript and creating Figure 1. All AI-assisted code was reviewed and tested by the authors, and all AI-assisted text and figures were reviewed and revised for consistency with the described methods and results. The authors take responsibility for the final content of this work, including all text, code, data, figures, claims, and experimental results produced with the aid of generative AI.

### Ethics statement

This work performs secondary computational analyses of previously published human postmortem spatial transcriptomics datasets from the dorsolateral prefrontal cortex (Maynard et al., 2021; Batiuk et al., 2022). We did not recruit participants or collect new human-subject data. The original studies report the relevant ethical approvals and informed-consent procedures, and our analyses use their released data and annotations in accordance with the associated access and usage conditions.

The learned one-dimensional anatomical coordinate is intended as a computational representation of layered spatial organization and should not be interpreted as biological ground truth or as a clinical measure. Such a representation may obscure genuine branching, local niches, or biological programs orthogonal to the inferred depth. Reference layer annotations, cohort composition, disease status, and tissue sampling may also introduce annotation- or cohort-specific biases. We therefore

interpret agreement with reference annotations as evidence of anatomical alignment rather than proof of biological validity, and recommend considering the reported reliability, synchronization, and Hodge-residual diagnostics together with the learned coordinate.

## REPRODUCIBILITY STATEMENT

Appendix B specifies depth estimation and the label-free source gate. Appendix C gives benchmark software, decoder configurations, ordering-metric definitions, and complete DLPFC results; Appendix E describes the 11-section Batiuk subset and common scored spots. Appendix D documents synthetic generation, the official GASTON implementation, training budgets, shared decoder, and evaluation calibration. Section and aggregate values retain their supplied precision. The supplementary code includes a replication kit covering the two main numerical studies and the synthetic-data analysis.

## REFERENCES

Prabir Barooah and Joao P Hespanha. Estimation from relative measurements: Electrical analogy and large graphs. *IEEE Transactions on Signal Processing*, 56(6):2181–2193, 2008.

Mykhailo Y Batiuk, Teadora Tyler, Katarina Dragicevic, Shenglin Mei, Rasmus Rydbirk, Viktor Petukhov, Ruslan Deviatiiarov, Dora Sedmak, Erzsebet Frank, Virginia Feher, et al. Upper cortical layer–driven network impairment in schizophrenia. *Science advances*, 8(41):eabn8367, 2022.

Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

Uthsav Chitra, Brian J Arnold, Hirak Sarkar, Kohei Sanno, Cong Ma, Sereno Lopez-Darwin, and Benjamin J Raphael. Mapping the topography of spatial gene expression with interpretable deep learning. *Nature Methods*, 22(2):298–309, 2025a.

Uthsav Chitra, Shu Dan, Fenna Krienen, and Benjamin J Raphael. Gaston-mix: a unified model of spatial gradients and domains using spatial mixture-of-experts. *Bioinformatics*, 41(Supplement_1): i523–i532, 2025b.

Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.

Emma Dann, Neil C Henderson, Sarah A Teichmann, Michael D Morgan, and John C Marioni. Differential abundance testing on single-cell data using k-nearest neighbor graphs. *Nature biotechnology*, 40(2):245–253, 2022.

Kangning Dong and Shihua Zhang. Deciphering spatial domains from spatially resolved transcriptomics with an adaptive graph attention auto-encoder. *Nature communications*, 13(1):1739, 2022.

Wendell H. Fleming and Raymond Rishel. An integral formula for total gradient variation. *Archiv der Mathematik*, 11:218–222, 1960. doi: 10.1007/BF01236935.

Nicolás García Trillos and Dejan Slepčev. Continuum limit of total variation on point clouds. *Archive for rational mechanics and analysis*, 220(1):193–241, 2016.

Trevor Hastie and Werner Stuetzle. Principal curves. *Journal of the American statistical association*, 84(406):502–516, 1989.

Jian Hu, Xiangjie Li, Kyle Coleman, Amelia Schroeder, Nan Ma, David J Irwin, Edward B Lee, Russell T Shinohara, and Mingyao Li. Spagcn: Integrating gene expression, spatial location and histology to identify spatial domains and spatially variable genes by graph convolutional network. *Nature methods*, 18(11):1342–1351, 2021.

Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.

Yinuo Jin, Joshua D Myers, Presha Rajbhandari, Jia Yi Zhang, Kaylee W Fang, John Shaw Moazami, Noreen Hosny, Brent R Stockwell, and Elham Azizi. Lynx: a deep generative model for linking spatial dynamics and cell interactions in multimodal spatial data. *bioRxiv*, pp. 2026–07, 2026.

Zheng Li and Xiang Zhou. Bass: multi-scale and multi-sample analysis enables accurate cell type clustering and spatial domain detection in spatial transcriptomic studies. *Genome biology*, 23(1): 168, 2022.

Yahui Long, Kok Siong Ang, Mengwei Li, Kian Long Kelvin Chong, Raman Sethi, Chengwei Zhong, Hang Xu, Zhiwei Ong, Karishma Sachaphibulkij, Ao Chen, et al. Spatially informed clustering, integration, and deconvolution of spatial transcriptomics with graphst. *Nature communications*, 14 (1):1155, 2023.

Yue Lyu, Chong Wu, Wei Sun, and Ziyi Li. Regional analysis to delineate intrasample heterogeneity with regionalst. *Bioinformatics*, 40(4):btae186, 2024.

Cong Ma, Uthsav Chitra, Shirley Zhang, and Benjamin J Raphael. Belayer: Modeling discrete and continuous spatial variation in gene expression from spatially resolved transcriptomics. *Cell systems*, 13(10):786–797, 2022.

Pascal Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, pp. 1269–1283, 1990.

Kristen R Maynard, Leonardo Collado-Torres, Lukas M Weber, Cedric Uytingco, Brianna K Barry, Stephen R Williams, Joseph L Catallini, Matthew N Tran, Zachary Besich, Madhavi Tippani, et al. Transcriptome-scale spatial gene expression in the human dorsolateral prefrontal cortex. *Nature neuroscience*, 24(3):425–436, 2021.

Duy Pham, Xiao Tan, Brad Balderson, Jun Xu, Laura F Grice, Sohye Yoon, Emily F Willis, Minh Tran, Pui Yeng Lam, Arti Raghubar, et al. Robust mapping of spatiotemporal trajectories and cell–cell interactions in healthy and diseased tissues. *Nature communications*, 14(1):7739, 2023.

Honglei Ren, Benjamin L Walker, Zixuan Cang, and Qing Nie. Identifying multicellular spatiotemporal organization of cells with spaceflow. *Nature communications*, 13(1):4076, 2022.

Filippo Santambrogio. $l^1$ and $l^\infty$ theory. In *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*, pp. 87–119. Springer, 2015.

Lulu Shang and Xiang Zhou. Spatially aware dimension reduction for spatial transcriptomics. *Nature communications*, 13(1):7203, 2022.

Amit Singer and H-T Wu. Vector diffusion maps and the connection laplacian. *Communications on pure and applied mathematics*, 65(8):1067–1144, 2012.

Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

Cole Trapnell, Davide Cacchiarelli, Jonna Grimsby, Prapti Pokharel, Shuqiang Li, Michael Morse, Niall J Lennon, Kenneth J Livak, Tarjei S Mikkelsen, and John L Rinn. The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nature biotechnology*, 32(4):381–386, 2014.

An Wang, Donald Geman, Uthsav Chitra, and Laurent Younes. Mapping spatial gradients in spatial transcriptomics data with score matching. *bioRxiv*, pp. 2025–11, 2025.

Qi-Ming Wang. Isoparametric functions on riemannian manifolds. i. *Mathematische Annalen*, 277 (4):639–646, 1987.

Hang Xu, Huazhu Fu, Yahui Long, Kok Siong Ang, Raman Sethi, Kelvin Chong, Mengwei Li, Rom Uddamvathanak, Hong Kai Lee, Jingjing Ling, et al. Unsupervised spatially embedded deep representation of spatial transcriptomics. *Genome medicine*, 16(1):12, 2024.

Edward Zhao, Matthew R Stone, Xing Ren, Jamie Guenthoer, Kimberly S Smythe, Thomas Pulliam, Stephen R Williams, Cedric R Uytingco, Sarah EB Taylor, Paul Nghiem, et al. Spatial transcriptomics at subspot resolution with bayesspace. *Nature biotechnology*, 39(11):1375–1384, 2021.

# A PROOFS AND SUPPORTING RESULTS

This appendix gives all supporting statements and the proofs of the three main theorems. It first identifies the depth foliation and records local perturbation, orientation, transnormality, and Hodge stability. The subsequent sections develop geometric-bias bounds, the detailed hypotheses and proof of ordering recovery, canonical mass coordinates and transport bounds, and the population extension with learned integration speed.

## A.1 FOLIATIONS AND POPULATION IDENTIFIABILITY

**Assumption 1** (Regular single-index field)**.** $\Omega \subset \mathbb{R}^2$ is open and connected; $d_0 \in C^2(\Omega)$ has no critical points; $m \in C^1(\mathbb{R}; \mathbb{R}^p)$ satisfies $m'(d_0(s)) \neq 0$; and each level set of $d_0$ is connected.

**Theorem 4** (Identifiable depth foliation)**.** *Under Assumption 1, the population Jacobian of $z = m \circ d_0$ has rank one and uniquely determines the projector $P_0$ in Eq. (4). If $z = \widetilde{m} \circ \widetilde{d}$ is another regular single-index representation, then $\widetilde{d} = g \circ d_0$ for a continuously differentiable strictly monotone function $g$.*

The theorem identifies the family of connected layer contours and their ordering up to reversal. A monotone reparameterization relabels the leaves of this foliation while preserving their geometry, so expression alone does not identify an absolute metric scale. The projector combines the normal information shared across features; an individual gene can be locally uninformative even when the multivariate derivative is nonzero.

## A.2 STABILITY OF PROJECTIVE LINE ESTIMATION

At a fixed location write $J_0 = an^\top$, where $\|n\| = 1$ and $\kappa = \|a\| = \sigma_1(J_0) > 0$.

**Theorem 5** (Projector perturbation)**.** *Let $\widehat{J} = J_0 + E$ with $\|E\|_2 \leq \kappa/2$, and let $\widehat{P}$ be the leading right-singular projector of $\widehat{J}$. Then*

$$\left\| \widehat{P} - nn^\top \right\|_F \leq \frac{2\sqrt{2}\,\|E\|_2}{\kappa}. \tag{15}$$

The bound is a direct rank-one specialization of singular-subspace perturbation. It shows that directional stability depends on derivative-estimation error relative to transcriptomic signal strength $\kappa$. The error $E$ may arise from neural differentiation or local affine estimation; the result does not require a particular parameterization. It is conditional on derivative accuracy, which feature reconstruction alone does not establish. Eigen-contrast measures directional anisotropy; because it is scale-invariant, it does not by itself certify absolute signal strength or derivative accuracy.

## A.3 ORIENTATION SYNCHRONIZATION AND TOPOLOGY

A continuous projector field is a map into $\mathbb{RP}^1$. On a simply connected tissue region it admits a continuous lift to a unit vector field. A general line field on a multiply connected region can have an orientation obstruction, although a hole alone does not imply one. Under Assumption 1, the true field is already orientable through $\nabla d_0 / \|\nabla d_0\|$; synchronization estimates this lift from independently signed local axes. Its graph counterpart is consistency of the relative signs around cycles. The proposition uses the unnormalized connection Laplacian; normalization changes the magnitudes of the noiseless null vector but preserves the signs recovered by rounding.

**Proposition 6** (Exact graph orientation)**.** *Suppose $G$ is connected and there are unit vectors $n_i$ with $n_i^\top n_j > 0$ for every $(i,j) \in E$. If the observed eigenvectors satisfy $u_i = \tau_i n_i$ for unknown $\tau_i \in \{\pm 1\}$, then $g_{ij} = \tau_i \tau_j$. The unnormalized connection Laplacian has a one-dimensional nullspace spanned by $\tau$. For its normalized version the nullvector is $D^{1/2}\tau$, where $D$ contains the positive weighted degrees. Both yield the same rounded signs, so spectral synchronization recovers all orientations up to one global sign.*

With noisy connections, eigenspace perturbation scales inversely with the relevant connection-Laplacian spectral gap. Exact recovery of rounded signs additionally requires control of individual

eigenvector entries. A small gap or inconsistent cycle signs therefore diagnoses unstable orientation, but does not by itself identify an anatomical defect. Orientability is distinct from integrability: a consistently oriented field may still have nonzero circulation.

### A.4 TRANSNORMAL GEOMETRY AND EXACT INTEGRATION

Because $d_0$ is identifiable only up to monotone transformation, HODGEDEPTH integrates a unit normal field to seek an arc-length-like depth coordinate. Whether this normalization preserves the original depth contours depends on the integrability of the unit normal field. On a region $\Omega$ where $d_0$ is twice continuously differentiable and $\|\nabla d_0\| > 0$, write $r = \|\nabla d_0\|$, $n = \nabla d_0/r$, and $t = (-n_2, n_1)$ for the unit tangent to a layer. Then,

$$\operatorname{curl} n = \partial_1 n_2 - \partial_2 n_1 = t^\top \nabla \log r. \tag{16}$$

Thus variation of gradient magnitude along a layer obstructs exact integration of its unit normals. A sufficient condition for exact integration is the *transnormal condition* (Wang, 1987). Let $I$ be an open interval containing $d_0(\Omega)$, and let $q : I \to (0, \infty)$ be continuous. Suppose

$$\|\nabla d_0(s)\| = q(d_0(s)), \quad s \in \Omega. \tag{17}$$

This condition makes the gradient magnitude depend only on depth and hence constant along each level set. The following corollary gives an explicit unit-gradient representative.

**Corollary 7** (Transnormal depth)**.** *Under Assumption 1 and the transnormal condition (17), fix $t_0 \in I$ and define*

$$d_\star(s) = \int_{t_0}^{d_0(s)} \frac{1}{q(u)}\, du. \tag{18}$$

*Then $d_\star$ is a strictly increasing reparameterization of $d_0$ and $\nabla d_\star(s) = \nabla d_0(s)/\|\nabla d_0(s)\|$. Hence, the unit normal field oriented in the direction of increasing $d_0$ is globally integrable, with potential $d_\star$.*

The coordinate $d_\star$ has unit gradient, preserves the original foliation, and increases at unit rate along its oriented normal trajectories. Signed distance to a smooth reference contour provides an example within a sufficiently small tubular neighborhood where it is smooth. The explicit potential above establishes global integrability without an additional topological assumption. For a general oriented field, local curl-free behavior must also be accompanied by zero circulation around noncontractible loops to yield a global potential. Outside the transnormal setting, Hodge projection supplies the closest exact edge field in the specified weighted norm, while need not preserve the original depth contours exactly.

### A.5 DISCRETE HODGE PROJECTION AND GLOBAL ERROR

For an edge vector $v$, let $\|v\|_W = \left\|W^{1/2} v\right\|_2$, and for a centered node vector $x$, let $\|x\|_L^2 = x^\top L x$. Let $\Pi_\nabla = B L^\dagger B^\top W$ be the weighted projection onto $\operatorname{im}(B)$, where $L^\dagger$ is the Moore–Penrose pseudoinverse.

**Theorem 8** (Hodge projection, nuisance rejection, and stability)**.** *Let $G$ be connected and suppose*

$$\omega = B d_\star + \xi + \eta, \quad B^\top W \eta = 0, \tag{19}$$

*where $\mathbf{1}^\top d_\star = 0$. Then the solution of Eq. (9) obeys*

$$\widehat{d} - d_\star = L^\dagger B^\top W \xi, \tag{20}$$

$$\left\|\widehat{d} - d_\star\right\|_L = \|\Pi_\nabla \xi\|_W \le \|\xi\|_W, \tag{21}$$

$$\left\|\widehat{d} - d_\star\right\|_2 \le \lambda_2(L)^{-1/2} \|\xi\|_W. \tag{22}$$

*The non-gradient component $\eta$ is removed exactly.*

The weighted orthogonal decomposition is $\mathbb{R}^{|E|} = \operatorname{im}(B) \oplus_W \ker(B^\top W)$. Only the gradient component of edge noise affects the recovered potential; the complementary component is visible

in the integration residual. Gradient-like errors can therefore bias depth even when the residual is small. The algebraic connectivity $\lambda_2(L)$ controls conversion from edge error to nodewise depth error, conditionally on the chosen graph and weights. These bounds concern the centered potential before its final range normalization.

**Corollary 9** (Direction-to-depth error)**.** *Assume a smooth integrable unit field $n_\star = \nabla d_\star$, and let $\ell_{ij} = \|s_j - s_i\|$. If the synchronized estimates satisfy $\|n_i - n_\star(s_i)\| \le \delta_i$, then the trapezoidal edge error satisfies*

$$|\omega_{ij} - (d_\star(s_j) - d_\star(s_i))| \le \tfrac{1}{2}\ell_{ij}(\delta_i + \delta_j) + C\ell_{ij}^3, \tag{23}$$

*when $d_\star$ has bounded third derivatives along edges. Combining Eqs. (22) and (23) yields an end-to-end depth bound controlled by local line estimation, orientation accuracy, graph resolution, and the graph spectral gap.*

**Interpretation.** The representation target is a foliation; its identifiable differential object is a projective normal field. Synchronization chooses an orientation, transnormal geometry provides conditions for a unit-speed potential, and discrete Hodge projection enforces global path independence. This separates the recovery of layer geometry from the choice of depth scale. Three distinct spectral quantities govern the analysis: local singular-value separation controls axis estimation, the connection-Laplacian gap controls orientation stability, and graph algebraic connectivity controls depth sensitivity. The results provide conditional geometric guarantees rather than a convergence theorem for neural training or a correctness theorem for source inference and endpoint selection.

### A.6 Proof of the identifiability theorem

By the chain rule, $J_z(s) = m'(d_0(s))\nabla d_0(s)^\top$. Both factors are nonzero under Assumption 1, so $J_z$ has rank one. Moreover,

$$J_z(s)^\top J_z(s) = \|m'(d_0(s))\|^2 \nabla d_0(s)\nabla d_0(s)^\top,$$

whose normalized leading eigenprojector is exactly $P_0(s)$.

Now suppose $z = \widetilde{m} \circ \widetilde{d}$ is another regular representation. The nullspace of $J_z(s)$ is the tangent line to a level set of $d_0$ and also to a level set of $\widetilde{d}$. Hence $\widetilde{d}$ is constant on every connected level set of $d_0$. Define $g(t)$ as this constant on $d_0^{-1}(t)$. Then $\widetilde{d} = g \circ d_0$. The submersion theorem provides local sections of $d_0$; composing $\widetilde{d}$ with these sections shows that $g$ is continuously differentiable. Differentiating gives $\nabla\widetilde{d} = g'(d_0)\nabla d_0$. Regularity makes $g'$ nonzero, and connectedness of the depth range forces its sign to be constant; therefore $g$ is strictly monotone. □

### A.7 Proof of the projector perturbation theorem

The singular values of $J_0 = an^\top$ are $\kappa$ and $0$. Wedin's sin-theta bound, applied to the one-dimensional leading right-singular subspace, gives

$$\sin\angle(\widehat{n}, n) \le \frac{\|E\|_2}{\kappa - \|E\|_2} \le \frac{2\,\|E\|_2}{\kappa}.$$

For unit vectors, the Frobenius distance between their rank-one projectors is $\sqrt{2}\sin\angle(\widehat{n}, n)$. Combining the two identities proves Eq. (15). □

### A.8 Proof of the graph-orientation proposition

Since $u_i = \tau_i n_i$ and $n_i^\top n_j > 0$ on every edge,

$$g_{ij} = \operatorname{sign}(u_i^\top u_j) = \tau_i\tau_j.$$

Let $A_g$ be the signed affinity with entries $a_{ij}g_{ij}$ and $D_{ii} = \sum_j a_{ij}$. For the vector $\tau = (\tau_1, \dots, \tau_n)^\top$, $(D - A_g)\tau = 0$. Conjugating by $\operatorname{diag}(\tau)$ transforms $D - A_g$ into the ordinary weighted graph Laplacian. Because $G$ is connected, that Laplacian has a one-dimensional nullspace. Thus the unnormalized connection Laplacian's nullspace is spanned by $\tau$. The normalized matrix $D^{-1/2}(D - A_g)D^{-1/2}$ has nullvector $D^{1/2}\tau$. Since the degrees are positive, both vectors have the same signs, and the implemented normalized spectral rounding recovers them up to a global sign. □

### A.9 Proof of the transnormal-depth corollary

First, the identity in Eq. (16) follows by setting $r = \|\nabla d_0\|$ and differentiating:

$$\partial_1\left(\frac{\partial_2 d_0}{r}\right) - \partial_2\left(\frac{\partial_1 d_0}{r}\right) = \frac{(\partial_1 d_0)(\partial_2 r) - (\partial_2 d_0)(\partial_1 r)}{r^2} = t^\top \nabla \log r.$$

The mixed derivatives cancel because $d_0 \in C^2$.

Let $G'(t) = 1/q(t)$ and $d_\star = G \circ d_0$. The chain rule gives

$$\nabla d_\star = G'(d_0)\nabla d_0 = \frac{\nabla d_0}{q(d_0)} = \frac{\nabla d_0}{\|\nabla d_0\|}.$$

Since $q > 0$, $G$ is strictly increasing. □

### A.10 Proof of the Hodge-projection theorem

The normal equations for Eq. (9) are $L\widehat{d} = B^\top W\omega$ on the centered subspace. Substituting the decomposition of $\omega$ gives

$$L\widehat{d} = Ld_\star + B^\top W\xi,$$

because $B^\top W\eta = 0$. Applying $L^\dagger$ on the centered subspace proves the first identity. Let $e = \widehat{d} - d_\star$. Then

$$Be = BL^\dagger B^\top W\xi = \Pi_\nabla \xi.$$

Weighted orthogonal projection is nonexpansive, so $\|e\|_L = \|Be\|_W = \|\Pi_\nabla \xi\|_W \le \|\xi\|_W$. Finally, $e^\top Le \ge \lambda_2(L)\,\|e\|_2^2$ for centered $e$, which gives Eq. (22). □

### A.11 Proof of the direction-to-depth corollary

The estimated edge integral differs from the trapezoidal rule applied to $n_\star$ by at most

$$\frac{1}{2}\ell_{ij}\left(\|n_i - n_\star(s_i)\| + \|n_j - n_\star(s_j)\|\right).$$

Because $n_\star = \nabla d_\star$, the exact line integral along the edge equals $d_\star(s_j) - d_\star(s_i)$. The trapezoidal error for the scalar integrand $\nabla d_\star(s(t))^\top(s_j - s_i)$ is $O(\ell_{ij}^3)$ under bounded third derivatives of $d_\star$. Combining the two bounds yields Eq. (23); applying Theorem 8 gives the nodewise statement. □

### A.12 Geometric bias and contour recovery

In addition to Assumption 1, suppose that $\Omega$ is bounded, connected, and Lipschitz, that $d_0$ extends smoothly to a neighborhood of its closure, and that $0 < r_- \le r = \|\nabla d_0(s)\| \le r_+ < \infty$ for $s \in \Omega$. Let $w : \Omega \to (0, \infty)$ be a fixed measurable weight satisfying $0 < w_- \le w(s) \le w_+ < \infty$ almost everywhere. Set $I = d_0(\Omega)$. For each $t \in I$, define the depth-$t$ contour as the level set $\Sigma_t = \{s \in \Omega : d_0(s) = t\}$. Each contour is connected under Assumption 1. Define

$$\mathcal{J}(u) = \int_\Omega w(s)\,\|\nabla u(s) - n(s)\|^2\,ds, \quad u_H = \operatorname*{arg\,min}_{u \in H^1(\Omega),\ \int_\Omega wu = 0} \mathcal{J}(u), \quad n = \frac{\nabla d_0}{r}, \tag{24}$$

where $H^1(\Omega) = \{u \in L^2(\Omega) : \partial_j u \in L^2(\Omega),\ j = 1, 2\}$ with the derivatives understood in the weak sense. For scalar or vector fields, write $\|f\|_w^2 = \int_\Omega w(s)\,\|f(s)\|^2\,ds$. The competing class $u = h(d_0)$ preserves every true contour. We allow locally absolutely continuous $h : I \to \mathbb{R}$ such that $h \circ d_0 \in H^1(\Omega)$, with the same weighted-mean centering. Every such function is constant on each true depth contour. When $h$ is strictly increasing, it preserves the true contours and their ordering exactly.

For almost every $t \in I$, define

$$A_j(t) = \int_{\Sigma_t} w(s)r(s)^{j-1}\,d\mathcal{H}^1(s), \qquad j = 0, 1, 2, \tag{25}$$

where $\mathcal{H}^1$ is arclength measure, and set $d\mu_t = w\, d\mathcal{H}^1/(A_0 r)$. The optimizer of Theorem 1 is unique up to its centering constant and has

$$h_F'(t) = \frac{A_1(t)}{A_2(t)}, \qquad \frac{1}{r_+} \le h_F'(t) \le \frac{1}{r_-} \quad \text{almost everywhere.} \tag{26}$$

Thus $h_F$ is absolutely continuous, Lipschitz, and strictly increasing; $u_F = h_F(d_0)$ has exactly the contours and ordering of $d_0$. The defect in Eq. (11) measures within-contour variation of gradient magnitude and is invariant under smooth increasing relabelings with positive derivative. It vanishes exactly when $r = q(d_0)$ for a positive continuous $q$.

**Corollary 10** (Approximate geometry recovery)**.** *Let $\widehat{n} \in L^2(\Omega, w; \mathbb{R}^2)$ be an estimated oriented field, and let $\widehat{u}$ be the weighted-mean-zero minimizer of the estimated-field analogue of Eq. (24), i.e., $\int_\Omega w(s)\, \|\nabla u(s) - \widehat{n}(s)\|^2\, ds$ over $u \in H^1(\Omega)$. If $\|\widehat{n} - n\|_w \le \delta$, then*

$$\|\nabla(\widehat{u} - u_F)\|_w \le \delta + \sqrt{\mathcal{D}_F - \mathcal{J}(u_H)}. \tag{27}$$

*Suppose additionally that, for some $\epsilon \ge 0$ and positive continuous function $q : I \to (0, \infty)$, $|r(s)/q(d_0(s)) - 1| \le \epsilon$ for $s \in \Omega$. Let $U = G(d_0)$, with $G' = 1/q$ and the same centering. Then $U \in H^1(\Omega)$ and*

$$\|\nabla(\widehat{u} - U)\|_w \le \delta + \epsilon\sqrt{\int_\Omega w}, \;\; \|\widehat{u} - U\|_w \le C_{\Omega,w}\left(\delta + \epsilon\sqrt{\int_\Omega w}\right), \tag{28}$$

*where $C_{\Omega,w}$ is the weighted Poincaré constant: $\|v\|_w \le C_{\Omega,w}\, \|\nabla v\|_w$ for every $v \in H^1(\Omega)$ with $\int_\Omega wv = 0$.*

These bounds separate field-estimation error from the geometric error introduced by unit-normal integration. They also yield a conditional-variance interpretation of contour coherence.

Suppose that, for almost every $t \in I$, the probability measure $\mu_t$ satisfies $\mathrm{Var}_{\mu_t}(v) \le C_{\mathrm{leaf}}^2 \int_{\Sigma_t} \|\nabla_{\Sigma_t} v\|^2\, d\mu_t$ for $v \in H^1(\Sigma_t)$ and a constant $C_{\mathrm{leaf}}$ independent of $t$. Here $\nabla_{\Sigma_t}$ denotes the tangential gradient, and the inequality is imposed on the entire depth contour $\Sigma_t$. Using the Sobolev restrictions to almost every depth contour, define $\overline{u}(t) = \int_{\Sigma_t} \widehat{u}(s)\, d\mu_t(s)$.

Then

$$\|\widehat{u} - \overline{u}(d_0)\|_w \le C_{\mathrm{leaf}}\left\{\delta + \sqrt{\mathcal{D}_F - \mathcal{J}(u_H)}\right\}. \tag{29}$$

Thus, locations on the same true contour receive similar depths in a weighted average sense when both error sources are small. Uniform contour location or topology guarantees would require stronger control. The finite-graph extension in Appendix A.12 separates these same error sources from edge discretization, without invoking a graph-to-continuum limit.

**Existence and admissibility.** Under the assumptions preceding Theorem 1, the weighted Poincaré inequality makes $\mathcal{J}$ coercive on the weighted-mean-zero subspace of $H^1(\Omega)$. The direct method and strict convexity in gradients give a unique centered minimizer $u_H$. Equivalently, gradients of $H^1$ functions form a closed subspace of weighted $L^2$, and $\nabla u_H$ is the orthogonal projection of $n$ onto that subspace. All projection statements below condition on the weight $w$.

The regularity and lower gradient bound permit the coarea formula (Fleming & Rishel, 1960) and the chain rule for Lipschitz compositions. Level integrals are finite for almost every $t$, since their integrals are bounded by the volume of $\Omega$, the weight bounds, and the bounds on $r$. A nonempty regular contour has positive length, so the relevant integrals are positive almost everywhere on $I$. These facts suffice even when the level integrals are not continuous.

**Proof of Theorem 1.** For an admissible locally absolutely continuous $h$, the chain rule on compact depth subintervals and coarea give

$$\begin{aligned}\mathcal{J}(h \circ d_0) &= \int_\Omega w\big(h'(d_0)r - 1\big)^2 \, ds \\ &= \int_I \int_{\Sigma_t} \frac{w}{r}\big(h'(t)r - 1\big)^2 \, d\mathcal{H}^1 \, dt \\ &= \int_I \left[A_2 h'^2 - 2A_1 h' + A_0\right] dt \\ &= \int_I A_2 \left(h' - \frac{A_1}{A_2}\right)^2 dt + \int_I \left(A_0 - \frac{A_1^2}{A_2}\right) dt.\end{aligned}$$

The ratio $A_2/A_1$ is a weighted average of $r$ along the contour under the measure with density $w$ relative to arclength. Consequently $r_- \le A_2/A_1 \le r_+$ almost everywhere. Its reciprocal defines a bounded positive derivative $h'_F$, and integration on the bounded interval $I$ gives a Lipschitz, strictly increasing $h_F$. Since $d_0$ is smooth with bounded gradient, $h_F(d_0) \in H^1(\Omega)$. Subtracting its weighted mean supplies the stated gauge. Completion of the square proves optimality and uniqueness up to constants, even if decreasing or nonmonotone competitors are admitted.

The conditional contour probability measure is

$$d\mu_t(s) = \frac{w(s)}{r(s)A_0(t)} \, d\mathcal{H}^1(s).$$

It satisfies $\mathbb{E}_{\mu_t}[r] = A_1/A_0$ and $\mathbb{E}_{\mu_t}[r^2] = A_2/A_0$. Substitution proves the variance expression in Eq. (11). This expression is nonnegative. It vanishes exactly when $r$ is constant almost everywhere on almost every connected contour. Continuity and regular foliation coordinates then imply zero tangential derivative of $r$ throughout $\Omega$. Connectedness of each level set yields $r = q(d_0)$ for a positive function $q$. Conversely, this condition makes every conditional variance zero.

For invariance, let $\widetilde{d} = g(d_0)$ with $g' > 0$ smooth. The leaves are unchanged and $\widetilde{r} = g'(t)r$. At $\widetilde{t} = g(t)$, the transformed integrals satisfy

$$\widetilde{A}_0(\widetilde{t}) = \frac{A_0(t)}{g'(t)}, \quad \widetilde{A}_1(\widetilde{t}) = A_1(t), \quad \widetilde{A}_2(\widetilde{t}) = g'(t)A_2(t).$$

Thus $\widetilde{h}'_F(g(t))g'(t) = h'_F(t)$, so the composed coordinate is unchanged up to a constant. Changing variables in the defect integral proves invariance of $\mathcal{D}_F$. The same physical weight is used before and after relabeling.

Finally, the normal equation for $u_H$ is

$$\int_\Omega w(\nabla u_H - n)^\top \nabla v \, ds = 0 \ \text{ for every } v \in H^1(\Omega).$$

Take $v = u_F - u_H$. Expanding $\mathcal{J}(u_F)$ and canceling the cross term gives

$$\mathcal{J}(u_F) = \mathcal{J}(u_H) + \|\nabla(u_F - u_H)\|_w^2,$$

which proves the exact bias identity. The gradient discrepancy contains both normal and tangential components; the identity does not equate it solely with tangential distortion. □

**Proof of Corollary 10.** Let $\mathsf{P}_w$ denote the weighted orthogonal projection onto gradients. Then

$$\nabla\widehat{u} - \nabla u_H = \mathsf{P}_w(\widehat{n} - n), \quad \|\nabla\widehat{u} - \nabla u_H\|_w \le \delta.$$

The triangle inequality and Eq. (12) prove Eq. (27). The approximate-transnormal condition implies $G' = 1/q \le (1+\epsilon)/r_-$, so $G$ is Lipschitz and $U = G(d_0)$ is admissible. Moreover,

$$\nabla U = \frac{r}{q(d_0)} n, \quad \|n - \nabla U\|_w \le \epsilon \sqrt{\int_\Omega w}.$$

Because $\nabla U$ belongs to the projection space,

$$\nabla\widehat{u} - \nabla U = \mathsf{P}_w(\widehat{n} - \nabla U).$$

Nonexpansiveness and the triangle inequality prove the gradient bound in Eq. (28). Both potentials have zero weighted mean, so weighted Poincaré proves its second bound. The term involving $\epsilon$ can be replaced by the sharper quantity $\|r/q(d_0) - 1\|_w$. □

**Contour-coherence bound.** Write $T = (-n_2, n_1)$ for the unit tangent. The sufficient leafwise assumption used in Eq. (29) is

$$\mathrm{Var}_{\mu_t}(f) \le C_{\mathrm{leaf}}^2 \mathbb{E}_{\mu_t}\left[\left|T^\top \nabla f\right|^2\right] \quad \text{for almost every } t \tag{30}$$

for admissible restrictions of $H^1(\Omega)$ functions. Coarea and this inequality give

$$\begin{aligned}\|\widehat{u} - \overline{u}(d_0)\|_w^2 &= \int_I A_0(t)\,\mathrm{Var}_{\mu_t}(\widehat{u})\,dt \\ &\le C_{\mathrm{leaf}}^2 \int_\Omega w\left|T^\top \nabla \widehat{u}\right|^2 ds \\ &\le C_{\mathrm{leaf}}^2 \left\|\nabla(\widehat{u} - u_F)\right\|_w^2,\end{aligned}$$

because $T^\top \nabla u_F = 0$. Apply Eq. (27). This proves an average within-contour bound. Neither it nor an $H^1$ error bound in two dimensions alone establishes uniform contour displacement or topology preservation.

There is a useful connection to Eq. (16). If the ratio of the largest to the smallest $r$ on each contour is at most $K$, applying Eq. (30) to $r$ and using $T^\top \nabla r = r \operatorname{curl} n$ yields

$$\mathcal{D}_F \le C_{\mathrm{leaf}}^2 K^2 \left\|\operatorname{curl} n\right\|_w^2. \tag{31}$$

Indeed, divide the resulting conditional variance bound by $\mathbb{E}_{\mu_t}[r^2]$, bound the within-contour squared magnitude ratio by $K^2$, and integrate against $A_0(t)dt$. The uniform leafwise Poincaré assumption is substantive; long or poorly connected contours can weaken this guarantee.

**A variable-width example with perfect local axes.** For $\beta > 0$, consider

$$\Omega = \{(x,y) : 0 < y < 1,\ 0 < x < 1 + \beta y\},\ \ d_0(x,y) = \frac{x}{1+\beta y},\ \ w = 1.$$

Its contours are straight but nonparallel. Let $t = d_0$, $b(t) = \sqrt{1+\beta^2 t^2}$, and $c_\beta = \log(1+\beta)/\beta$. Direct differentiation gives

$$\nabla d_0 = \frac{(1, -\beta t)^\top}{1+\beta y},\ \ r = \frac{b(t)}{1+\beta y},\ \ n = \frac{(1,-\beta t)^\top}{b(t)},\ \ \operatorname{curl} n = -\frac{\beta}{(1+\beta y)b(t)}.$$

Parametrizing a contour by $(x,y) = (t(1+\beta y), y)$ gives $d\mathcal{H}^1 = b(t)dy$. Therefore

$$A_0 = 1 + \frac{\beta}{2},\ \ A_1 = b(t),\ \ A_2 = b(t)^2 c_\beta,\ \ h_F(t) = \frac{\operatorname{arsinh}(\beta t)}{\beta c_\beta} + C,$$

and

$$\mathcal{D}_F = 1 + \frac{\beta}{2} - \frac{\beta}{\log(1+\beta)} = \frac{\beta^2}{12} + O(\beta^3) > 0. \tag{32}$$

The strict positivity follows also from nonconstant $r$ along every contour.

Here the unconstrained Hodge solution must depart from the true foliation. Its natural boundary condition is $(\nabla u_H - n)^\top \nu = 0$, with outward unit normal $\nu$. But at the interior of the boundary segment $y = 0$,

$$(\nabla u_F - n)^\top \nu = \left(\frac{1}{c_\beta} - 1\right)\frac{\beta t}{b(t)} > 0 \ \ (t > 0).$$

Thus $u_F$ is not a stationary point of the unrestricted problem, and $\mathcal{J}(u_H) < \mathcal{D}_F$. If $u_H$ were constant on every true contour, it would belong to the restricted class, contradicting optimality of $u_F$ there. An exact feature field $m(d_0)$ with nonzero derivative supplies perfect local axes in this example. The resulting contour distortion is caused by imposing unit speed, rather than by local estimation error.

**Finite-graph normalization bias.** The population statements above do not establish a graph-to-continuum limit. A direct finite-graph bound nevertheless separates the same errors. Suppose $U = G(d_0)$ is $C^3$ on a neighborhood of every straight edge segment, let $U_i = U(s_i)$, and assume

$$\|\widehat{n}_i - n(s_i)\| \le \delta_i, \quad \left| \frac{r(s_i)}{q(d_0(s_i))} - 1 \right| \le \epsilon_i.$$

For the implemented trapezoidal increment,

$$|\widehat{\omega}_{ij} - (U_j - U_i)| \le \frac{\ell_{ij}}{2} (\delta_i + \delta_j + \epsilon_i + \epsilon_j) + C\ell_{ij}^3 =: b_{ij}. \tag{33}$$

To prove this, compare each endpoint normal with $\nabla U = (r/q(d_0))n$, and then bound the trapezoidal error for the exact line integral of $\nabla U$. For any fixed connected weighted graph, Theorem 8 consequently gives, after common centering,

$$\left\| \widehat{d} - U \right\|_L \le \|b\|_W, \quad \left\| \widehat{d} - U \right\|_2 \le \lambda_2(L)^{-1/2} \|b\|_W.$$

The three contributions in Eq. (33) are orientation-estimation error, unit-normal model bias, and edge discretization. Smoothness along the entire edge segment is required; for an edge crossing a hole, regularity only on the observed tissue does not supply this assumption.

### A.13 Ordering recovery: hypotheses and proof

This subsection specifies the graph construction and local-error hypotheses used in Theorem 2, and proves its simultaneous ordering guarantee. All statements are deterministic for the realized graph and weights. Matrix norms without a subscript $F$ are operator norms, and vector norms are Euclidean.

**Graph and trapezoidal integration.** Let $G = (V, E)$ be a connected undirected graph with $n \ge 2$ vertices and positive spatial weights, and let $L_g$ be its weighted Laplacian. Orient its edges arbitrarily, with incidence convention $(Bd)_e = d_j - d_i$ for $e = (i, j)$. Write $x_e = s_j - s_i$ and $\ell_e = \|x_e\|_2$, and define $\mathsf{A} : \mathbb{R}^{2n} \to \mathbb{R}^{|E|}$ by

$$(\mathsf{A}v)_e = \tfrac{1}{2}(v_i + v_j)^\top x_e. \tag{34}$$

Both endpoint blocks have the same sign. Errors on incident edges therefore share contributions from the same estimated normal.

**Projector refinement and synchronization.** At each node with positive fidelity $c_i^P$, let $\widetilde{P}_i$ be the rank-one orthogonal projector onto a leading right singular vector of $\widehat{J}_i$. At zero-fidelity nodes, choose any rank-one orthogonal projector. The componentwise solve in Eq. (62) is

$$H = (C_P + \lambda_P L_g + \epsilon_P I)^{-1} C_P, \qquad C_P = \mathrm{diag}(c_i^P). \tag{35}$$

Assume $\lambda_P > 0$, $\epsilon_P > 0$, $c_i^P \ge 0$, and $\sum_i c_i^P > 0$; the implementation uses $\epsilon_P = 10^{-8}$. Put $q_i = \sum_j H_{ij}$ and $Q_i = \sum_j H_{ij}\widetilde{P}_j$. As shown below, $H$ is entrywise nonnegative and $q_i > 0$, so $Q_i/q_i$ is a convex combination of the raw projectors and has the same leading eigenspace as $Q_i$.

Let $v_i$ be a computed unit leading eigenvector of $Q_i$, with arbitrary sign. Let $S$ be the symmetric signed adjacency used in synchronization, with positive edge magnitudes and $\mathrm{sign}(S_{ij}) = \mathrm{sign}(v_i^\top v_j)$ on every edge of $G$. For a positive diagonal matrix $D_S$, let $\xi$ be an eigenvector of $D_S^{-1/2} S D_S^{-1/2}$ associated with its largest algebraic eigenvalue. The synchronized normals are $\widehat{n}_i = \mathrm{sign}(\xi_i) v_i$. The hypotheses below ensure that these signs are nonzero. This formulation includes degree normalization and any other positive diagonal normalization.

**Hodge solution and pair influences.** For the realized integration weights $W = \mathrm{diag}(a_e)$, with $a_e > 0$, put $L = B^\top W B$. The centered Hodge estimator is

$$\widehat{d} = L^\dagger B^\top W \mathsf{A} \widehat{n}, \qquad \mathbf{1}^\top \widehat{d} = 0, \tag{36}$$

where $L^\dagger$ is the Moore–Penrose pseudoinverse. For distinct vertices $u, v$, define

$$\begin{aligned} h_{uv} &= e_v - e_u, \quad & j^{uv} &= W B L^\dagger h_{uv}, \\ \zeta^{uv} &= \mathsf{A}^\top j^{uv}, \quad & R_{uv} &= h_{uv}^\top L^\dagger h_{uv}, \end{aligned} \tag{37}$$

where $e_i$ is the $i$th coordinate vector. Here $j^{uv}$ is the electrical unit flow from $u$ to $v$ with conductances $a_e$, and $R_{uv}$ is effective resistance (Barooah & Hespanha, 2008). For fixed geometry and weights, $\zeta_i^{uv} \in \mathbb{R}^2$ measures the sensitivity of the recovered pair difference to node $i$'s normal.

**Target field and local-error hypotheses.** Let $d_\star$ be $C^3$ on an open neighborhood containing every graph-edge segment, with $n_\star = \nabla d_\star$ and $\|n_\star\| = 1$ there. Assume that the operator norm of its third derivative is at most $M$ on those segments. Write $n_i^\star = n_\star(s_i)$ and $P_i^\star = n_i^\star (n_i^\star)^\top$. At each node with $c_i^P > 0$, assume

$$J_i^\star = b_i (n_i^\star)^\top, \qquad \kappa_i = \|b_i\|_2 > 0, \qquad \left\| \widehat{J}_i - J_i^\star \right\|_2 \le \eta_i \le \kappa_i / 2.$$

Set $p_i = 2\eta_i/\kappa_i$ at these nodes and $p_i = 1$ at zero-fidelity nodes. Define

$$\begin{aligned} t_i &= \frac{\sum_j H_{ij} \left( p_j + \left\| P_j^\star - P_i^\star \right\|_2 \right)}{q_i}, \\ \varepsilon_i &= \arcsin\left( \frac{t_i}{1 - t_i} \right), \qquad\qquad \delta_i = 2 \sin(\varepsilon_i / 2), \end{aligned} \tag{38}$$

and suppose $t_i < 1/2$ for every node. The numerator separates propagated Jacobian error from smoothing bias due to different population orientations. Require, on every synchronization edge,

$$\alpha_{ij} + \varepsilon_i + \varepsilon_j < \pi/2, \qquad \alpha_{ij} = \arccos\left( (n_i^\star)^\top n_j^\star \right). \tag{39}$$

**Proof: projector error and the stabilized smoother.** For rank-one orthogonal projectors in $\mathbb{R}^2$,

$$\|P - P'\|_F = \sqrt{2} \, \|P - P'\|_2 , \qquad \|P - P'\|_2 \le 1.$$

Thus Theorem 5 gives $\left\| \widetilde{P}_i - P_i^\star \right\|_2 \le p_i$ at positive-fidelity nodes; at zero-fidelity nodes the bound holds with $p_i = 1$.

Let $T = C_P + \lambda_P L_g + \epsilon_P I$. It is an irreducible, strictly diagonally dominant $M$-matrix: its off-diagonal entries are nonpositive, the underlying positive-edge graph is connected, and its diagonal dominance margin is $c_i^P + \epsilon_P > 0$. Its inverse is entrywise positive. Therefore $H = T^{-1} C_P$ is nonnegative and every row sum $q_i$ is positive. Moreover,

$$H\mathbf{1} = \mathbf{1} - \epsilon_P T^{-1} \mathbf{1},$$

so $q_i \le 1$. If $c_j^P = 0$, the entire $j$th column of $H$ vanishes: an unsupported raw projector contributes no input error to the fused matrices.

The normal equations in Eq. (62) give $Q_i = \sum_j H_{ij} \widetilde{P}_j$ componentwise, and hence

$$\begin{aligned} \|Q_i - q_i P_i^\star\|_2 &\le \sum_j H_{ij} \left( \left\| \widetilde{P}_j - P_j^\star \right\|_2 + \left\| P_j^\star - P_i^\star \right\|_2 \right) \\ &\le q_i t_i. \end{aligned}$$

Comparing with $q_i P_i^\star$ retains the numerical ridge without treating its row scale as angular bias.

**Refined-axis error.** Write $Q_i / q_i = P_i^\star + F_i$, with $\|F_i\|_2 \le t_i < 1/2$. Weyl's inequality gives a leading eigenvalue $\widehat{\lambda}_{i1} \ge 1 - t_i$ and a second eigenvalue at most $t_i$, so the leading eigenspace is one-dimensional. Let $\overline{v}_i$ denote its target-aligned unit representative, with $(n_i^\star)^\top \overline{v}_i > 0$. Projecting its eigenvector equation onto the orthogonal complement of $n_i^\star$ gives

$$\widehat{\lambda}_{i1} (I - P_i^\star) \overline{v}_i = (I - P_i^\star) F_i \overline{v}_i.$$

Consequently, for $\theta_i = \angle(\overline{v}_i, n_i^\star)$,

$$\sin\theta_i \le \frac{t_i}{1 - t_i}, \qquad \theta_i \le \varepsilon_i, \qquad \|\overline{v}_i - n_i^\star\|_2 = 2 \sin(\theta_i / 2) \le \delta_i.$$

**Exact orientation after spectral rounding.** Write the computed eigenvectors as $v_i = \tau_i \overline{v}_i$, with $\tau_i \in \{-1, 1\}$. The angular triangle inequality and Eq. (39) imply

$$\angle(\overline{v}_i, \overline{v}_j) \le \varepsilon_i + \alpha_{ij} + \varepsilon_j < \pi/2.$$

Therefore $\overline{v}_i^\top \overline{v}_j > 0$ and $\operatorname{sign}(v_i^\top v_j) = \tau_i \tau_j$ on every edge. Let $A_+$ be the unsigned adjacency with entries $|(S)_{ij}|$, and put $T_\tau = \operatorname{diag}(\tau_i)$. Then $S = T_\tau A_+ T_\tau$, so

$$D_S^{-1/2} S D_S^{-1/2} = T_\tau \big(D_S^{-1/2} A_+ D_S^{-1/2}\big) T_\tau.$$

For every positive diagonal $D_S$, the matrix in parentheses is symmetric, nonnegative, and irreducible. By Perron–Frobenius, its largest algebraic eigenvalue is simple and has a strictly positive eigenvector $z$. Hence $\xi = cT_\tau z$ for some nonzero scalar $c$, and spectral rounding gives $\operatorname{sign}(\xi_i) = \sigma\tau_i$ with the single global sign $\sigma = \operatorname{sign}(c)$. It follows that

$$\sigma \widehat{n}_i = \overline{v}_i, \qquad \|\sigma \widehat{n}_i - n_i^\star\|_2 \le \delta_i \quad \text{for every } i.$$

Degree normalization is the special case $D_{S,ii} = \sum_j (A_+)_{ij}$, for which the Perron eigenvalue is $1$ and $z = D_S^{1/2}\mathbf{1}$. The general argument does not require this normalization or an additional entrywise perturbation bound.

**Trapezoidal discrepancy and pair influence.** Let $d_\star^V$ be the centered vector of sampled target depths, stack $n_i^\star$ into $n_\star^V$, and define the globally aligned solution and normal error by

$$\widehat{d}^{\,\mathrm{or}} = \sigma \widehat{d}, \qquad e_i^n = \sigma \widehat{n}_i - n_i^\star, \qquad r = \mathsf{A} n_\star^V - B d_\star^V.$$

For an edge segment $s(t) = s_i + t x_e$, the integral of $f_e(t) = \nabla d_\star(s(t))^\top x_e$ is $d_\star(s_j) - d_\star(s_i)$. Since $|f_e''(t)| = |D^3 d_\star(s(t))[x_e, x_e, x_e]| \le M\ell_e^3$, the trapezoidal remainder on $[0,1]$ gives

$$|r_e| \le M\ell_e^3/12. \tag{40}$$

By Eq. (34), the aligned edge field satisfies exactly

$$\sigma \mathsf{A} \widehat{n} = B d_\star^V + \mathsf{A} e^n + r.$$

The centered normal equations, or Theorem 8, yield

$$\widehat{d}^{\,\mathrm{or}} - d_\star^V = L^\dagger B^\top W (\mathsf{A} e^n + r). \tag{41}$$

Multiplication by $h_{uv}^\top$ gives the exact pair identity

$$\begin{aligned} h_{uv}^\top (\widehat{d}^{\,\mathrm{or}} - d_\star^V) &= (j^{uv})^\top (\mathsf{A} e^n + r) \\ &= \sum_i (\zeta_i^{uv})^\top e_i^n + (j^{uv})^\top r. \end{aligned} \tag{42}$$

Using $\|e_i^n\|_2 \le \delta_i$ and Eq. (40) proves Eq. (13), simultaneously for all pairs and for the same global sign. Every positive true depth difference larger than its error radius remains positive, proving the ordering statement. □

**Effective-resistance envelope.** The sharper pair bound in Eq. (13) retains cancellation among incident-edge contributions to each $\zeta_i^{uv}$. A simpler envelope is

$$\begin{aligned} E_{uv} \le \sqrt{R_{uv}} \Bigg[ &\left\{ \sum_{e=(i,j)} \frac{a_e \ell_e^2}{4} (\delta_i + \delta_j)^2 \right\}^{1/2} \\ &+ \frac{M}{12} \left\{ \sum_e a_e \ell_e^6 \right\}^{1/2} \Bigg]. \end{aligned} \tag{43}$$

Indeed, the definition of $j^{uv}$ gives

$$B^\top j^{uv} = h_{uv}, \qquad (j^{uv})^\top W^{-1} j^{uv} = h_{uv}^\top L^\dagger h_{uv} = R_{uv}.$$

Since both endpoint blocks of $\mathsf{A}$ equal $x_e^\top/2$,

$$\zeta_i^{uv} = \tfrac{1}{2} \sum_{e \ni i} j_e^{uv} x_e.$$

No additional incidence sign appears. This sum is invariant to an edge reversal because both $x_e$ and $j_e^{uv}$ reverse. The triangle inequality and weighted Cauchy–Schwarz give

$$\sum_i \|\zeta_i^{uv}\|_2 \, \delta_i \le \tfrac{1}{2} \sum_{e=(i,j)} |j_e^{uv}| \ell_e (\delta_i + \delta_j)$$
$$\le \sqrt{R_{uv}} \left\{ \sum_{e=(i,j)} \frac{a_e \ell_e^2}{4} (\delta_i + \delta_j)^2 \right\}^{1/2}.$$

The same argument bounds $\sum_e |j_e^{uv}| \ell_e^3$ and proves Eq. (43). Finally, $h_{uv} \perp \mathbf{1}$ and $\|h_{uv}\|_2^2 = 2$ imply $R_{uv} \le 2/\lambda_2(L)$, yielding a global spectral envelope. Neither bound assumes independent edge errors.

**Scope and probabilistic interpretation.** The deterministic result allows data-dependent confidence weights. If the simultaneous Jacobian bounds and geometric margins hold on an event of probability at least $1 - \beta$, the simultaneous ordering guarantee holds with at least the same probability. Because the margins involve population quantities, this implication alone does not provide an observable confidence certificate. The result concerns the centered potential before range normalization; positive affine normalization preserves its ordering.

Choosing $d_\star$ as the transnormal representative in Corollary 7 makes it a strictly increasing reparameterization of $d_0$, so the guarantee recovers the ordering of $d_0$. Nontransnormal geometry introduces the additional geometric bias analyzed in Theorem 1. The E1/E2 construction and masked-fidelity variants, including the documented E3 solve, are covered when their realized graph is connected and their smoothing and synchronization steps have the forms above. The result does not establish the correctness of mask construction, source inference, or endpoint selection.

### A.14 Conditional uncertainty with shared normal errors

In this subsection, $\widehat{d}$ and the estimated normals denote the estimates after the common global alignment in Theorem 2. The shared-node structure also determines uncertainty propagation. Condition on a sigma-field $\mathcal{F}$ that fixes the graph, $W$, $\mathsf{A}$, and the sampled target vectors $d_\star^V$ and $n_\star^V$, hence also $r$. For the actual conditional law of the stacked oriented-normal error, write

$$\mu = \mathbb{E}[e^n \mid \mathcal{F}], \;\; \Sigma = \mathrm{Cov}(e^n \mid \mathcal{F}), \;\; K = L^\dagger B^\top W \mathsf{A}.$$

Assuming finite second moments, Eq. (41) implies

$$\mathbb{E}[\widehat{d} - d_\star^V \mid \mathcal{F}] = K\mu + L^\dagger B^\top W r, \tag{44}$$
$$\mathrm{Cov}(\widehat{d} \mid \mathcal{F}) = K \Sigma K^\top, \tag{45}$$
$$\mathrm{Var}(\widehat{d}_v - \widehat{d}_u \mid \mathcal{F}) = (\zeta^{uv})^\top \Sigma \zeta^{uv}. \tag{46}$$

These are exact linear propagation identities; $\Sigma$ includes dependencies from shared observations, overlapping local neighborhoods, and projector refinement. Conditioning on estimated weights does not make those errors independent or centered. Likewise, holding weights fixed numerically in a resampling exercise does not by itself establish the conditional law assumed here. Nonlinear range normalization changes the covariance, although it preserves pair ordering.

For interpretation in two spatial dimensions, let $t_i^\star$ be the unit tangent and write a small oriented angular error as

$$n_i = \cos(\theta_i) n_i^\star + \sin(\theta_i) t_i^\star.$$

Then $e_i^n = \theta_i t_i^\star - \theta_i^2 n_i^\star/2 + O(|\theta_i|^3)$. To first order, with $g_i^{uv} = (t_i^\star)^\top \zeta_i^{uv}$, the pair variance is $(g^{uv})^\top \mathrm{Cov}(\theta \mid \mathcal{F}) g^{uv}$. An edge parallel to both endpoint normals has zero first-order angular error in its trapezoidal increment; oblique edges need not. This first-order statement is not an exact covariance formula for finite angular errors.

The familiar identity $\mathrm{Var}(\widehat{d}_v - \widehat{d}_u) = \sigma^2 R_{uv}$ would require the appropriate edge-error covariance, for example $\mathsf{A}\Sigma\mathsf{A}^\top = \sigma^2 W^{-1}$ in the conditional setting above. Such a covariance is not implied by this estimator. Its confidence weights are not established inverse noise variances, and incident edges reuse estimated normals. Therefore neither minimum-variance optimality nor monotone variance reduction from adding edges follows from the electrical-flow representation alone.

### A.15 CANONICAL MASS COORDINATES AND TRANSPORT BOUNDS

We use the usual generalized quantile $Q_\nu(v) = \inf\{t : F_\nu(t) \geq v\}$ for $0 < v < 1$, extended to the endpoints when the support is compact. For scalar probability measures with finite $p$th moments, $W_p$ denotes the Wasserstein distance with cost $|x - y|^p$. Increasing quantile matching is optimal in one dimension (Santambrogio, 2015).

**Measured foliations and canonical coordinates.** Foliation recovery leaves a monotone ambiguity in the numerical depth scale. A spatial probability measure $\mu$ fixes an interpretable coordinate within this equivalence class: the fraction of tissue mass below each contour. Equal spot weights, tissue area, and cell mass define different measures and therefore different coordinates. We fix both $\mu$ and the orientation of increasing depth. For a scalar field $d$, write $F_d(t) = \mu\{s : d(s) \leq t\}$ and let $d_\#\mu$ denote its depth distribution. The probability integral transform and one-dimensional monotone transport provide the calibration (Santambrogio, 2015).

**Proposition 11** (Canonical coordinate of a measured foliation)**.** *Under Assumption 1, suppose $\Omega$ is open, $\mu$ has full support and a positive density on $\Omega$, and $(d_0)_\#\mu$ has a positive continuous density $f_0$ on $I = d_0(\Omega)$. Let $\nu$ have a positive continuous density on a compact interval $J$, with quantile function $Q_\nu$. Both coordinates*

$$r_0(s) = F_{d_0}(d_0(s)), \quad u_\nu(s) = Q_\nu(r_0(s)) \tag{47}$$

*are invariant under increasing reparameterizations of $d_0$. The coordinate $u_\nu$ is the unique $C^1$ coordinate with the same oriented leaves and distribution $\nu$ under $\mu$.*

Thus an oriented foliation, a spatial measure, and a prescribed law jointly identify a numerical coordinate. Choosing $\nu = \mathrm{Unif}[0, 1]$ gives the canonical mass coordinate $r_0$. Without orientation, the two choices are $Q_\nu(r_0)$ and $Q_\nu(1 - r_0)$. The prescribed law fixes a convention; it is not a physical depth law inferred from expression. The unit-speed representative in Corollary 7 instead fixes a geometric scale through the ambient metric. Mass calibration is an optional theoretical output, not a transformation applied in the reported experiments.

Across sections, equal mass coordinates identify equal cumulative mass under the chosen measures. Interpreting these positions as anatomically homologous requires comparable tissue coverage and biological mass. Differences in cell abundance or sampled tissue extent can change this correspondence even when both foliations are recovered accurately.

**Recovery from estimated depths.** Consider $N$ locations and a fixed population oracle $d_\star$. Set $x_i = d_\star(s_i)$, $y_i = \widehat{d}_i$, and $\|v\|_{2,N} = (N^{-1}\sum_i v_i^2)^{1/2}$. Define

$$E_N = \inf_{c \in \mathbb{R}} \|y + c\mathbf{1} - x\|_{2,N}\,, \tag{48}$$

$$\Delta_N = \sup_t \left| N^{-1} \sum_{i=1}^{N} \mathbf{1}\{x_i \leq t\} - F_{d_\star}(t) \right|. \tag{49}$$

Here $E_N$ measures depth error after removing the arbitrary offset; $\Delta_N$ measures sampling discrepancy along the oracle coordinate.

Let $R_i$ be the fitted-depth rank used in Theorem 3, and set $\widehat{r}_i = R_i/N$ and $r_{\star,i} = F_{d_\star}(x_i)$.

**Proof of Proposition 11.** The assumed positive continuous densities make $F_{d_0} : I \to (0, 1)$ and $Q_\nu : (0, 1) \to \mathrm{int}(J)$ continuously differentiable and strictly increasing. Consequently $u_\nu$ is $C^1$, has the same oriented leaves as $d_0$, and has no critical points. Since $F_{d_0}(d_0(S))$ is uniform when $S \sim \mu$, its composition with $Q_\nu$ has law $\nu$.

Let $\widetilde{d} = h \circ d_0$ for an increasing reparameterization $h$. The sublevel sets satisfy

$$\{s : \widetilde{d}(s) \leq h(t)\} = \{s : d_0(s) \leq t\}, \quad F_{\widetilde{d}}(h(t)) = F_{d_0}(t).$$

Thus both $r_0$ and $u_\nu$ are invariant to $h$.

Any other $C^1$ coordinate $v$ with the same oriented leaves has the form $v = g \circ d_0$ with $g$ strictly increasing. To see this directly, $v$ is constant on each connected leaf, and the submersion property

supplies a local transverse section on which $g$ is $C^1$. The local definitions agree, giving a global function on $I$. If $v_{\#}\mu = \nu$, then

$$F_\nu(g(t)) = \mu\{s : v(s) \le g(t)\} = F_{d_0}(t).$$

Inverting $F_\nu$ gives $g(t) = Q_\nu(F_{d_0}(t))$, proving uniqueness. The equality holds throughout $\Omega$; alternatively an almost-everywhere equality extends by continuity and full support. Reversing orientation replaces $r_0$ by $1 - r_0$, since the oracle law is atomless. □

**Population stability of mass calibration.** The following lemma isolates the transport argument before its empirical version.

**Lemma 12** (Population rank stability)**.** *Let $d, \widehat{d} \in L^p(\mu)$, $1 \le p < \infty$, and suppose both scalar laws are atomless. If $F_d$ is $M$-Lipschitz, then*

$$\left\| F_{\widehat{d}} \circ \widehat{d} - F_d \circ d \right\|_{L^p(\mu)} \le 2M \inf_{c \in \mathbb{R}} \left\| \widehat{d} + c - d \right\|_{L^p(\mu)}. \tag{50}$$

*If $Q_\nu$ is $K$-Lipschitz, applying it to both mass coordinates multiplies the bound by $K$.*

**Proof.** First fix the additive gauge and set $T = Q_{d_{\#}\mu} \circ F_{\widehat{d}}$. The map $T$ transports $\widehat{d}_{\#}\mu$ to $d_{\#}\mu$ by increasing matching. Therefore

$$\begin{aligned} \left\| T \circ \widehat{d} - \widehat{d} \right\|_{L^p(\mu)} &= W_p(\widehat{d}_{\#}\mu, d_{\#}\mu) \\ &\le \left\| \widehat{d} - d \right\|_{L^p(\mu)}. \end{aligned}$$

The inequality uses the coupling induced by the common location $s$. Continuity of $F_d$ and the probability integral transform give $F_d(T(\widehat{d})) = F_{\widehat{d}}(\widehat{d})$ almost everywhere. Hence

$$\begin{aligned} \left\| F_{\widehat{d}} \circ \widehat{d} - F_d \circ d \right\|_{L^p(\mu)} &\le M \left\| T \circ \widehat{d} - d \right\|_{L^p(\mu)} \\ &\le M \big( W_p(\widehat{d}_{\#}\mu, d_{\#}\mu) + \left\| \widehat{d} - d \right\|_{L^p(\mu)} \big) \\ &\le 2M \left\| \widehat{d} - d \right\|_{L^p(\mu)}. \end{aligned}$$

Replacing $\widehat{d}$ by $\widehat{d} + c$ does not change its mass coordinate. Taking the infimum over $c$ proves Eq. (50). The target-law statement follows from Lipschitz continuity of $Q_\nu$. □

**Proof of Theorem 3.** Let $F = F_{d_\star}$ and $F_N^\star(t) = N^{-1} \sum_i \mathbf{1}\{x_i \le t\}$. Choose an offset attaining $E_N$ and absorb it into $y$; this leaves all estimated ranks unchanged. First resolve estimated ties into ordinal ranks $R_i \in \{1, \dots, N\}$ by any fixed ordering. Write $x_{(1)} \le \cdots \le x_{(N)}$ and assign

$$t_i = x_{(R_i)}.$$

With $\nu_N^\star = N^{-1} \sum_i \delta_{x_i}$ and $\widehat{\nu}_N = N^{-1} \sum_i \delta_{y_i}$, optimality of increasing matching gives

$$\|t - y\|_{2,N} = W_2(\nu_N^\star, \widehat{\nu}_N) \le \|x - y\|_{2,N} = E_N.$$

Consequently $\|t - x\|_{2,N} \le 2E_N$.

Oracle ties cause no difficulty. For every $i$,

$$F_N^\star(t_i-) \le \frac{R_i}{N} \le F_N^\star(t_i).$$

Both endpoints differ from $F(t_i)$ by at most $\Delta_N$: this is immediate for the right endpoint and follows by taking a left limit for the other, since $F$ is continuous. Therefore

$$\left| \frac{R_i}{N} - F(x_i) \right| \le \Delta_N + |F(t_i) - F(x_i)| \le \Delta_N + M|t_i - x_i|.$$

Taking the normalized Euclidean norm and using the preceding transport bound proves Eq. (14) for ordinal ranks.

To handle average ranks at estimated ties, average the ordinal rank vectors over all permutations within each tied block. Their average is exactly the average-rank vector. Every ordinal tie breaking obeys the same bound because the scalar matching remains optimal and $E_N, \Delta_N$ do not change. Convexity of the norm therefore proves the stated bound for average ranks. This argument is deterministic and makes no independence assumption about $y$, including its dependence on the observed locations and features. □

The ranks in Theorem 3 use the values $1/N, \dots, 1$. If centered plotting positions $(R_i - 1/2)/N$ are used instead, the triangle inequality adds $1/(2N)$ to the bound. An upper empirical-CDF rank differs from an average rank at estimated ties, so it should not be silently substituted for the stated convention.

**Composition with Hodge integration and oracle sampling.** For a correctly oriented Hodge oracle, let $\Pi_\nabla = BL^\dagger B^\top W$ denote the weighted gradient projection and $\|v\|_W^2 = v^\top W v$. Then Theorem 8 gives

$$\|\widehat{r} - r_\star\|_{2,N} \le \frac{2M\,\|\Pi_\nabla \xi\|_W}{\sqrt{N\lambda_2(L)}} + \Delta_N. \tag{51}$$

Under Corollary 7, $r_\star = r_0$ because $d_\star$ is an increasing reparameterization of $d_0$. The result therefore connects local integration errors and graph connectivity to recovery of a coordinate invariant to monotone depth reparameterization. A $K$-Lipschitz target quantile multiplies the bound by $K$.

Under Theorem 8, choose a common orientation and center the sampled oracle vector. Let $e = \widehat{d} - d_\star$. The theorem and the Poincaré inequality give

$$E_N \le \frac{\|e\|_2}{\sqrt{N}} \le \frac{\|e\|_L}{\sqrt{N\lambda_2(L)}} = \frac{\|\Pi_\nabla \xi\|_W}{\sqrt{N\lambda_2(L)}}.$$

The arbitrary centering offset does not affect $E_N$ or ranks. Substitution into Eq. (14) proves Eq. (51). If $d_\star$ is an increasing reparameterization of $d_0$, their population mass coordinates agree by Proposition 11.

When $s_1, \dots, s_N$ are independent with law $\mu$ and $d_\star$ is a fixed population field, the oracle values $x_i = d_\star(s_i)$ are independent with CDF $F_{d_\star}$. The Dvoretzky–Kiefer–Wolfowitz–Massart inequality (Massart, 1990) therefore yields

$$\Pr\left\{\Delta_N > \sqrt{\frac{\log(2/\alpha)}{2N}}\right\} \le \alpha. \tag{52}$$

This applies to the oracle empirical CDF; it asserts no empirical-process bound for the fitted depths. If a separate event controls the Hodge error with probability at least $1 - \beta$, a union bound gives both controls with probability at least $1 - \alpha - \beta$, without requiring independence between the events. For fixed spatial designs or dependent sampling, Eq. (14) remains valid with its deterministic $\Delta_N$.

### A.15.1 Transport of the joint spatial-depth representation

Marginal depth distributions discard spatial correspondence. The transport–$L^2$ distance of García Trillos & Slepčev (2016) compares a function together with its spatial measure:

$$d_{TL^2}\big((\mu, f), (\vartheta, g)\big)^2 = \inf_{\pi \in \Gamma(\mu,\vartheta)} \int \big(\|s - t\|^2 + |f(s) - g(t)|^2\big)\, d\pi(s,t), \tag{53}$$

where $\Gamma(\mu, \vartheta)$ denotes spatial couplings. This is the $W_2$ distance between the graph measures $(\mathrm{Id}, f)_\#\mu$ and $(\mathrm{Id}, g)_\#\vartheta$ in the product space. Spatial and depth units are fixed in this definition.

**Corollary 13** (Joint spatial-depth recovery)**.** *Under Theorem 3, suppose $\mu$ has finite second spatial moment, the sample sites are distinct, and the population oracle $d_\star$ is $C_d$-Lipschitz in Euclidean distance. Let $\mu_N = N^{-1}\sum_i \delta_{s_i}$ and $r_\star = F_{d_\star} \circ d_\star$. Then*

$$\begin{aligned} d_{TL^2}\big((\mu_N, \widehat{r}), (\mu, r_\star)\big) \le\ & 2ME_N + \Delta_N \\ & + \sqrt{1 + M^2 C_d^2}\, W_2(\mu_N, \mu). \end{aligned} \tag{54}$$

*If $Q_\nu$ is $K$-Lipschitz, the corresponding calibrated coordinates satisfy*

$$d_{TL^2}\big((\mu_N, Q_\nu \circ \widehat{r}), (\mu, Q_\nu \circ r_\star)\big) \le K(2ME_N + \Delta_N) + \sqrt{1 + K^2 M^2 C_d^2}\, W_2(\mu_N, \mu). \tag{55}$$

**Proof.** Insert $(\mu_N, r_\star|_{\{s_i\}})$ in the triangle inequality. Coupling each sample site to itself bounds the first term by $\big\|\widehat{r} - r_\star|_{\{s_i\}}\big\|_{2,N}$, which is at most $2ME_N + \Delta_N$ by Theorem 3. For the second term, $r_\star$ is $MC_d$-Lipschitz. Any spatial coupling $\pi \in \Gamma(\mu_N, \mu)$ therefore gives

$$\int \big(\|s-t\|^2 + |r_\star(s) - r_\star(t)|^2\big)\, d\pi \le (1 + M^2 C_d^2) \int \|s-t\|^2 \, d\pi.$$

Taking the infimum and then the square root proves Eq. (54). Applying the same argument to $Q_\nu \circ r_\star$, whose Lipschitz constant is at most $KMC_d$, proves Eq. (55). □

This deterministic decomposition separates coordinate error from scalar and spatial sampling discrepancies. It is not a convergence theorem for the implemented graph construction or for neural derivative estimation. A unit gradient alone also does not imply a global Euclidean Lipschitz constant of one on a nonconvex domain; the stated $C_d$ assumption is explicit.

**Depth histograms and the spatial measure.** For a regular $C^1$ depth $d$ and $\mu(ds) = p_\mu(s)\, ds$, the coarea formula (Fleming & Rishel, 1960) gives the scalar density

$$f_d(t) = \int_{d^{-1}(t)} \frac{p_\mu(s)}{\|\nabla d(s)\|}\, d\mathcal{H}^1(s) \quad \text{for almost every } t. \tag{56}$$

Thus a depth histogram combines contour length, spatial sampling or biological mass, and the speed of the chosen coordinate. Under an increasing differentiable reparameterization $h$ with $h' > 0$, its density transforms as $f_{h\circ d}(h(t)) = f_d(t)/h'(t)$. Layer-conditioned histograms can reveal the ordering of reference layers along a learned coordinate, but histogram peaks alone do not identify physical layer thickness or verify the spatial geometry of its contours.

### A.15.2 A POPULATION EXTENSION WITH LEARNED INTEGRATION SPEED

The mass constraint also supports a substantive extension beyond unit-normal integration. Let $n_0 = \nabla d_0 / \|\nabla d_0\|$ and let $w$ be a bounded, strictly positive weight. Consider

$$\min_{\substack{u \in C^1(\Omega),\ a \in C^0(\Omega),\ a>0 \\ u_\# \mu = \nu}} \int_\Omega w(s)\, \|\nabla u(s) - a(s) n_0(s)\|^2 \, d\mu(s). \tag{57}$$

**Proposition 14** (Exact integration with a learned speed)**.** *Under the assumptions of Proposition 11, Eq. (57) has minimum zero and the unique coordinate $u = u_\nu$. Its unique speed is*

$$a_\nu(s) = \frac{f_0(d_0(s))}{f_\nu(u_\nu(s))}\, \|\nabla d_0(s)\|. \tag{58}$$

*No transnormality assumption is required.*

The positive speed supplies the integrating factor needed to preserve the original contours, while the nondegenerate target law fixes the scale and excludes constant collapse. This population extension is not implemented in the reported experiments. Its exact identification result does not establish stability under noisy normals or computational convergence. An empirical implementation must also accommodate atomic spot measures, for example through a discrete target law or a soft transport penalty. Post-hoc calibration of the current Hodge depth preserves its estimated contours and cannot repair contour distortion introduced during integration.

**Proof of Proposition 14.** Write $F_0 = F_{d_0}$. Positive continuous source and target densities imply that $u_\nu = Q_\nu \circ F_0 \circ d_0$ is $C^1$ and

$$\nabla u_\nu(s) = \frac{f_0(d_0(s))}{f_\nu(u_\nu(s))} \nabla d_0(s) = a_\nu(s) n_0(s).$$

The displayed $a_\nu$ is continuous and strictly positive. By Proposition 11, $(u_\nu)_{\#}\mu = \nu$, so this pair is feasible and has objective zero. Since the objective is nonnegative, zero is its minimum.

For any minimizing pair $(u, a)$, positivity of $w$ implies $\nabla u = a n_0$ almost everywhere with respect to $\mu$. The residual is continuous, and $\mu$ has full support, so the equality holds throughout $\Omega$. Every tangent derivative of $u$ along a leaf of $d_0$ is zero. Connected regular leaves imply $u = g \circ d_0$, and local transverse sections show that $g \in C^1(I)$. The gradient identity yields

$$g'(d_0(s)) = \frac{a(s)}{\|\nabla d_0(s)\|} > 0.$$

Thus $g$ is increasing. The constraint $u_{\#}\mu = \nu$ now gives $F_\nu(g(t)) = F_0(t)$, so $g = Q_\nu \circ F_0$. Hence $u = u_\nu$ uniquely, and $a = \|\nabla u_\nu\| = a_\nu$ is also unique. No step requires $\|\nabla d_0\|$ to be constant along a leaf. □

The regularity assumption on the scalar density $f_0$ is explicit because a positive spatial density alone does not guarantee it: the geometry of the domain can change the measure of depth cross sections. The speed proposition establishes population identification with exact normals. A stability theorem for perturbed normals would require additional control of derivatives or integrating factors. In particular, optimizing over $a > 0$ does not by itself provide a uniform coercivity bound.

# B Additional implementation details

The main framework distinguishes Hodge Depth and Source–Eikonal depth. This appendix describes implementations and selection rules in numerical realizations.

## B.1 Spatial graph and local regression

The inputs are locations $s_i \in \mathbb{R}^2$ and feature vectors $z_i \in \mathbb{R}^p$, as in Section 3.1. We construct a symmetrized nearest-neighbor graph $G = (V, E)$ from physical coordinates and retain its largest connected component. In the v1.3 configuration, the integration graph has six nearest neighbors before symmetrization. Neighbor ties are resolved deterministically by observation identifier. Coordinates are centered and divided by the median nearest-neighbor spacing; below, $s_i$ denotes these rescaled coordinates. For $\ell_{ij} = \|s_j - s_i\|_2$, the spatial graph weights are $w_{ij} = \exp[-\ell_{ij}^2/(2\bar{\ell}^2)]$, where $\bar{\ell}$ is the median graph-edge length.

Local feature variation is estimated using a separate, larger neighborhood $\mathcal{N}_k(i)$. The v1.3 setting uses $k = 48$ neighbors and the first $p = 14$ standardized expression PCs. With $x_{ij} = s_j - s_i$ and $h_i$ the distance to the $k$th neighbor, define

$$K_{ij} = \exp\left[-\frac{1}{2}\left(\frac{\|x_{ij}\|_2}{0.65 h_i}\right)^2\right].$$

The implemented local fit uses differences from the target observation:

$$\widehat{J}_i = \operatorname*{arg\,min}_{J \in \mathbb{R}^{p\times 2}} \sum_{j \in \mathcal{N}_k(i)} K_{ij} \, \|(z_j - z_i) - J x_{ij}\|_2^2 + \rho_i \, \|J\|_F^2 . \tag{59}$$

Here $\rho_i > 0$ is a small ridge parameter scaled to the local spatial second moment. The target value $z_i$ supplies the intercept; no additional intercept is estimated. Thus $\widehat{J}_i$ estimates the $p \times 2$ spatial Jacobian in Eq. (3).

## B.2 Implemented eigen-contrast and support weights

**Geometric interpretation.** The matrix $J_z^\top J_z$ describes feature change along a spatial displacement $v$ through $v^\top J_z^\top J_z v = \|J_z v\|^2$. Under the regular single-index model it vanishes along layer tangents and has one nonzero normal direction. It is therefore a rank-one, degenerate pullback of feature-space geometry.

Let $\widehat{\sigma}_{i1} \geq \widehat{\sigma}_{i2}$ be the singular values of $\widehat{J}_i$. Its leading right-singular vector $\tilde{n}_i \in \mathbb{R}^2$ estimates an unoriented normal axis, represented by $\widetilde{P}_i = \tilde{n}_i \tilde{n}_i^\top$. The raw reliability is

$$r_i = \mathrm{clip}_{[10^{-3},1]}\left(\frac{\widehat{\sigma}_{i1}^2 - \widehat{\sigma}_{i2}^2}{\max\{\widehat{\sigma}_{i1}^2 + \widehat{\sigma}_{i2}^2, 10^{-12}\}}\right). \tag{60}$$

This score measures directional anisotropy. It does not, by itself, distinguish a weak directional signal from anisotropic noise. In particular, when $m'(d_0) = 0$ throughout a neighborhood, local expression does not identify the normal direction. Spatial regularization can interpolate information from surrounding locations, but the local identifiability argument requires nonzero feature change.

**Neighborhood-asymmetry correction.** E1 uses $c_i = r_i$ as node confidence. E2 additionally measures whether the neighborhood surrounds the target symmetrically:

$$A_i = \frac{\left\|\sum_j K_{ij} x_{ij}\right\|_2}{\sum_j K_{ij} \left\|x_{ij}\right\|_2}, \quad c_i = r_i(1 - A_i)^2. \tag{61}$$

The sums are over $\mathcal{N}_k(i)$, with a numerical denominator floor. By the triangle inequality, $0 \leq A_i \leq 1$. A one-sided neighborhood receives less influence in smoothing, synchronization, and integration. This correction uses sampling geometry and does not identify anatomical layer boundaries.

### B.3 Support-harmonic refinement (E3)

E3 further identifies locations at which local estimation relies on extrapolation. Write $\pi_{ij} = K_{ij} / \sum_j K_{ij}$,

$$\mu_i = \sum_j \pi_{ij} x_{ij}, \quad C_i = \sum_j \pi_{ij}(x_{ij} - \mu_i)(x_{ij} - \mu_i)^\top,$$

and define $\xi_i = \mu_i^\top (C_i + \rho_i^S I)^{-1} \mu_i$, with a small positive ridge $\rho_i^S$. Locations with $\xi_i \geq 1$ seed an unsupported set $\mathcal{U}$, expanded by one hop along immediate spatial edges. E3 sets the projector-fidelity weight $c_i^P$ to zero on $\mathcal{U}$ and to $c_i$ elsewhere. E1 and E2 use $c_i^P = c_i$ everywhere. Subsequent synchronization and integration retain $c_i$. E3 therefore reconstructs local directions through spatial support; it does not force neighboring depths to be equal. We retain it as a boundary sensitivity.

### B.4 Projector smoothing and orientation

The vectors $\tilde{n}_i$ and $-\tilde{n}_i$ describe the same axis, so vector averaging can cancel compatible local estimates. We instead smooth their projectors:

$$\min_{Q_1,\ldots,Q_n} \sum_i c_i^P \left\|Q_i - \widetilde{P}_i\right\|_F^2 + \lambda_P \sum_{(i,j)\in E} w_{ij} \left\|Q_i - Q_j\right\|_F^2 + \epsilon_P \sum_i \left\|Q_i\right\|_F^2. \tag{62}$$

Each $Q_i$ is a symmetric $2 \times 2$ matrix. The small stabilizer is $\epsilon_P = 10^{-8}$ in the implementation. This convex quadratic problem is solved componentwise using the graph Laplacian. At masked E3 locations, the solution is approximately a weighted harmonic extension from supported neighbors. We then take a leading unit eigenvector $u_i$ of $Q_i$ and reproject to $\widehat{P}_i = u_i u_i^\top$.

To orient these axes, let $g_{ij} = \mathrm{sign}(u_i^\top u_j)$, taking $g_{ij} = 1$ when the inner product is zero. The implemented edge confidence is

$$a_{ij} = w_{ij}\big(0.05 + 0.95\sqrt{c_i c_j}\big) \times \big(0.1 + 0.9|u_i^\top u_j|\big). \tag{63}$$

Its positive floors preserve weak constraints at low-confidence locations. We seek signs using the objective in Equation (7). For the spectral relaxation, let $\mathcal{A}_{ij} = a_{ij} g_{ij}$ on graph edges and $D_{ii} = \sum_j a_{ij}$. We round the leading eigenvector of $D^{-1/2}\mathcal{A}D^{-1/2}$ by its signs, equivalently using the smallest eigenvector of the normalized connection Laplacian. The oriented normals are $n_i = \sigma_i u_i$. One overall sign remains undetermined.

### B.5 HODGE SOLVE, RESIDUAL, AND NORMALIZATION

For connected $G$ and positive edge weights, $L = B^\top W B$ has only the constant null direction. With $L^\dagger$ denoting its Moore–Penrose pseudoinverse,

$$\widehat{d} = L^\dagger B^\top W \omega. \tag{64}$$

The fitted increments $B\widehat{d}$ form the weighted orthogonal projection of $\omega$ onto $\mathrm{im}(B)$, the edge flows generated by scalar node potentials. The normalized residual

$$R_H = \frac{\left\| W^{1/2}(\omega - B\widehat{d}) \right\|_2^2}{\left\| W^{1/2}\omega \right\|_2^2} \tag{65}$$

quantifies the remaining nonintegrable component, with a denominator floor in numerical evaluation. Residuals computed with different confidence weights are not directly comparable measures of anatomical accuracy.

E1–E3 return the normalized coordinate

$$\phi_i = \frac{\widehat{d}_i - \min_j \widehat{d}_j}{\max_j \widehat{d}_j - \min_j \widehat{d}_j}, \tag{66}$$

with E2 and E3 resolving the global sign by a coordinate-based convention that uses no annotations. E1 retains its numerical solver sign. Neither convention assigns an anatomical interpretation to the endpoints. A zero-range solution is degenerate. The term *signed Hodge* refers to its oriented increments, even though the exported coordinate lies in $[0, 1]$.

### B.6 SOURCE INFERENCE AND LABEL-FREE SELECTION

**Source–Eikonal coordinate.** A signed coordinate orders locations across the tissue. In recurrent layered organization, similar expression profiles may instead occur on both sides of an internal source sheet. Source–Eikonal depth represents distance from that sheet, allowing the two sides to share a depth. This changes the global coordinate interpretation rather than repairing a local boundary estimate.

We scan centers $t \in \{0.15, 0.16, \ldots, 0.85\}$ along the support-corrected Hodge coordinate $\phi$. At equal offsets on either side of $t$, Gaussian windows along $\phi$ give mean expression profiles $\mu_r^-(t)$ and $\mu_r^+(t)$. We choose the eligible center minimizing the reflection ratio

$$M(t) = \frac{\sum_r v_r \left\| \mu_r^+ - \mu_r^- \right\|_2^2}{\frac{1}{2}\sum_r v_r \left( \left\| \mu_r^+ - \bar{\mu} \right\|_2^2 + \left\| \mu_r^- - \bar{\mu} \right\|_2^2 \right)}, \tag{67}$$

where $v_r$ is the smaller effective sample size of the paired windows and $\bar{\mu}$ is their weighted pooled mean. The primary profile uses 14 standardized PCs. Eligibility requires at least 15% of spots on each side, eight supported offset pairs, and effective sample size at least 20 per window. These requirements prevent selection from inadequately sampled sides of the tissue.

For the selected center $t_\star$, define the source band

$$\mathcal{S} = \{i : |\phi_i - t_\star| \le h_S\}, \;\; h_S = \mathrm{median}_{(i,j)\in E} |\phi_j - \phi_i|.$$

If this band is empty, the implementation uses the $\max(10, \lfloor n/100 \rfloor)$ spots closest to $t_\star$. Multi-source Dijkstra propagation computes the distances and normalization in Eq. (10), where each graph-edge cost is its physical length divided by the median edge length. Thus $\psi$ is an isotropic graph-geodesic distance from the inferred source band. The calculation applies to the entire analyzed component. The source candidate is available for comparison even when the gate retains Hodge Depth.

**Label-free endpoint gate.** The gate returns $d_i^{\mathrm{sel}} = \psi_i$ only when all four prespecified conditions hold: (i) the primary reflection ratio is below one; (ii) the same inequality holds at 4, 8, 14, and 20 PCs, with the center reselected at each scale; (iii) piecewise-linear expression reconstruction with a fixed $K_{\mathrm{ref}}$ intervals has lower squared error along $\psi$ than along $\phi$; and (iv) five-fold spline-regression prediction favors the source coordinate by a fixed evidence margin. Otherwise it returns $d_i^{\mathrm{sel}} = \phi_i$.

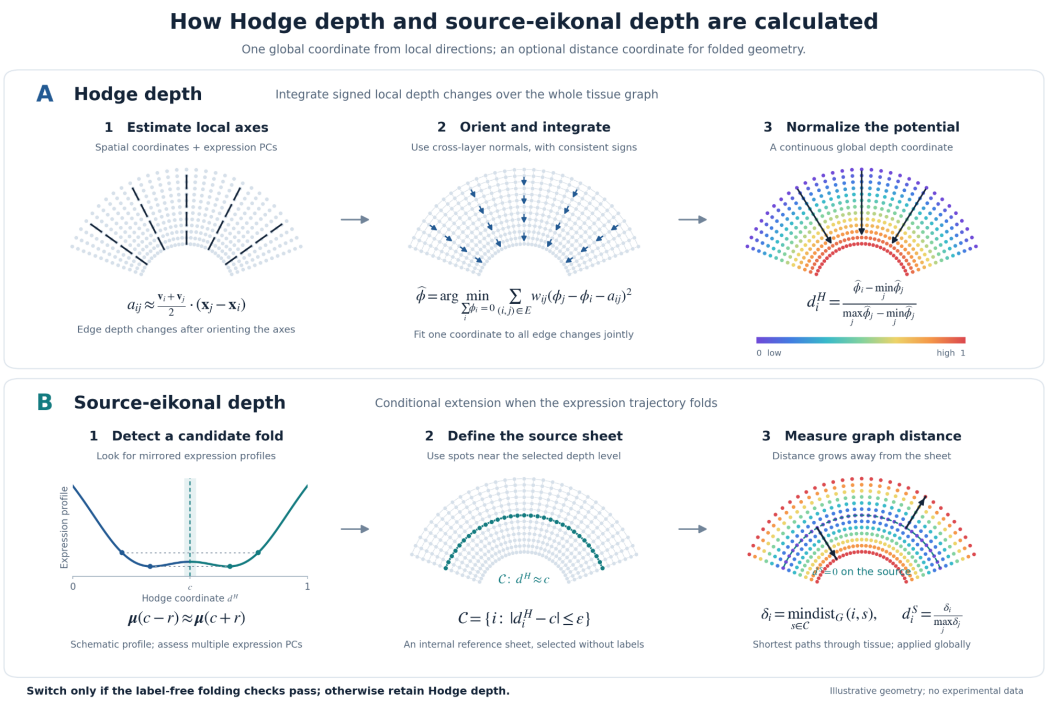


Figure 4: Calculation of Hodge Depth and Source–Eikonal depth. Local normal axes are oriented consistently and integrated into a global potential. When the label-free gate supports recurrent molecular organization, graph distance from an inferred internal source sheet supplies the complementary unsigned coordinate.

For the fourth condition, source-center selection and expression regression are repeated on each training fold. If $\Delta_f$ is the signed-coordinate prediction error minus the source-coordinate error on fold $f$, the rule is

$$\bar{\Delta} - 2.1318467863 \, \frac{s_\Delta}{\sqrt{5}} > 0.$$

The globally fitted PCs and Hodge coordinate remain fixed. This is a conditional prediction heuristic, not a calibrated confidence bound or independent validation of the full encoder. The modular implementation uses $K_{\mathrm{ref}} = 7$, independently of the downstream cluster count $K$. If no eligible source center exists, source analysis is unavailable and the implementation reports this failure; this differs from a valid candidate failing the selection conditions. Neither the source-depth construction nor the gate refits the local normal field or the Hodge potential.

### B.7 Neural gradient estimation and its affine specialization

The general construction uses differentiable multivariate fields $f_{\theta_i} : \Omega \to \mathbb{R}^p$. A parameter-sharing set $\Theta_{\mathrm{share}}$ specifies whether all locations use one network, separate local networks, or partially shared parameters. A common estimation objective is

$$(\widehat{\theta}_i)_i = \operatorname*{arg\,min}_{(\theta_i)_i \in \Theta_{\mathrm{share}}} \sum_i \sum_{j \in \mathcal{N}_k(i)} K_{ij} \, \| f_{\theta_i}(s_j) - z_j \|_2^2 + \lambda_\nabla \mathcal{R}_\nabla((f_{\theta_i})_i), \tag{68}$$

with $\widehat{J}_i = D f_{\widehat{\theta}_i}(s_i)$. The reconstruction weights may use the spatial kernel in Appendix B.1. The regularizer controls derivatives or model complexity. For twice-differentiable neural fields, an integrated squared Hessian is one possible derivative-variation penalty. With complete parameter sharing, the objective fits one global feature field; automatic differentiation supplies its $p \times 2$ Jacobian at each location. No anatomical labels or scalar depth bottleneck enter this objective.

For the evaluated specialization, use separate affine networks $f_{\theta_i}(s) = z_i + A_i(s - s_i)$, fix the intercept at the observed $z_i$, and choose $\lambda_\nabla \mathcal{R}_\nabla = \sum_i \rho_i \|A_i\|_F^2$. The objective then separates over locations and is exactly the weighted ridge estimator in Eq. (59), with $\widehat{J}_i = \widehat{A}_i$. The local slopes are allowed to differ, so this specialization does not impose a single globally linear expression field. A Hessian-only penalty would vanish for affine fields; the explicit ridge penalty is therefore needed to recover the implemented estimator.

Both parameterizations feed the same eigen-contrast, projector refinement, synchronization, and Hodge solve. Although every differentiable feature output has a gradient, the normalized leading axis of the multivariate Jacobian need not be integrable. The geometric stages consequently remain relevant with a shared neural field. The perturbation bounds apply conditional on their stated derivative and orientation error assumptions; accurate feature reconstruction alone does not establish derivative accuracy. All reported HodgeDepth experiments use the local affine specialization.

### B.8 Feature construction and downstream decoding

The DLPFC fixed pipeline combines HVG-scaled PCA20 with the gate-selected coordinate in a tied-covariance Gaussian mixture, followed by a 50-neighbor vote. The synthetic study uses a common one-dimensional decoder, while the Batiuk study uses the encoder–decoder combinations in Appendix E.

The v1.3 geometric settings use $\lambda_P = 0.01$. The original E1 implementation retains its earlier unsupervised tuning over neighborhood size, feature dimension, and smoothing. Diameter contrasts, multiscale neighborhoods, and noise-calibrated confidence are further exploratory refinements. Subsequent depth fusion produces a separate layer score: geometric depth can vary within a layer even when expected expression is constant.

### B.9 Depth visualization and endpoint assignments

Figure 5 extends the compact main-text figure with reference maps and both depth candidates for the same four sections. The gate returns Source–Eikonal depth for 151509 and Hodge Depth for 151669, 151671, and 151673. Its four fixed criteria are specified in Appendix B.6; no reference-layer label or downstream ARI enters these endpoint decisions. The unselected candidate panels are diagnostic comparisons, not additional selected outputs.

Reference maps show measured spot locations. Colors are shared with the histograms, which condition on the reference layer and normalize each layer separately to unit area. Legend counts describe the reference annotations, not predicted cluster sizes. Gray, unannotated spots appear only in the spatial maps. Histograms use 30 common bins over $[0, 1]$, with common vertical limits across methods within a section. The supplied signed coordinates and source-eikonal scale and orientation are retained; reference labels are not used to refit the displayed coordinates.

## C DLPFC benchmark: configurations and complete results

**Benchmark software.** We benchmarked nine spatial domain identification methods using the authors' publicly released software implementations. The Python packages included GraphST (v1.1.1), STAGATE_pyG (v1.0.0), SEDR (v1.0.0), SpaGCN (v1.2.7), SpaceFlow (v1.0.3), and GASTON (gaston-spatial v0.0.2), obtained from their official GitHub repositories. STAGATE was evaluated using the authors' PyTorch Geometric implementation. The R packages included SpatialPCA (v1.3.0) and BASS (v1.1.0.017), obtained from the authors' official repositories, and BayesSpace (v1.20.2), obtained from Bioconductor.

### C.1 Evaluation scope and configuration families

The benchmark compares predicted spatial domains with manual annotations on 12 DLPFC Visium sections. For partition-agreement evaluation, layer labels are treated as nominal categories. ARI, NMI, adjusted mutual information (AMI), and the other partition-agreement measures do not require an ordering of layer names. One-to-one matched accuracy aligns predicted and reference category

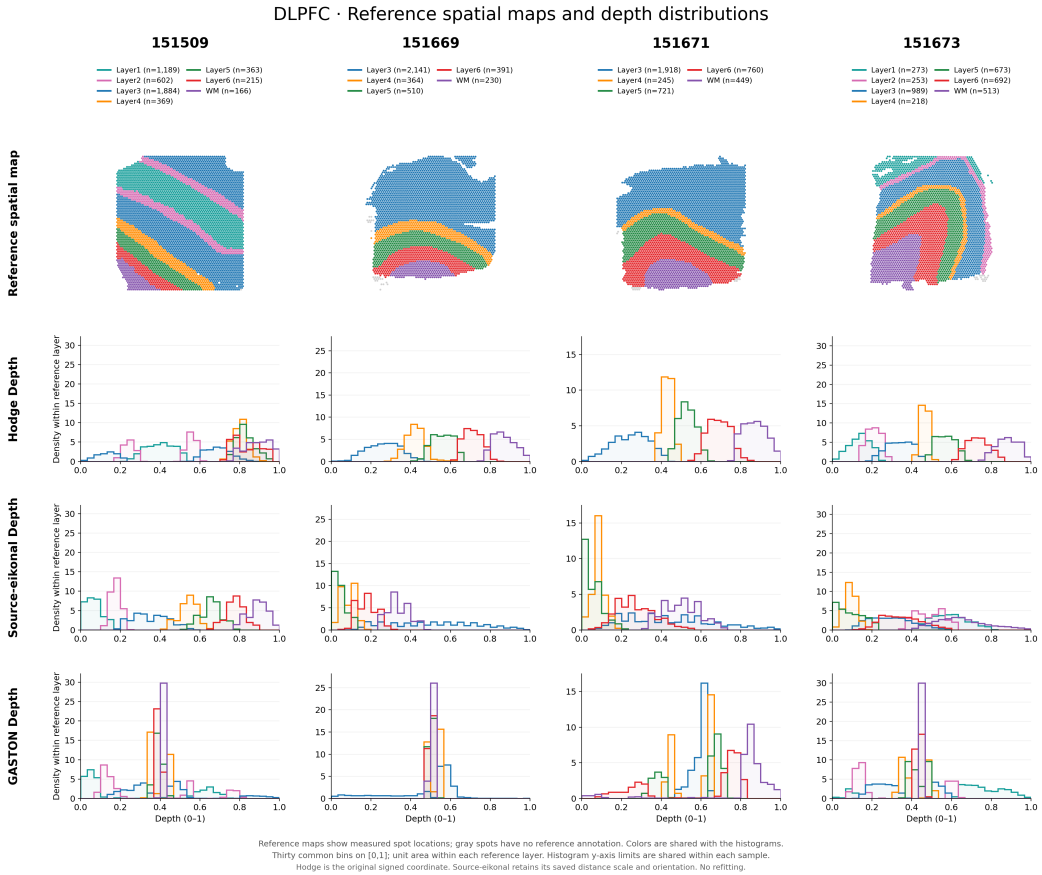


Figure 5: Reference anatomy and depth distributions in four DLPFC sections. Rows show reference spatial maps, signed Hodge depth, source-eikonal depth, and GASTON isodepth. Colors identify reference layers; gray map spots lack annotations. Histograms use 30 common bins on $[0, 1]$, unit area within each layer, and shared $y$-axis limits within each section. **The label-free gate selects the source-eikonal row for 151509 and the signed-Hodge row for the other three sections.** Both candidate rows are shown for comparison. Reference labels do not enter the gate. Source-eikonal values retain their saved scale and orientation without refitting.

IDs for evaluation only. The separate coordinate-fidelity analysis in Section 5.1 uses the anatomical order L1 through L6 to WM. Means and medians summarize sections with equal weight.

**HodgeDepth configurations.** The encoder-and-decoder configuration considers the candidate family

$$\{\mathrm{E2}, \mathrm{E5}\} \times \{\text{joint GMM}\}.$$

E2 supplies asymmetry-corrected signed Hodge depth, whereas E5 supplies source-eikonal distance. The joint mixture heads combine expression PCs and depth. All partition-agreement and spatial-coherence metrics are evaluated on that same partition; the ordering metrics use the coordinate directly. The selection between E2 and E5 is data-driven and label-free. The pipeline uses the gate-selected encoder, a tied-covariance Gaussian mixture on HVG-scaled expression PCA20 together with depth, and a 50-nearest-neighbor vote.

**Spot sets and implementation records.** HodgeDepth uses all analyzed spots, with 12 spatially disconnected spots excluded in section 151673. STAGATE denotes the PyTorch Geometric implementation, and SpaGCN denotes the variant without histological image features.

Table 3: DLPFC benchmark across 12 sections. Values are section means. ARI rank is the mean rank among these ten methods; First counts sections with the highest ARI. Bold marks the best value in each column. Appendix C reports the full results.

| Method | ARI ↑ | NMI ↑ | Acc. ↑ | ARI rank ↓ | First |
|---|---|---|---|---|---|
| **HodgeDepth** | **0.618** | **0.695** | **0.737** | **1.75** | **6** |
| GraphST | 0.523 | 0.657 | 0.678 | 3.92 | 1 |
| SEDR | 0.533 | 0.663 | 0.683 | 4.08 | 1 |
| STAGATE (PyG) | 0.509 | 0.642 | 0.679 | 4.33 | 1 |
| BASS | 0.515 | 0.650 | 0.671 | 4.58 | 1 |
| SpatialPCA | 0.503 | 0.640 | 0.664 | 5.50 | 0 |
| BayesSpace | 0.469 | 0.608 | 0.617 | 6.50 | 1 |
| SpaGCN (no HE) | 0.434 | 0.567 | 0.603 | 7.00 | 0 |
| GASTON | 0.420 | 0.592 | 0.575 | 7.50 | 1 |
| SpaceFlow | 0.305 | 0.458 | 0.463 | 9.83 | 0 |

Table 4: Agreement with manual layers. Entries are section means with mean ranks in parentheses among the ten methods. Bold marks the best mean and the lowest mean rank separately, including ties at the displayed precision.

| Method | ARI ↑ | NMI ↑ | AMI ↑ | Acc. ↑ | FMI ↑ | Homog. ↑ | Compl. ↑ | Purity ↑ |
|---|---|---|---|---|---|---|---|---|
| HD | **0.618** (**1.75**) | **0.695** (**1.92**) | **0.694** (**1.83**) | **0.737** (**2.00**) | **0.711** (**1.58**) | **0.692** (**3.00**) | **0.701** (**1.17**) | **0.799** (**2.67**) |
| GraphST | 0.523 (3.92) | 0.657 (3.46) | 0.656 (3.46) | 0.678 (4.00) | 0.636 (3.92) | 0.668 (3.75) | 0.647 (3.25) | 0.780 (3.75) |
| SEDR | 0.533 (4.08) | 0.663 (3.54) | 0.663 (3.50) | 0.683 (4.00) | 0.643 (3.83) | 0.675 (4.88) | 0.655 (3.50) | 0.780 (4.08) |
| STAGATE | 0.509 (4.33) | 0.642 (4.46) | 0.641 (4.50) | 0.679 (4.25) | 0.621 (4.62) | 0.669 (3.79) | 0.618 (4.79) | 0.772 (4.79) |
| BASS | 0.515 (4.58) | 0.650 (4.79) | 0.649 (4.75) | 0.671 (4.62) | 0.625 (4.58) | 0.681 (3.79) | 0.623 (5.33) | 0.794 (3.71) |
| SpatialPCA | 0.503 (5.50) | 0.640 (5.21) | 0.639 (5.21) | 0.664 (5.25) | 0.619 (5.33) | 0.655 (5.04) | 0.626 (5.21) | 0.766 (5.75) |
| BayesSpace | 0.469 (6.50) | 0.608 (6.71) | 0.607 (6.75) | 0.617 (6.79) | 0.587 (6.75) | 0.623 (6.75) | 0.595 (6.67) | 0.728 (7.25) |
| SpaGCN | 0.434 (7.00) | 0.567 (7.96) | 0.566 (7.96) | 0.603 (7.00) | 0.561 (6.96) | 0.586 (7.79) | 0.552 (7.75) | 0.725 (7.00) |
| GASTON | 0.420 (7.50) | 0.592 (6.96) | 0.591 (7.04) | 0.575 (7.42) | 0.552 (7.58) | 0.616 (6.62) | 0.571 (7.33) | 0.728 (6.33) |
| SpaceFlow | 0.305 (9.83) | 0.458 (10.00) | 0.457 (10.00) | 0.463 (9.67) | 0.450 (9.83) | 0.495 (9.58) | 0.435 (10.00) | 0.649 (9.67) |

Tables 4 and 5 summarize the remaining comparisons. Higher values indicate greater agreement for ARI, NMI, AMI, matched accuracy, homogeneity, completeness, purity, and Fowlkes–Mallows index (FMI). Following the benchmark convention, lower spatial chaos (CHAOS), abnormal-spot fraction (PAS), and local inverse Simpson index (LISI) values indicate greater spatial coherence. These spatial and expression statistics do not use reference labels in their definitions.

Aggregate means, medians, ranks, and win counts for the ten-method comparison follow the supplied summary table. Section values are retained at their reported precision; near ties need not reproduce the aggregate ranks after rounding.

### C.2 Continuous depth evaluation

For the continuous depth evaluation results reported in Table 2, each coordinate is oriented once per section by the sign of its Spearman correlation, for evaluation only. Adjacent-pair AUC is $P(d_{\ell+1} > d_\ell)$ for adjacent anatomical layers (0.5 denotes no separation). Its mean and minimum are computed over evaluable pairs within each section and then averaged across sections. *Pairs ordered* reports the total correctly ordered pairs out of 64 evaluable section–layer-pair combinations; all other entries are equal-weight section means. Monotone $R^2$ is obtained by isotonic regression of the ordinal layer index on the coordinate. Higher is better throughout;

### C.3 Complete Br5595 domain-recovery comparison

Figures 6 and 7 reproduce every section and method from the supplied Br5595 comparison, with the reference column repeated for readability. The main-text figure shows 151670 from Br5595 and 151674 from Br8100, comparing GraphST (spatial coherence), SEDR (strongest mean ARI), and GASTON (continuous depth). The additional maps below cover all four Br5595 sections; complete

Table 5: Spatial statistics. Entries are section means with mean ranks in parentheses among the ten methods. Arrows indicate the preferred direction in the benchmark. Bold marks the best mean and the lowest mean rank separately, including ties at the displayed precision.

| Method | CHAOS ↓ | PAS ↓ | LISI mean ↓ |
|---|---|---|---|
| HD | **0.0582** (**2.17**) | **0.0110** (**1.67**) | **1.225** (2.58) |
| GraphST | **0.0582** (2.67) | 0.0125 (2.17) | **1.225** (**2.42**) |
| SEDR | 0.0583 (3.46) | 0.0197 (4.92) | 1.274 (6.00) |
| STAGATE | 0.0589 (6.96) | 0.0337 (7.17) | 1.298 (8.00) |
| BASS | 0.0594 (8.62) | 0.0381 (8.17) | 1.310 (8.17) |
| SpatialPCA | 0.0583 (3.88) | 0.0205 (4.71) | 1.263 (5.54) |
| BayesSpace | 0.0589 (7.38) | 0.0334 (7.50) | 1.259 (5.17) |
| SpaGCN | 0.0595 (8.83) | 0.0718 (9.92) | 1.348 (9.42) |
| GASTON | 0.0584 (4.38) | 0.0172 (3.33) | 1.236 (3.67) |
| SpaceFlow | 0.0588 (6.67) | 0.0266 (5.46) | 1.262 (4.04) |

quantitative results cover all 12 sections. The panels are typeset from unchanged regions of the supplied image. No predictions, colors, or ARIs are recomputed. Predicted cluster IDs are arbitrary; matching colors across methods need not denote the same anatomical layer. STAGATE uses PyG, SpaGCN omits histology, and BASS is fitted jointly within the donor, as in the original figure. The original HodgeDepth heading "selected" denotes the reported label-free configuration, not selection by reference-label ARI.

### C.4 Section-level results

The following tables report section-level results for the ten-method comparison. HD denotes the HodgeDepth. The remaining abbreviations are GS (GraphST), SE (SEDR), ST (STAGATE with PyTorch Geometric), BA (BASS), SP (SpatialPCA), BS (BayesSpace), SG (SpaGCN without histology), GA (GASTON), and SF (SpaceFlow). Bold marks the best values at the reported precision, including ties, according to the direction indicated for each metric.

### C.5 Depth-encoder ablations on DLPFC

We examine coordinate source, sign assignment, and increment normalization on the 12 DLPFC sections. Encoder variants use the frozen configuration and cached preprocessing, modifying either sign assignment or increment magnitude while holding the remaining steps fixed. Graph-based reference coordinates use the encoder's weighted 6-NN graph; spatial PC1 is computed directly from spot coordinates. We evaluate ordering against L1$<\cdots<$L6$<$WM on annotated spots. Every coordinate uses the same tied-covariance GMM on 20 PCs of HVG-scaled expression (divided by $\sqrt{20}$) and standardized depth, with annotation-derived $K$ and a 50-NN vote.

**Coordinate source.** Originial attains mean Spearman $|\rho| = 0.885$ and monotone $R^2 = 0.883$, compared with $0.633$–$0.663$ and $0.458$–$0.548$ for spatial PC1, Fiedler, and diffusion-map coordinates (Table 17). Boundary distance performs less well ($|\rho| = 0.222$). E2 exceeds each tested geometric coordinate in Spearman correlation on all 12 sections.

**Sign assignment.** Replacing spectral synchronization with independent random signs reduces mean $|\rho|$ to $0.285$. Greedy propagation reaches $0.790$; its deficit is concentrated in sections 151508 ($0.155$ versus $0.901$ for HD) and 151676 ($0.566$ versus $0.977$). These observations support coherent orientation and robustness to the observed greedy failures; the aggregate spectral–greedy difference is inconclusive (two-sided Wilcoxon $p = 0.176$).

**Increment normalization.** Scaling synchronized directions by local Jacobian magnitude yields mean $|\rho| = 0.873$ and monotone $R^2 = 0.888$, versus $0.885$ and $0.883$ for unit-direction increments. E2 improves Spearman correlation on 10 of 12 sections, while fixed-head ARI is similar ($0.546$ versus $0.544$). This supports a modest benefit in average ordering fidelity rather than uniform improvement across outcomes.

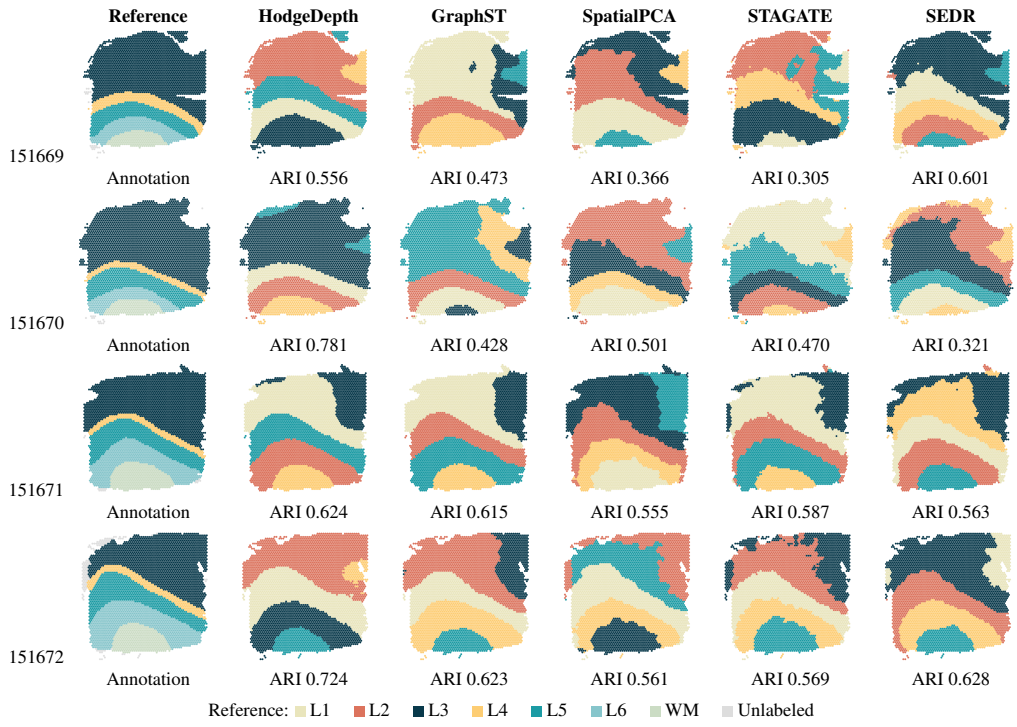


Figure 6: Br5595 domain recovery, part 1: reference annotations, HodgeDepth, GraphST, SpatialPCA, STAGATE (PyG), and SEDR. All four donor sections are shown.

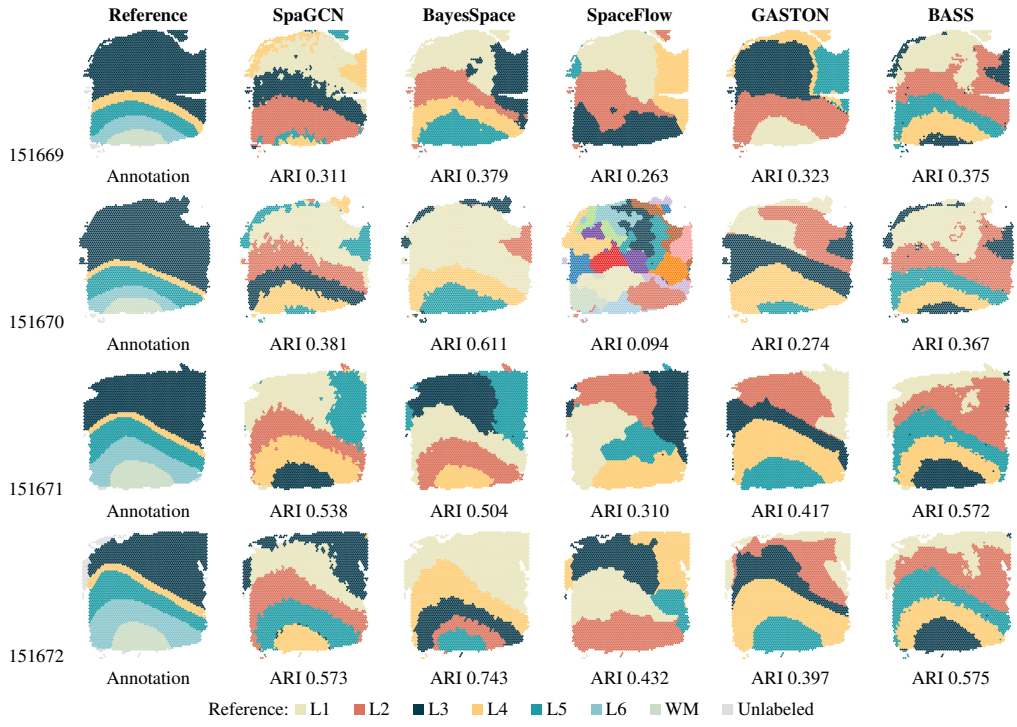


Figure 7: Br5595 domain recovery, part 2: reference annotations, SpaGCN without histology, BayesSpace, SpaceFlow, GASTON, and BASS fitted jointly within donor. All four donor sections are shown.

Table 6: Adjusted Rand index (ARI). Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | 0.568 | 0.437 | 0.491 | **0.588** | 0.561 | 0.544 | 0.439 | 0.441 | 0.502 | 0.405 |
| 151508 | 0.520 | 0.479 | 0.497 | 0.526 | **0.567** | 0.503 | 0.419 | 0.418 | 0.437 | 0.317 |
| 151509 | **0.602** | 0.497 | 0.492 | 0.492 | 0.463 | 0.457 | 0.351 | 0.470 | 0.454 | 0.205 |
| 151510 | **0.566** | 0.485 | 0.559 | 0.466 | 0.478 | 0.428 | 0.458 | 0.462 | 0.354 | 0.385 |
| 151669 | 0.556 | 0.473 | **0.601** | 0.305 | 0.375 | 0.366 | 0.379 | 0.311 | 0.323 | 0.263 |
| 151670 | **0.781** | 0.428 | 0.321 | 0.470 | 0.367 | 0.501 | 0.611 | 0.381 | 0.274 | 0.094 |
| 151671 | **0.624** | 0.615 | 0.563 | 0.587 | 0.572 | 0.555 | 0.504 | 0.538 | 0.417 | 0.310 |
| 151672 | 0.724 | 0.623 | 0.628 | 0.569 | 0.575 | 0.561 | **0.743** | 0.573 | 0.397 | 0.432 |
| 151673 | 0.666 | 0.648 | 0.574 | 0.604 | 0.553 | 0.536 | 0.554 | 0.383 | **0.673** | 0.369 |
| 151674 | **0.653** | 0.428 | 0.606 | 0.381 | 0.569 | 0.509 | 0.368 | 0.454 | 0.361 | 0.264 |
| 151675 | 0.565 | **0.574** | 0.573 | 0.573 | 0.545 | 0.492 | 0.375 | 0.467 | 0.401 | 0.318 |
| 151676 | **0.593** | 0.584 | 0.492 | 0.551 | 0.549 | 0.589 | 0.422 | 0.308 | 0.442 | 0.301 |
| Mean | **0.618** | 0.523 | 0.533 | 0.509 | 0.515 | 0.503 | 0.469 | 0.434 | 0.420 | 0.305 |
| Median | **0.597** | 0.491 | 0.561 | 0.538 | 0.551 | 0.506 | 0.430 | 0.448 | 0.409 | 0.314 |

Table 7: Normalized mutual information (NMI). Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | 0.656 | 0.653 | 0.664 | **0.711** | 0.675 | 0.652 | 0.635 | 0.580 | 0.665 | 0.539 |
| 151508 | 0.625 | 0.640 | 0.655 | 0.657 | **0.669** | 0.619 | 0.567 | 0.541 | 0.603 | 0.464 |
| 151509 | **0.663** | 0.659 | 0.658 | 0.655 | 0.628 | 0.629 | 0.556 | 0.629 | 0.628 | 0.413 |
| 151510 | **0.678** | 0.640 | 0.670 | 0.634 | 0.636 | 0.603 | 0.572 | 0.587 | 0.589 | 0.515 |
| 151669 | **0.659** | 0.591 | 0.637 | 0.510 | 0.556 | 0.576 | 0.590 | 0.481 | 0.519 | 0.394 |
| 151670 | **0.710** | 0.558 | 0.517 | 0.569 | 0.539 | 0.593 | 0.612 | 0.508 | 0.515 | 0.372 |
| 151671 | **0.736** | 0.722 | 0.692 | 0.692 | 0.665 | 0.657 | 0.636 | 0.654 | 0.618 | 0.424 |
| 151672 | **0.734** | 0.719 | 0.693 | 0.687 | 0.676 | 0.682 | 0.717 | 0.681 | 0.590 | 0.562 |
| 151673 | 0.721 | **0.740** | 0.705 | 0.717 | 0.693 | 0.702 | 0.693 | 0.528 | 0.721 | 0.506 |
| 151674 | **0.749** | 0.613 | 0.723 | 0.500 | 0.703 | 0.628 | 0.521 | 0.551 | 0.533 | 0.393 |
| 151675 | 0.691 | 0.657 | **0.692** | 0.684 | 0.675 | 0.628 | 0.600 | 0.584 | 0.532 | 0.476 |
| 151676 | **0.718** | 0.687 | 0.655 | 0.685 | 0.687 | 0.705 | 0.592 | 0.478 | 0.587 | 0.435 |
| Mean | **0.695** | 0.657 | 0.663 | 0.642 | 0.650 | 0.640 | 0.608 | 0.567 | 0.592 | 0.458 |
| Median | **0.701** | 0.655 | 0.667 | 0.671 | 0.672 | 0.629 | 0.596 | 0.566 | 0.589 | 0.450 |

Table 8: Adjusted mutual information (AMI). Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | 0.655 | 0.652 | 0.663 | **0.711** | 0.674 | 0.651 | 0.634 | 0.579 | 0.664 | 0.538 |
| 151508 | 0.624 | 0.639 | 0.654 | 0.656 | **0.668** | 0.619 | 0.566 | 0.540 | 0.602 | 0.463 |
| 151509 | **0.663** | 0.659 | 0.657 | 0.655 | 0.627 | 0.628 | 0.555 | 0.628 | 0.627 | 0.411 |
| 151510 | **0.677** | 0.639 | 0.669 | 0.633 | 0.635 | 0.602 | 0.571 | 0.586 | 0.588 | 0.514 |
| 151669 | **0.658** | 0.590 | 0.636 | 0.509 | 0.555 | 0.575 | 0.589 | 0.480 | 0.518 | 0.393 |
| 151670 | **0.709** | 0.558 | 0.516 | 0.569 | 0.538 | 0.592 | 0.611 | 0.507 | 0.514 | 0.369 |
| 151671 | **0.736** | 0.722 | 0.692 | 0.691 | 0.665 | 0.656 | 0.635 | 0.654 | 0.618 | 0.423 |
| 151672 | **0.733** | 0.718 | 0.693 | 0.686 | 0.676 | 0.682 | 0.717 | 0.681 | 0.590 | 0.562 |
| 151673 | 0.720 | **0.739** | 0.704 | 0.716 | 0.693 | 0.701 | 0.692 | 0.527 | 0.720 | 0.505 |
| 151674 | **0.749** | 0.612 | 0.723 | 0.499 | 0.702 | 0.627 | 0.520 | 0.549 | 0.532 | 0.392 |
| 151675 | 0.690 | 0.656 | **0.692** | 0.683 | 0.675 | 0.627 | 0.599 | 0.583 | 0.531 | 0.474 |
| 151676 | **0.717** | 0.686 | 0.654 | 0.684 | 0.686 | 0.704 | 0.591 | 0.476 | 0.585 | 0.434 |
| Mean | **0.694** | 0.656 | 0.663 | 0.641 | 0.649 | 0.639 | 0.607 | 0.566 | 0.591 | 0.457 |
| Median | **0.700** | 0.654 | 0.666 | 0.669 | 0.671 | 0.627 | 0.595 | 0.564 | 0.589 | 0.449 |

Table 9: One-to-one matched accuracy. Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | 0.720 | 0.528 | 0.604 | **0.777** | 0.737 | 0.715 | 0.565 | 0.602 | 0.639 | 0.530 |
| 151508 | 0.671 | 0.613 | 0.669 | 0.688 | **0.739** | 0.656 | 0.577 | 0.604 | 0.594 | 0.469 |
| 151509 | **0.694** | 0.674 | 0.676 | 0.669 | 0.616 | 0.637 | 0.569 | 0.647 | 0.602 | 0.393 |
| 151510 | **0.709** | 0.674 | 0.705 | 0.665 | 0.678 | 0.602 | 0.638 | 0.624 | 0.480 | 0.492 |
| 151669 | 0.731 | 0.651 | **0.799** | 0.543 | 0.578 | 0.570 | 0.578 | 0.537 | 0.526 | 0.457 |
| 151670 | **0.783** | 0.667 | 0.533 | 0.704 | 0.602 | 0.699 | 0.744 | 0.631 | 0.452 | 0.243 |
| 151671 | **0.751** | **0.751** | 0.679 | 0.738 | 0.723 | 0.728 | 0.662 | 0.691 | 0.559 | 0.479 |
| 151672 | **0.833** | 0.751 | 0.779 | 0.691 | 0.709 | 0.687 | 0.803 | 0.696 | 0.554 | 0.553 |
| 151673 | 0.783 | 0.790 | 0.700 | 0.775 | 0.673 | 0.686 | 0.692 | 0.582 | **0.798** | 0.545 |
| 151674 | **0.783** | 0.583 | 0.724 | 0.527 | 0.663 | 0.642 | 0.523 | 0.587 | 0.533 | 0.443 |
| 151675 | 0.681 | **0.738** | 0.714 | 0.686 | 0.662 | 0.616 | 0.484 | 0.575 | 0.587 | 0.485 |
| 151676 | 0.710 | 0.718 | 0.616 | 0.690 | 0.670 | **0.729** | 0.566 | 0.455 | 0.581 | 0.466 |
| Mean | **0.737** | 0.678 | 0.683 | 0.679 | 0.671 | 0.664 | 0.617 | 0.603 | 0.575 | 0.463 |
| Median | **0.725** | 0.674 | 0.690 | 0.689 | 0.671 | 0.671 | 0.577 | 0.603 | 0.570 | 0.474 |

Table 10: Homogeneity. Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | 0.595 | 0.629 | 0.627 | **0.726** | 0.685 | 0.662 | 0.630 | 0.532 | 0.673 | 0.549 |
| 151508 | 0.578 | 0.639 | 0.677 | 0.673 | **0.692** | 0.617 | 0.564 | 0.554 | 0.619 | 0.471 |
| 151509 | 0.640 | **0.687** | 0.636 | 0.682 | 0.663 | 0.649 | 0.583 | 0.641 | 0.663 | 0.439 |
| 151510 | **0.702** | 0.676 | 0.686 | 0.671 | 0.671 | 0.606 | 0.579 | 0.602 | 0.635 | 0.558 |
| 151669 | **0.710** | 0.623 | 0.687 | 0.570 | 0.624 | 0.612 | 0.675 | 0.526 | 0.573 | 0.426 |
| 151670 | **0.732** | 0.611 | 0.600 | 0.643 | 0.630 | 0.658 | 0.633 | 0.593 | 0.589 | 0.629 |
| 151671 | **0.774** | 0.758 | 0.735 | 0.722 | 0.707 | 0.694 | 0.674 | 0.695 | 0.659 | 0.455 |
| 151672 | 0.715 | **0.734** | 0.697 | 0.717 | 0.705 | 0.712 | 0.704 | 0.707 | 0.608 | 0.578 |
| 151673 | 0.714 | 0.726 | 0.705 | **0.736** | 0.707 | 0.705 | 0.707 | 0.542 | 0.717 | 0.514 |
| 151674 | **0.733** | 0.613 | 0.711 | 0.513 | 0.712 | 0.616 | 0.534 | 0.563 | 0.543 | 0.403 |
| 151675 | **0.694** | 0.634 | 0.687 | 0.686 | 0.684 | 0.619 | 0.588 | 0.588 | 0.511 | 0.477 |
| 151676 | **0.722** | 0.688 | 0.648 | 0.684 | 0.697 | 0.708 | 0.601 | 0.486 | 0.598 | 0.441 |
| Mean | **0.692** | 0.668 | 0.675 | 0.669 | 0.681 | 0.655 | 0.623 | 0.586 | 0.616 | 0.495 |
| Median | **0.712** | 0.657 | 0.687 | 0.683 | 0.689 | 0.653 | 0.615 | 0.575 | 0.613 | 0.474 |

Table 11: Completeness. Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | **0.730** | 0.678 | 0.706 | 0.697 | 0.665 | 0.642 | 0.640 | 0.636 | 0.656 | 0.529 |
| 151508 | **0.680** | 0.640 | 0.634 | 0.641 | 0.647 | 0.622 | 0.570 | 0.529 | 0.587 | 0.457 |
| 151509 | **0.688** | 0.634 | 0.681 | 0.631 | 0.596 | 0.610 | 0.531 | 0.618 | 0.597 | 0.389 |
| 151510 | **0.655** | 0.608 | 0.654 | 0.602 | 0.604 | 0.599 | 0.564 | 0.573 | 0.549 | 0.479 |
| 151669 | **0.614** | 0.561 | 0.593 | 0.461 | 0.501 | 0.543 | 0.523 | 0.443 | 0.474 | 0.366 |
| 151670 | **0.689** | 0.514 | 0.454 | 0.511 | 0.471 | 0.540 | 0.592 | 0.445 | 0.458 | 0.264 |
| 151671 | **0.701** | 0.690 | 0.654 | 0.664 | 0.629 | 0.623 | 0.602 | 0.618 | 0.582 | 0.397 |
| 151672 | **0.753** | 0.704 | 0.690 | 0.659 | 0.650 | 0.655 | 0.731 | 0.658 | 0.573 | 0.548 |
| 151673 | 0.728 | **0.754** | 0.704 | 0.699 | 0.680 | 0.699 | 0.679 | 0.515 | 0.725 | 0.498 |
| 151674 | **0.766** | 0.613 | 0.736 | 0.488 | 0.694 | 0.641 | 0.509 | 0.538 | 0.523 | 0.384 |
| 151675 | 0.689 | 0.681 | **0.697** | 0.683 | 0.667 | 0.637 | 0.612 | 0.580 | 0.555 | 0.475 |
| 151676 | **0.715** | 0.686 | 0.661 | 0.685 | 0.678 | 0.702 | 0.582 | 0.470 | 0.575 | 0.429 |
| Mean | **0.701** | 0.647 | 0.655 | 0.618 | 0.623 | 0.626 | 0.595 | 0.552 | 0.571 | 0.435 |
| Median | **0.695** | 0.659 | 0.671 | 0.650 | 0.649 | 0.630 | 0.587 | 0.555 | 0.574 | 0.443 |

Table 12: Cluster purity. Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | 0.731 | 0.670 | 0.671 | **0.778** | 0.763 | 0.715 | 0.673 | 0.637 | 0.699 | 0.625 |
| 151508 | 0.709 | 0.727 | 0.731 | 0.759 | **0.782** | 0.718 | 0.681 | 0.685 | 0.720 | 0.622 |
| 151509 | 0.745 | 0.740 | 0.745 | 0.744 | 0.710 | 0.757 | 0.697 | 0.737 | **0.783** | 0.600 |
| 151510 | **0.809** | 0.781 | 0.784 | 0.758 | 0.761 | 0.710 | 0.694 | 0.730 | 0.725 | 0.666 |
| 151669 | **0.827** | 0.818 | 0.826 | 0.731 | 0.815 | 0.781 | 0.820 | 0.738 | 0.763 | 0.674 |
| 151670 | 0.843 | 0.843 | 0.836 | 0.834 | 0.846 | 0.824 | **0.847** | 0.826 | 0.818 | 0.823 |
| 151671 | **0.903** | 0.892 | 0.890 | 0.885 | 0.877 | 0.861 | 0.853 | 0.862 | 0.731 | 0.665 |
| 151672 | 0.863 | **0.881** | 0.862 | 0.861 | 0.861 | 0.854 | 0.844 | 0.866 | 0.690 | 0.719 |
| 151673 | 0.788 | 0.807 | 0.784 | 0.812 | 0.794 | 0.752 | 0.780 | 0.645 | **0.818** | 0.676 |
| 151674 | **0.815** | 0.690 | 0.797 | 0.624 | 0.803 | 0.742 | 0.604 | 0.697 | 0.645 | 0.527 |
| 151675 | **0.769** | 0.745 | 0.741 | 0.724 | 0.761 | 0.718 | 0.590 | 0.705 | 0.623 | 0.619 |
| 151676 | **0.784** | 0.768 | 0.698 | 0.749 | 0.757 | 0.765 | 0.656 | 0.574 | 0.718 | 0.567 |
| Mean | **0.799** | 0.780 | 0.780 | 0.772 | 0.794 | 0.766 | 0.728 | 0.725 | 0.728 | 0.649 |
| Median | **0.798** | 0.774 | 0.784 | 0.758 | 0.788 | 0.754 | 0.696 | 0.718 | 0.722 | 0.645 |

Table 13: Fowlkes–Mallows index (FMI). Higher is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | **0.668** | 0.547 | 0.596 | 0.658 | 0.636 | 0.622 | 0.541 | 0.561 | 0.587 | 0.505 |
| 151508 | 0.640 | 0.580 | 0.588 | 0.613 | **0.646** | 0.601 | 0.531 | 0.524 | 0.540 | 0.444 |
| 151509 | **0.708** | 0.610 | 0.619 | 0.606 | 0.582 | 0.581 | 0.493 | 0.590 | 0.576 | 0.371 |
| 151510 | **0.665** | 0.600 | 0.661 | 0.584 | 0.595 | 0.562 | 0.583 | 0.584 | 0.491 | 0.517 |
| 151669 | 0.711 | 0.659 | **0.742** | 0.525 | 0.576 | 0.579 | 0.575 | 0.534 | 0.543 | 0.500 |
| 151670 | **0.872** | 0.645 | 0.554 | 0.664 | 0.585 | 0.687 | 0.770 | 0.596 | 0.521 | 0.292 |
| 151671 | **0.730** | 0.722 | 0.682 | 0.702 | 0.689 | 0.677 | 0.638 | 0.663 | 0.572 | 0.487 |
| 151672 | 0.800 | 0.719 | 0.726 | 0.675 | 0.680 | 0.669 | **0.815** | 0.679 | 0.547 | 0.574 |
| 151673 | 0.727 | 0.712 | 0.649 | 0.672 | 0.629 | 0.619 | 0.631 | 0.487 | **0.734** | 0.478 |
| 151674 | **0.713** | 0.528 | 0.674 | 0.480 | 0.640 | 0.601 | 0.469 | 0.542 | 0.465 | 0.381 |
| 151675 | 0.636 | **0.652** | 0.644 | 0.643 | 0.618 | 0.579 | 0.484 | 0.553 | 0.516 | 0.430 |
| 151676 | **0.661** | 0.654 | 0.580 | 0.627 | 0.622 | 0.657 | 0.515 | 0.420 | 0.532 | 0.415 |
| Mean | **0.711** | 0.636 | 0.643 | 0.621 | 0.625 | 0.619 | 0.587 | 0.561 | 0.552 | 0.450 |
| Median | **0.709** | 0.649 | 0.647 | 0.635 | 0.625 | 0.610 | 0.558 | 0.557 | 0.542 | 0.461 |

Table 14: Spatial chaos (CHAOS). Lower is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | **0.0563** | 0.0564 | 0.0564 | 0.0569 | 0.0573 | **0.0563** | 0.0567 | 0.0576 | **0.0563** | 0.0567 |
| 151508 | **0.0551** | **0.0551** | 0.0552 | 0.0554 | 0.0573 | 0.0553 | 0.0554 | 0.0567 | 0.0554 | 0.0552 |
| 151509 | **0.0535** | **0.0535** | **0.0535** | 0.0539 | 0.0545 | **0.0535** | 0.0537 | 0.0548 | **0.0535** | 0.0536 |
| 151510 | **0.0547** | **0.0547** | **0.0547** | 0.0555 | 0.0564 | **0.0547** | 0.0556 | 0.0572 | **0.0547** | 0.0551 |
| 151669 | 0.0609 | **0.0608** | 0.0614 | 0.0623 | 0.0623 | **0.0608** | 0.0610 | 0.0620 | 0.0616 | 0.0619 |
| 151670 | **0.0603** | **0.0603** | 0.0604 | 0.0617 | 0.0624 | 0.0604 | 0.0624 | 0.0615 | 0.0605 | 0.0620 |
| 151671 | **0.0576** | **0.0576** | **0.0576** | 0.0577 | 0.0590 | **0.0576** | 0.0579 | 0.0585 | 0.0578 | 0.0580 |
| 151672 | **0.0581** | **0.0581** | **0.0581** | **0.0581** | 0.0590 | 0.0582 | **0.0581** | 0.0587 | **0.0581** | **0.0581** |
| 151673 | 0.0603 | **0.0601** | **0.0601** | **0.0601** | 0.0606 | 0.0605 | 0.0611 | 0.0621 | **0.0601** | 0.0606 |
| 151674 | **0.0594** | **0.0594** | **0.0594** | 0.0610 | 0.0600 | 0.0595 | 0.0607 | 0.0606 | 0.0601 | 0.0603 |
| 151675 | **0.0601** | 0.0604 | 0.0604 | 0.0612 | 0.0610 | 0.0602 | 0.0613 | 0.0611 | 0.0603 | 0.0609 |
| 151676 | 0.0621 | **0.0620** | 0.0621 | 0.0625 | 0.0628 | 0.0624 | 0.0627 | 0.0635 | 0.0623 | 0.0636 |
| Mean | **0.0582** | **0.0582** | 0.0583 | 0.0589 | 0.0594 | 0.0583 | 0.0589 | 0.0595 | 0.0584 | 0.0588 |
| Median | **0.0587** | **0.0587** | **0.0587** | 0.0591 | 0.0595 | 0.0588 | 0.0594 | 0.0597 | 0.0591 | 0.0592 |

Table 15: Abnormal-spot fraction (PAS). Lower is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | **0.0031** | 0.0083 | 0.0140 | 0.0388 | 0.0390 | 0.0218 | 0.0282 | 0.0663 | 0.0130 | 0.0253 |
| 151508 | 0.0130 | 0.0132 | 0.0246 | 0.0502 | 0.0413 | 0.0173 | 0.0271 | 0.0737 | **0.0087** | 0.0187 |
| 151509 | **0.0065** | 0.0119 | 0.0186 | 0.0345 | 0.0393 | 0.0177 | 0.0378 | 0.0681 | 0.0098 | 0.0228 |
| 151510 | 0.0121 | 0.0104 | 0.0166 | 0.0425 | 0.0488 | 0.0181 | 0.0445 | 0.1299 | **0.0067** | 0.0319 |
| 151669 | **0.0109** | **0.0109** | 0.0128 | 0.0317 | 0.0429 | 0.0117 | 0.0246 | 0.0688 | 0.0243 | 0.0161 |
| 151670 | **0.0054** | 0.0077 | 0.0243 | 0.0223 | 0.0423 | 0.0106 | 0.0206 | 0.0557 | 0.0157 | 0.0663 |
| 151671 | **0.0058** | 0.0071 | 0.0148 | 0.0246 | 0.0345 | 0.0136 | 0.0221 | 0.0470 | 0.0144 | 0.0136 |
| 151672 | 0.0107 | **0.0072** | 0.0092 | 0.0299 | 0.0336 | 0.0177 | 0.0217 | 0.0446 | 0.0147 | 0.0132 |
| 151673 | **0.0121** | 0.0146 | 0.0176 | 0.0278 | 0.0302 | 0.0294 | 0.0467 | 0.0874 | 0.0434 | 0.0376 |
| 151674 | 0.0180 | 0.0174 | 0.0185 | 0.0378 | 0.0329 | 0.0278 | 0.0343 | 0.0776 | **0.0166** | 0.0270 |
| 151675 | **0.0161** | 0.0214 | 0.0373 | 0.0309 | 0.0331 | 0.0348 | 0.0571 | 0.0682 | 0.0212 | 0.0253 |
| 151676 | **0.0176** | 0.0194 | 0.0275 | 0.0338 | 0.0390 | 0.0260 | 0.0361 | 0.0746 | 0.0182 | 0.0217 |
| Mean | **0.0110** | 0.0125 | 0.0197 | 0.0337 | 0.0381 | 0.0205 | 0.0334 | 0.0718 | 0.0172 | 0.0266 |
| Median | 0.0115 | **0.0114** | 0.0181 | 0.0328 | 0.0390 | 0.0179 | 0.0312 | 0.0685 | 0.0152 | 0.0241 |

Table 16: Mean LISI. Lower is preferred. Bold marks the best value in each row.

| Section | HD | GS | SE | ST | BA | SP | BS | SG | GA | SF |
|---|---|---|---|---|---|---|---|---|---|---|
| 151507 | **1.175** | 1.198 | 1.232 | 1.303 | 1.322 | 1.280 | 1.276 | 1.306 | 1.238 | 1.250 |
| 151508 | **1.169** | 1.225 | 1.303 | 1.314 | 1.313 | 1.228 | 1.235 | 1.373 | 1.233 | 1.227 |
| 151509 | 1.242 | 1.243 | 1.224 | 1.291 | 1.327 | 1.274 | 1.285 | 1.347 | **1.218** | 1.274 |
| 151510 | 1.277 | 1.266 | 1.313 | 1.339 | 1.346 | 1.233 | 1.291 | 1.523 | **1.223** | 1.294 |
| 151669 | 1.174 | **1.167** | 1.217 | 1.233 | 1.281 | 1.182 | 1.211 | 1.330 | 1.208 | 1.168 |
| 151670 | **1.175** | 1.187 | 1.265 | 1.223 | 1.298 | 1.194 | 1.180 | 1.310 | 1.220 | 1.527 |
| 151671 | 1.165 | **1.158** | 1.222 | 1.213 | 1.261 | 1.195 | 1.193 | 1.224 | 1.174 | 1.161 |
| 151672 | **1.151** | 1.157 | 1.179 | 1.236 | 1.244 | 1.230 | 1.187 | 1.245 | 1.233 | 1.165 |
| 151673 | 1.305 | **1.266** | 1.329 | 1.393 | 1.328 | 1.348 | 1.373 | 1.422 | 1.389 | 1.314 |
| 151674 | 1.272 | 1.249 | 1.301 | 1.277 | 1.329 | 1.310 | 1.255 | 1.395 | **1.217** | 1.278 |
| 151675 | 1.290 | 1.273 | 1.355 | 1.373 | 1.317 | 1.313 | 1.337 | 1.344 | **1.233** | 1.254 |
| 151676 | 1.301 | 1.314 | 1.348 | 1.376 | 1.350 | 1.371 | 1.282 | 1.352 | 1.244 | **1.237** |
| Mean | **1.225** | **1.225** | 1.274 | 1.298 | 1.310 | 1.263 | 1.259 | 1.348 | 1.236 | 1.262 |
| Median | **1.209** | 1.234 | 1.283 | 1.297 | 1.320 | 1.254 | 1.265 | 1.345 | 1.228 | 1.252 |

Table 17: Focused depth-encoder ablations on 12 DLPFC sections. Values are section means; adjacent-pair AUC is first averaged over adjacent layer pairs within each section. Higher is better for all columns.

| Variant | Spearman $\lvert\rho\rvert$ | Mean adj. AUC | Monotone $R^2$ |
|---|---|---|---|
| **HodgeDepth** | 0.885 | 0.901 | 0.883 |
| *Encoder variants* | | | |
| Greedy sign propagation | 0.790 | 0.868 | 0.824 |
| Random signs | 0.285 | 0.560 | 0.162 |
| Magnitude-weighted increments | 0.873 | 0.882 | 0.888 |
| *Geometric coordinates* | | | |
| Fiedler vector | 0.661 | 0.665 | 0.530 |
| Diffusion-map coordinate 1 | 0.658 | 0.662 | 0.522 |
| Graph distance from boundary | 0.222 | 0.572 | 0.096 |
| Spatial-coordinate PC1 | 0.633 | 0.649 | 0.458 |

Table 18: Recovery on the six stepwise synthetic examples. Depth error is affine-calibrated NRMSE in percent; lower is better. ARI uses the shared five-cluster decoder; higher is better. Bold marks the better value for each metric. The mean weights datasets equally.

| | | Depth error (%) ↓ | | ARI ↑ | |
|---|---|---|---|---|---|
| Geometry | Spots | Hodge Distance | GASTON | Hodge Distance | GASTON |
| Straight | 3,264 | **1.73** | 13.09 | **0.934** | 0.586 |
| Curved | 3,264 | **1.91** | 5.62 | **0.928** | 0.640 |
| Concentric | 9,302 | 1.90 | **1.31** | **0.804** | 0.684 |
| Wavy | 3,264 | **2.02** | 11.37 | **0.950** | 0.541 |
| Curved with holes | 3,040 | **2.27** | 6.71 | **0.921** | 0.574 |
| Disconnected wavy | 2,592 | **2.24** | 5.74 | **0.953** | 0.700 |
| Unweighted mean | — | **2.01** | 7.31 | **0.915** | 0.621 |

# D SYNTHETIC STUDY: PROTOCOL, DIAGNOSTICS, AND FOLLOW-UP

## D.1 ILLUSTRATIVE RECOVERY IN CONTROLLED LAYERED GEOMETRIES

We test whether a spatial coordinate can recover layered geometry when expression varies systematically between layers but is uninformative within their interiors. The benchmark comprises straight, curved, concentric, and wavy layers, together with curved layers containing holes and disconnected wavy layers. These six stepwise examples are supplemented by two smooth-expression controls and one previously identified failure case.

This synthetic study focuses on recovery of a global one-dimensional depth coordinate. We therefore restrict the comparison to GASTON (Chitra et al., 2025a), which also estimates a continuous global scalar index, termed isodepth, from spatial locations and expression measurements. The shared output allows direct evaluation against the known generating depth and comparison of layer recovery under the same decoder. Broader comparisons of spatial-domain clustering are considered in the real-data analysis.

**Data and protocol.** Each spot has a location $s_i \in \mathbb{R}^2$ and a continuous expression vector $z_i \in \mathbb{R}^{14}$. For the stepwise examples,

$$z_i = \mu_{\ell_i} + \varepsilon_i, \quad \varepsilon_i \overset{\text{iid}}{\sim} \mathcal{N}(0, 0.9^2 I_{14}), \tag{69}$$

where $\ell_i \in \{1, \dots, 5\}$ denotes one of five equal-width bands in the generating signed distance $d_0 \in [-0.8, 0.8]$. The 14 features represent an expression embedding rather than raw gene counts. Because each layer has a constant mean, continuous position within a layer must be inferred through geometric assumptions and spatial regularization. The encoder uses eigen-contrast to weight directional evidence when assembling local axes into global geometry.

We evaluate Hodge Depth and the official GASTON neural-index implementation on the same observations. GASTON initialization is selected by the lowest final expression reconstruction loss. Generating depths and labels are used only for evaluation. Graph construction, local estimation, network architecture, and training budgets are specified in Appendices B and D.

**Evaluation.** We report normalized depth RMSE as a percentage of the generating range, after the same affine calibration to truth for both methods. Calibration is performed separately within each connected component. For layer recovery, we apply a shared five-cluster $K$-means decoder to each native scalar coordinate and report adjusted Rand index (ARI). This evaluates the coordinates under a common decoder, rather than GASTON's native segmentation. Disconnected-component ARIs are averaged by component size, and relative offsets or signs between components are not evaluated.

**Recovery results.** Hodge Distance has lower physical-depth error in five of the six stepwise cases, with mean NRMSE of 2.01% versus 7.31% for GASTON (Table 18). It also has higher common-decoder ARI in all six cases, averaging 0.915 versus 0.621. Figure 8 compares the recovered layers with the generating boundaries. GASTON has lower depth error in the concentric case (1.31% versus

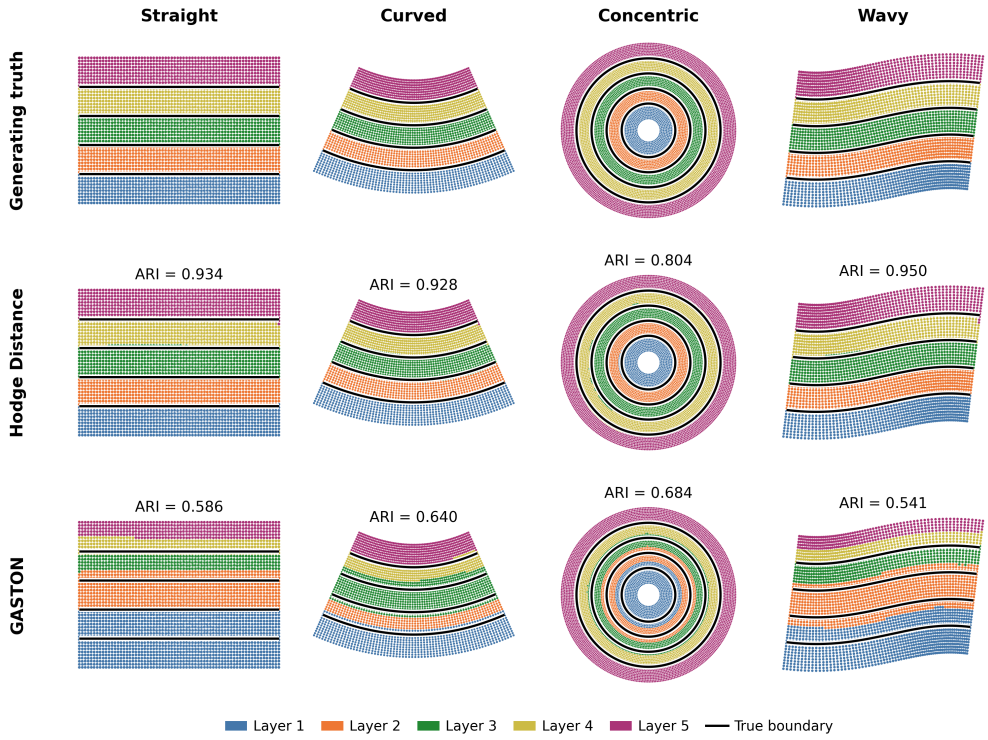


Figure 8: Layer recovery in four stepwise geometries. Rows show generating truth, Hodge Distance, and GASTON; columns show straight, curved, concentric, and wavy layers. Colors denote generating layers or clusters ordered by increasing affine-aligned depth for display. Solid black curves with white outlines mark the same true boundaries in every row, making disagreement with the recovered assignments visible.

1.90%), although its ARI is lower. These metrics measure different properties: $K$-means depends on the spacing and sampling density of the estimated coordinate, and even exact generating depth need not yield ARI 1 (Appendix D).

The two smooth-expression controls also favor Hodge Distance in depth error. The previously observed straight-with-holes failure favors GASTON (30.28% versus 6.66% NRMSE). Across the six main cases, mean local direction errors, measured from the fitted scalar fields using the same local-gradient diagnostic, are similar: $8.23°$ for Hodge Distance and $7.91°$ for GASTON. The gain in physical-depth recovery therefore coexists with comparable average local orientation accuracy. GASTON's reconstruction objective does not enforce physical spacing along its index, whereas the unit-normal Hodge construction explicitly favors that geometry. These results evaluate the complete encoder; attributing gains to individual stages requires matched ablations. The analytical variable-width example in Appendix A.12 distinguishes contour recovery from unit-speed parameterization. Each setting uses one data realization; the multiple GASTON starts address optimization on that realization and do not provide independent data replicates.

### D.2 Dataset scope and fitting

The six primary cases use stepwise means; curved and concentric controls use means that vary smoothly with depth at the same noise level and seed. The spot counts are given in Tables 18 and 19. In the disconnected wavy case, each of the two components contains 1,296 spots. The true layer cutpoints are $-0.48, -0.16, 0.16, 0.48$. The original sampling grids follow the generating geometry, with modest tubular curvature in the curved and wavy cases and annular sampling in the concentric case.

Table 19: Smooth-expression controls. Bold marks the better score.

| | | Depth error (%) ↓ | | ARI ↑ | |
|---|---|---|---|---|---|
| Case | Spots | Hodge Distance | GASTON | Hodge Distance | GASTON |
| Curved, smooth | 3,264 | **0.79** | 1.94 | **0.921** | 0.847 |
| Concentric, smooth | 9,302 | **0.46** | 0.90 | **0.575** | 0.491 |

Hodge Distance uses the support-corrected Hodge Depth implementation, with an edge-length cutoff of 1.8 median nearest-neighbor spacings. These reported runs use local weighted ridge regression, the affine specialization of the general Jacobian-estimation framework.

For GASTON, we use unchanged `gaston-spatial==0.0.2` neural-network code, with spatial architecture $2 \to 20 \to 20 \to 1$ and expression architecture $1 \to 20 \to 20 \to 14$. Training uses ReLU activations, full-batch Adam with learning rate $10^{-3}$, and the official columnwise standardization of coordinates and expression features. All 14 observed features enter directly, without PCA or count preprocessing. Each original dataset has 30 completed starts (seeds 0–29), for 270 fits in total.

### D.3 Evaluation definitions

Let $c(i)$ denote the connected component containing spot $i$. For each method independently, define

$$(\widehat{a}_c, \widehat{b}_c) = \operatorname*{argmin}_{a,b} \sum_{i \in c} (a\widehat{d}_i + b - d_{0i})^2, \tag{70}$$

$$\widetilde{d}_i = \widehat{a}_{c(i)}\widehat{d}_i + \widehat{b}_{c(i)}, \tag{71}$$

$$E_d = \frac{100}{R}\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\widetilde{d}_i - d_{0i})^2}, \tag{72}$$

where $R = 1.6$ for the original nine cases. This post-fit calibration removes sign, offset, and linear scale; it does not remove nonlinear distortion. For example, a 2% error is RMSE 0.032 in physical depth units, not a classification error rate.

The shared decoder applies one-dimensional $K$-means with $K = 5$, 50 initializations, and seed 20260924 to each method's native index within each component. We report nominal ARI against the generating layer labels; component scores are weighted by their spot counts. Colors are relabeled only for display.

## E Batiuk DLPFC: evaluation and configuration details

Table 20: ARI on the 11 analyzed Batiuk DLPFC sections. Means and medians summarize section-level ARIs. Bold marks the largest value in each row.

| | HodgeDepth | GraphST | GASTON |
|---|---|---|---|
| Mean ARI | **0.5686** | 0.4560 | 0.4407 |
| Median ARI | **0.5307** | 0.4361 | 0.4381 |
| Sections ranked first | **9** | 1 | 1 |

### E.1 Analysis subset, annotations, and evaluation sets

The Batiuk data are the Visium component of the public release associated with Batiuk et al. (2022), available from Zenodo record 6921620. The analysis uses the 11 sections listed in Table 21. MB7, MB19 and MB22, present in the original 14-section source inventory, are outside this analysis set. The analyzed subset contains five control sections (MB9, MB11, MB15, MB18-2, and MB21) and six schizophrenia sections (MB6, MB12, MB14, MB23, MB50, and MB52).

Table 21: Section-level comparisons on the 11-section Batiuk analysis subset. All methods use the same scored spots within a section. $K$ is the requested count. The mean row reports the total number of scored spots and unweighted section-mean ARIs. Bold marks the highest ARI in each row.

| Section | Scored spots | $K$ | HodgeDepth | GraphST | GASTON |
|---|---|---|---|---|---|
| MB6 | 2,629 | 7 | **0.6800** | 0.6038 | 0.6731 |
| MB9 | 378 | 7 | **0.7967** | 0.4203 | 0.2989 |
| MB11 | 1,219 | 7 | **0.4978** | 0.3441 | 0.2458 |
| MB12 | 292 | 7 | **0.6904** | 0.5905 | 0.6048 |
| MB14 | 1,301 | 7 | 0.4507 | **0.5009** | 0.4435 |
| MB15 | 505 | 7 | **0.6638** | 0.5457 | 0.5816 |
| MB18-2 | 916 | 7 | **0.5203** | 0.4175 | 0.3399 |
| MB21 | 2,933 | 7 | **0.4829** | 0.3868 | 0.4327 |
| MB23 | 1,250 | 7 | **0.4632** | 0.4274 | 0.2883 |
| MB50 | 2,857 | 7 | 0.5612 | 0.3339 | **0.5769** |
| MB52 | 4,224 | 7 | **0.4480** | 0.4448 | 0.3624 |
| Equal-weight mean | 18,504 | — | **0.5686** | 0.4560 | 0.4407 |

The section-level table reports 18,504 common annotated spots across the 11 analyzed sections. Before HodgeDepth fitting, coordinate-only quality control removes isolated spots and retains the largest component of a pixel-distance graph. The encoder may impose an additional connected-component restriction. Cross-method scoring uses the common annotated spots listed in Table 21: MB6 contributes 2,629 of its 2,633 annotated spots, and MB15 contributes 505 of 506; the other sections retain their reported annotated counts. The three methods retain their respective expression-preprocessing pipelines.

Counts and full-resolution pixel coordinates are joined by section identifier and barcode. Coordinates use pixel column as $x$ and pixel row as $y$. Reference labels are treated as nominal categories, and ARI is calculated directly from the predicted and reference partitions without relabeling predicted cluster IDs. Model fitting uses expression and coordinates.

**Cluster counts and annotation coverage.** Every section uses a requested $K = 7$, corresponding to L1–L6 and white matter. Six sections have all seven categories annotated, and MB9 contains annotations for L1–L3 only. Its ARI therefore describes agreement on that annotated subset. The requested count is a benchmark setting informed by the known annotation categories, not an estimate of the number of domains.

## F Broader related work and geometric context

**Manifold learning and trajectories.** Isomap approximates geodesic distances on a neighborhood graph (Tenenbaum et al., 2000); Laplacian eigenmaps use graph-Laplacian eigenfunctions (Belkin & Niyogi, 2003); and diffusion maps derive multiscale coordinates from a random walk (Coifman & Lafon, 2006). These methods construct global representations from local relationships. HodgeDepth adds an assumption tailored to layered data: local multivariate derivatives share a dominant spatial direction, providing expression-derived evidence for tissue-contour normals. Principal curves estimate continuous structure in feature space (Hastie & Stuetzle, 1989). Spatial depth instead describes corresponding positions in the physical arrangement of molecular variation.

**Multimodal dynamics and downstream summaries.** LYNX, a contemporaneous multimodal preprint, learns shared spatial dynamics through a heterogeneous-attention variational graph autoencoder and extracts continuous or branching trajectories for interaction analysis (Jin et al., 2026). RegionalST uses entropy-based summaries of regional cell-type composition to identify regions of interest (Lyu et al., 2024), while Milo tests differential abundance in cellular neighborhoods between conditions (Dann et al., 2022). A learned depth provides a spatial ordering along which composition and heterogeneity can be studied.

**Why local smoothness is insufficient.** For a spatial graph $G = (V, E)$, the penalty $\sum_{(i,j)\in E} w_{ij} \|u_i - u_j\|^2$ prefers low-frequency signals but does not identify which signal represents depth. The Fiedler vector depends on graph shape and bottlenecks and may align with the tissue's long axis. HODGEDEPTH addresses four additional requirements: expression alignment through the common direction of local multivariate change; orientation consistency across unoriented axes; integrability through fitted increments with zero circulation around graph cycles; and global reconstruction with sensitivity to directional error and graph connectivity. Intermediate diagnostics help reveal departures from the assumed scalar geometry.