HYPERBRIDGE: FROM SPARSE SPATIAL OMICS OBSERVATIONS TO DYNAMIC 3D TISSUE STATES

Anonymous authors
Paper under double-blind review

ABSTRACT

Spatial omics profiles molecular states within their native tissue context, yet dense measurements across three-dimensional space and biological time remain costly and experimentally challenging. Available datasets therefore contain sparse tissue sections or observations from only a few biological stages, leaving spatial regions and intermediate tissue states unmeasured. Although these forms of missingness arise in different experimental settings, both require inference of unobserved spatial–molecular distributions from sparse tissue snapshots. We introduce HyperBridge, a generative method that reconstructs unobserved tissue states as an unbalanced stochastic bridge between endpoint distributions. A dynamic three-dimensional hypergraph captures higher-order relationships among observed cells or spots and provides endpoint-conditioned structural signals for the bridge control, while activity probabilities allow represented mass to vary without requiring particle-wise correspondence. The same formulation supports reconstruction between tissue sections and between three-dimensional states at different biological stages. Experiments across diverse spatial-omics platforms show that HyperBridge improves held-out section reconstruction and preserves tissue organization. Large-scale reconstruction further produces coherent dense tissue atlases, while experiments on planarian regeneration demonstrate reconstruction and forecasting across unobserved biological stages.

1 INTRODUCTION

Spatial omics measures molecular programs in their native anatomical context and reveals how cells are organized within tissue (Bressan et al., 2023; Rao et al., 2021). This organization extends through three-dimensional space and changes during development, regeneration, and disease (Xu et al., 2024). However, spatial profiling is expensive and often destructive, so available datasets usually contain only a limited number of tissue sections or biological stages (Moses & Pachter, 2022). Reconstructing the missing states would connect measured observations into a coherent description of tissue organization and enable the study of molecular changes where direct measurements are unavailable. The central difficulty is that endpoint observations are sparse and unpaired, while their sampling support, resolution, and represented mass may differ; thus, the missing tissue state cannot be recovered by pointwise interpolation.

Existing methods for spatial reconstruction can be broadly divided into three categories. Registration and joint modeling methods place multiple sections in a common spatial coordinate system (Zeira et al., 2022; Wang et al., 2023). They address geometric misalignment but generally do not generate unobserved tissue states. Interpolation and generative methods infer molecular profiles, cells, or virtual sections between measured planes (Lin et al., 2026; Li et al., 2025; Deshpande et al., 2025). However, smooth interpolation or fixed local neighborhoods can weaken tissue boundaries and fail to represent relationships that change along the reconstruction path. Graph-based methods provide local spatial context (Hu et al., 2021), while graphs represent the relationships among multiple observations (Feng et al., 2019). Unbalanced transport allows the represented mass to vary between observations (Pariset et al., 2023). These ideas have not been integrated into a stochastic reconstruction process that allows both the relationships among particles and the represented mass to evolve. A method is therefore needed to generate unobserved tissue states while preserving changing spatial organization and accommodating unequal endpoint mass. Reconstruction across biological time is more difficult because measurements at different stages are usually obtained from different speci-

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.

mens, without one-to-one correspondence between cells (Schiebinger et al., 2019; Cui et al., 2023). Differentiation, proliferation, migration, and tissue remodeling can simultaneously alter molecular states, spatial organization, and the number of observed cells (Cui et al., 2023). Temporal reconstruction must therefore infer an unobserved three-dimensional tissue state from unpaired endpoint states while al-lowing changes in both organization and represented mass(Schiebinger et al., 2019). Since anatomical depth and biological time describe different axes, a spatiotemporal model should share structural information across them while allowing the dynamics specific to biological time to depart from those associated with spatial variation.

We introduce **HyperBridge**, which models spatial and temporal reconstruction as bridge learning between endpoint tissue-state distributions. An unbalanced stochastic bridge is used to accommodate differences in represented mass, and a three-dimensional hypergraph is used to capture higher-order relationships among particles across tissue sections. The hypergraph is updated as the generated state evolves, and its endpoint-conditioned structural message guides the learned control drift instead of relying on a fixed neighborhood structure. Intermediate bridge states are decoded into spatial coordinates and molecular measurements, making them directly comparable with observed tissue data. The same bridge construction can be used for spatial reconstruction between tissue sections and for temporal reconstruction between three-dimensional tissue states. The interpretation of the bridge coordinate and the endpoint observations changes with the reconstruction setting, while the state representation, dynamic hypergraph message, and controlled stochastic evolution remain shared. Our contributions are summarized as follows:

- Spatial and temporal reconstruction from sparse spatial omics observations is formulated as bridge learning between unpaired tissue-state distributions with potentially unequal represented mass.
- Higher-order relationships among particles across three-dimensional tissue sections are modeled with a dynamic hypergraph, and the resulting structural message is used to guide the learned control drift across spatial or biological coordinates.
- HyperBridge is evaluated on held-out sections, dense three-dimensional tissue reconstruction, and biological-stage reconstruction across cell- and spot-resolved datasets, where improved molecular fidelity, tissue organization, and mass consistency are observed over the evaluated baselines.

2 METHOD

2.1 PROBLEM FORMULATION

Considering spatial omics observations collected from a sparse set of tissue sections and, when available, biological stages. Let t denote a biological stage and h denote a section coordinate within that stage. An observed section is represented as $\mathcal{X}_{t,h} = \{(\mathbf{r}_i, \mathbf{x}_i, \mathbf{m}_i)\}_{i=1}^{N_{t,h}}$, where $\mathbf{r}_i$ is the registered spatial coordinate, $\mathbf{x}_i$ is the measured molecular profile, and $\mathbf{m}_i$ contains optional morphological features such as H&E information. An observation may correspond to an individual cell or a spatial spot, depending on the assay. The number of observations and their represented mass may vary across sections and stages.

Our goal is to infer an unobserved tissue state between two observed endpoint states. The endpoints are represented by finite nonnegative measures

$$\mu_\ell = \sum_{j=1}^{N_\ell} a_j^\ell \delta_{\mathbf{q}_j^\ell}, \qquad \ell \in \{0, 1\}, \tag{1}$$

where $a_j^\ell > 0$ is the represented mass of an endpoint observation and $\mathbf{q}_j^\ell = \left(\mathbf{Z}_j^\ell, \mathbf{r}_j^\ell\right)$ contains its latent molecular representation and spatial coordinate. The two endpoint measures may contain different numbers of observations and are not assumed to have particle-wise correspondence. Their total represented masses may also differ:

$$M(\mu_0) = \sum_{j=1}^{N_0} a_j^0, \qquad M(\mu_1) = \sum_{j=1}^{N_1} a_j^1. \tag{2}$$

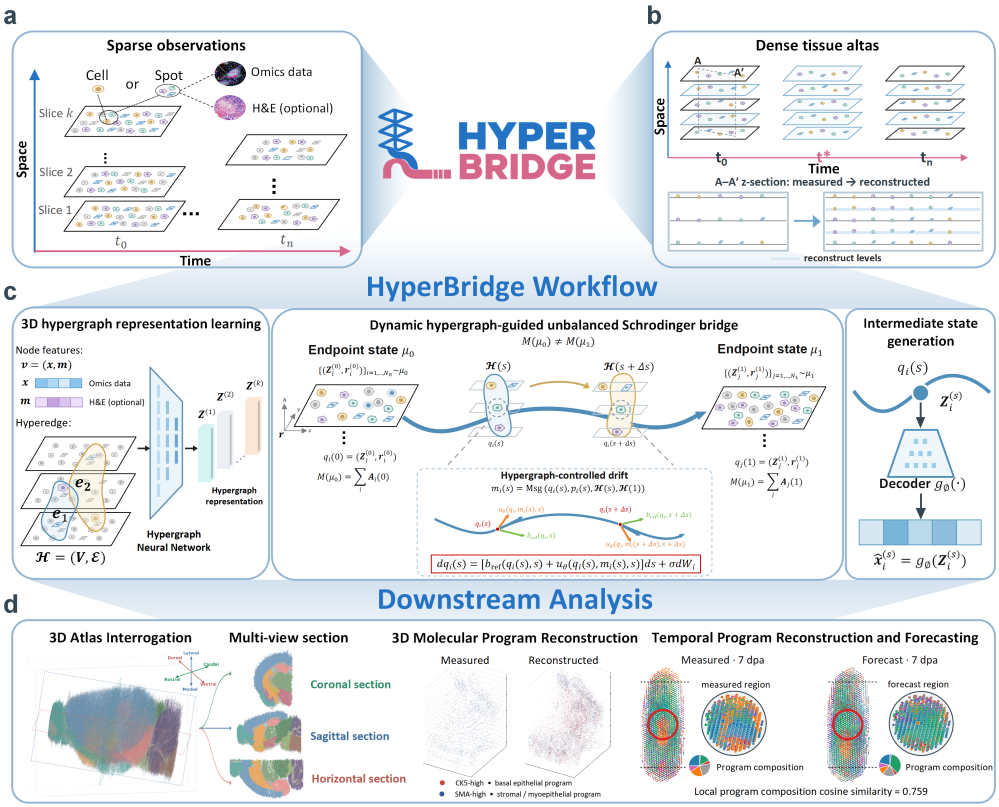

Figure 1: Overview of HyperBridge. (a) Sparse spatial-omics observations across tissue depth and biological stages. (b) Reconstructed dense tissue states and a cross-sectional view of spatial continuity. (c) Three-dimensional hypergraph encoding, dynamic hypergraph-guided unbalanced bridge, and decoding of intermediate states. (d) Downstream spatial and temporal analyses of the reconstructed tissue.

The bridge is simulated using K computational particle slots. Each slot carries a tissue-state trajectory

$$\mathbf{q}_i(s) = (\mathbf{Z}_i(s), \mathbf{r}_i(s)), \qquad i = 1, \dots, K, \quad s \in [0, 1], \tag{3}$$

together with an activity probability $p_i(s)$. The probability $p_i(s)$ represents the probability that slot i contributes to the tissue state at bridge coordinate s. The expected intermediate measure is

$$\rho_s = \sum_{i=1}^{K} \bar{a}_i \, p_i(s) \delta_{\mathbf{q}_i(s)}, \tag{4}$$

where $\bar{a}_i$ is the base mass associated with slot i. Thus, $M(\rho_s) = \sum_{i=1}^{K} \bar{a}_i \, p_i(s)$ can vary along the bridge even though the number of computational slots remains fixed. An inactive slot continues to evolve as a shadow trajectory but contributes no, or only a probability-weighted, mass to the current tissue state.

2.2 3D Hypergraph Representation Learning

We use a hypergraph to represent higher-order relationships among observations. For each observation, molecular and morphological features are first encoded as $\mathbf{Z}_i^{(0)} = E(\mathbf{x}_i, \mathbf{m}_i)$, where the morphology input is omitted when it is unavailable.

Let $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{B})$ denote a hypergraph with node set $\mathcal{V}$, hyperedge set $\mathcal{E}$, and incidence matrix $\mathbf{B}$. A hyperedge can connect more than two observations, and overlapping hyperedges allow one observation to participate in multiple tissue structures (Feng et al., 2019). We construct hyperedges

using spatial and molecular compatibility. For two observations i and j, their affinity is defined as

$$a_{ij} = \exp\left(-\frac{\|\mathbf{r}_i - \mathbf{r}_j\|_2^2}{2\ell_r^2} - \frac{\|\mathbf{Z}_i^{(0)} - \mathbf{Z}_j^{(0)}\|_2^2}{2\ell_Z^2}\right), \tag{5}$$

where ℓ_r and ℓ_Z control the spatial and molecular scales. For each anchor node, observations with the highest admissible affinities are grouped into one or more hyperedges.

The hypergraph representation is learned through normalized propagation:

$$\mathbf{Z}^{(k+1)} = \mathrm{LN}\left[\mathbf{Z}^{(k)} + \varphi_k\left(\mathbf{P}_{\mathcal{H}}\mathbf{Z}^{(k)}\mathbf{W}_k\right)\right], \tag{6}$$

where $\mathbf{P}_{\mathcal{H}}$ is the normalized hypergraph propagation operator and φ_k is a nonlinear transformation. The final representation $\mathbf{Z}_i$ combines the molecular state of observation i with information propagated through overlapping tissue structures.

This learned representation initializes the molecular component of the endpoint state. During bridge simulation, the same type of hypergraph construction is applied to the evolving particle slots, allowing the structural context to change with the generated tissue state.

2.3 Dynamic Hypergraph-guided Unbalanced Schrödinger Bridge

At bridge coordinate s, the current hypergraph is constructed from the slot states and their activity probabilities:

$$\mathcal{H}(s) = \mathcal{H}\left(\{\mathbf{q}_i(s), p_i(s)\}_{i=1}^K\right). \tag{7}$$

The activity probabilities weight the contribution of each slot when local tissue context is aggregated. The terminal tissue state provides an endpoint hypergraph $\mathcal{H}(1)$.

Let $\mathbf{c}_i(s)$ denote the current hypergraph context of slot i:

$$\mathbf{c}_i(s) = \mathrm{Agg}\left(\mathcal{H}(s), \{\mathbf{q}_j(s), p_j(s)\}_{j=1}^K\right)_i. \tag{8}$$

Because the current and terminal observations are not paired, the endpoint context for slot i is obtained through a soft, permutation-invariant alignment:

$$\widetilde{\mathbf{c}}_i(1 \mid s) = \sum_j \pi_{ij}(s)\mathbf{c}_j(1), \tag{9}$$

where $\pi_{ij}(s)$ measures the compatibility between the current state $\mathbf{q}_i(s)$ and the endpoint state $\mathbf{q}_j(1)$,

$$\pi_{ij}(s) = \frac{\exp\left(-d_\eta\big(\mathbf{q}_i(s), \mathbf{q}_j(1)\big)/\tau\right)}{\sum_{\ell=1}^K \exp\left(-d_\eta\big(\mathbf{q}_i(s), \mathbf{q}_\ell(1)\big)/\tau\right)}. \tag{10}$$

The endpoint-conditioned structural message is then defined as

$$\mathbf{m}_i(s) = \mathrm{Msg}\left(\mathbf{q}_i(s), p_i(s), \mathbf{c}_i(s), \widetilde{\mathbf{c}}_i(1 \mid s)\right) = \left[\mathbf{c}_i(s); \widetilde{\mathbf{c}}_i(1 \mid s) - \mathbf{c}_i(s)\right], \tag{11}$$

where $[\,\cdot\,;\,\cdot\,]$ denotes vector concatenation. The first term describes the current local tissue context, while the second term describes the structural change required to approach the endpoint organization. This message is used by the controlled drift and does not impose a one-to-one correspondence between endpoint particles.

A classical Schrödinger bridge is defined relative to a reference stochastic process (Tong et al., 2023). Here, the reference drift is augmented with a learned, endpoint-conditioned correction:

$$\mathrm{d}\mathbf{q}_i(s) = \left[\mathbf{b}_{\mathrm{ref}}\left(\mathbf{q}_i(s), s\right) + \mathbf{u}_\theta\left(\mathbf{q}_i(s), \mathbf{m}_i(s), s\right)\right]\mathrm{d}s + \sigma_s \mathrm{d}\mathbf{W}_i(s). \tag{12}$$

The state $\mathbf{q}_i(s)$ contains both the latent molecular representation and the spatial coordinate. Therefore, the controlled drift can modify both components according to the current state and its hypergraph context.

The represented mass is changed through active and inactive state probabilities. Let $\lambda_i^+(s)$ and $\lambda_i^-(s)$ denote activation and deactivation rates. The corresponding activity probability evolves according to

$$\frac{\mathrm{d}p_i(s)}{\mathrm{d}s} = \lambda_i^+(s)\big(1 - p_i(s)\big) - \lambda_i^-(s)p_i(s), \tag{13}$$

where the rates are predicted by an activity transition model:

$$\lambda_i^+(s), \lambda_i^-(s) = h_\psi\left(\mathbf{q}_i(s), s\right). \tag{14}$$

Equivalently, the implementation may update a binary state $\chi_i(s) \in \{0, 1\}$ using these transition probabilities, with $p_i(s)$ representing its expected active state. Inactive slots continue to follow Eq. equation 12 as shadow trajectories, but their contribution to ρ_s is reduced according to $p_i(s)$.

The activity process provides the source and sink mechanism of the unbalanced bridge. In particular,

$$\frac{\mathrm{d}}{\mathrm{d}s}M(\rho_s) = \sum_{i=1}^{K} \bar{a}_i \frac{\mathrm{d}p_i(s)}{\mathrm{d}s}. \tag{15}$$

Thus, unequal endpoint mass is represented through changes in the probability of active participation rather than through changes in the number of computational slots. This is consistent with unbalanced transport, which allows the total mass of the transported measures to vary (Pariset et al., 2023).

The current hypergraph is rebuilt as the particle states and activity probabilities evolve. This state dependence follows the idea of dynamic hypergraphs, in which the relationships among nodes are updated as node states change (Jiang et al., 2019).

2.4 Intermediate Tissue-State Generation

After the bridge is simulated to an intermediate coordinate s, the latent molecular state is decoded into the observed molecular space:

$$\widehat{\mathbf{x}}_i(s) = g_\xi\left(\mathbf{Z}_i(s)\right). \tag{16}$$

The spatial coordinate is directly obtained from $\mathbf{r}_i(s)$. The decoded intermediate tissue state is therefore represented as

$$\widehat{\mu}_s = \sum_{i=1}^{K} \bar{a}_i\, p_i(s) \delta_{(\mathbf{r}_i(s), \widehat{\mathbf{x}}_i(s))}. \tag{17}$$

For spatial reconstruction, evaluating Eq. equation 17 at different bridge coordinates produces virtual sections between observed planes. For temporal reconstruction, the same operation produces a predicted three-dimensional tissue state at an unobserved biological stage. In both cases, the generated state is a weighted coordinate–molecular measure that can be directly compared with observed data.

2.5 Overall Objective

The model is trained by combining a controlled-bridge objective with reconstruction supervision. The drift regularization term is defined as

$$\mathcal{L}_{\mathrm{drift}} = \mathbb{E}\left[\int_0^1 \sum_{i=1}^{K} \bar{a}_i p_i(s) \frac{\|\mathbf{u}_\theta(\mathbf{q}_i(s), \mathbf{m}_i(s), s)\|_2^2}{2\sigma_s^2}\,\mathrm{d}s\right]. \tag{18}$$

The activity transition model is regularized relative to its reference transition process:

$$\mathcal{L}_{\mathrm{act}} = \mathcal{R}\left(p_i(s), \lambda_i^+(s), \lambda_i^-(s)\right), \tag{19}$$

where $\mathcal{R}$ is evaluated from the active/inactive transition probabilities used by the shadow-step update.

Endpoint consistency is imposed through mass-relaxed endpoint divergences:

$$\mathcal{L}_{\mathrm{end}} = \lambda_0\, \mathrm{KL}_+\left(\rho_0 \| \mu_0\right) + \lambda_1\, \mathrm{KL}_+\left(\rho_1 \| \mu_1\right). \tag{20}$$

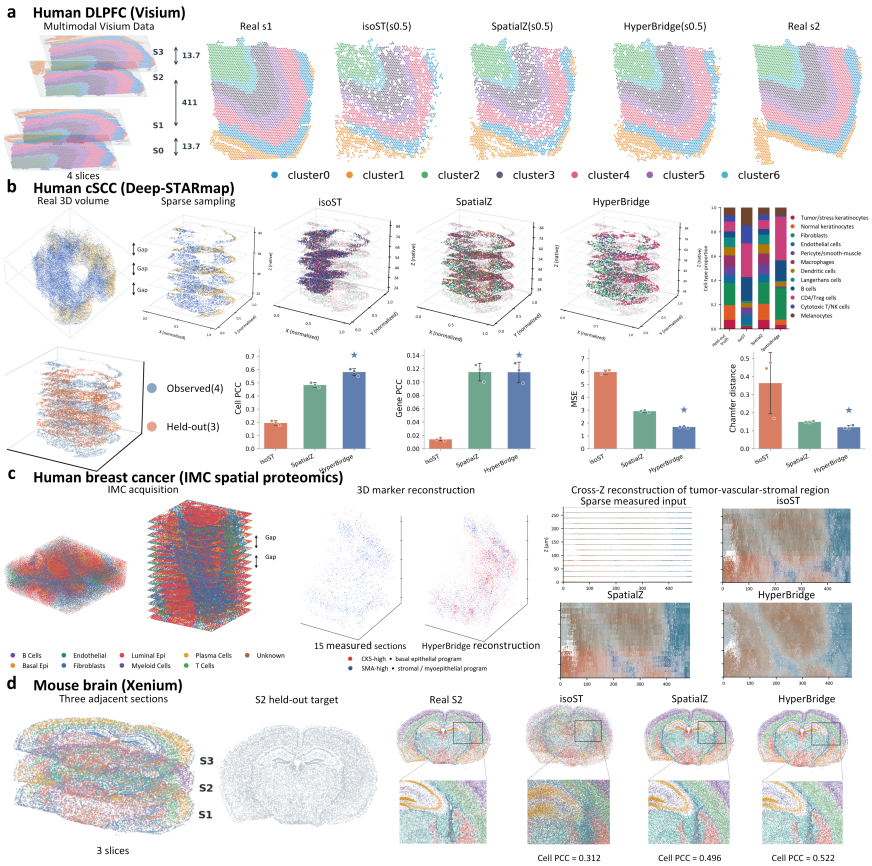

Figure 2: **Cross-spatial-omics platform benchmarking.** Held-out section reconstruction across four measurement technologies: (a) Visium, (b) Deep-STARmap, (c) IMC spatial proteomics, and (d) Xenium. Each panel compares the measured reference with predictions from the competing methods under the same observed-input and target-section split.

For a training state withheld from the conditioning observations, the decoded prediction is compared with the target measure using an unbalanced transport discrepancy:

$$\mathcal{L}_{\mathrm{rec}} = \mathbb{E}_{s^\star}\left[\mathcal{D}_{\mathrm{UOT}}\left(\widehat{\mu}_{s^\star}, \mu_{s^\star}\right)\right]. \tag{21}$$

The total objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{drift}} + \lambda_{\mathrm{act}}\mathcal{L}_{\mathrm{act}} + \mathcal{L}_{\mathrm{end}} + \lambda_{\mathrm{rec}}\mathcal{L}_{\mathrm{rec}}. \tag{22}$$

The target state used in $\mathcal{L}_{\mathrm{rec}}$ is excluded from the encoder hypergraph and endpoint conditioning to prevent information leakage.

The same formulation is used for both reconstruction settings. When the endpoints are observed tissue sections, the bridge coordinate represents relative anatomical position and the generated measures correspond to virtual sections. When the endpoints are three-dimensional tissue states measured at different biological stages, the bridge coordinate represents relative biological time and the generated measures correspond to unobserved stages.

3 Experiments

We organize the evaluation to examine the generality, scalability, and spatiotemporal modeling capacity of HyperBridge. We first assess cross-platform generalization through held-out-section reconstruction on human DLPFC Visium (Maynard et al., 2021), human cSCC Deep-STARmap (Sui et al., 2025), human breast-cancer imaging mass cytometry (Giesen et al., 2014), and mouse-brain Xenium data (10x Genomics, 2023), covering different molecular modalities, spatial resolutions, and observation units. We then evaluate large-scale reconstruction using Allen MERFISH data

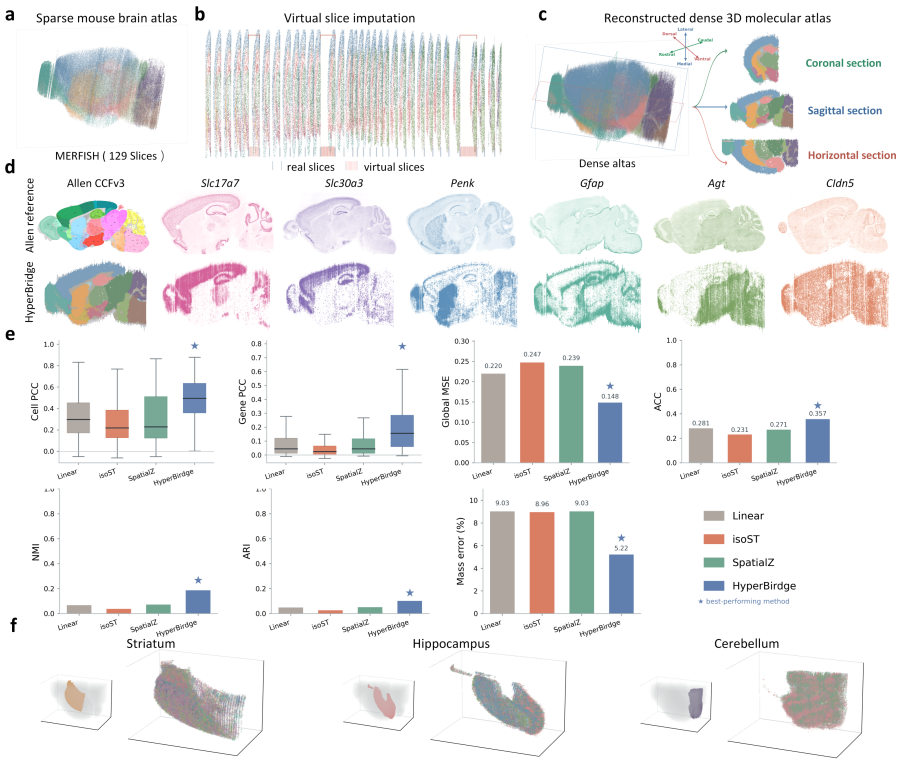

Figure 3: **Large-scale 3D atlas reconstruction.** (a–c) Allen MERFISH measured sections, virtual intermediate layers, and the resulting dense atlas; (d) marker fields; (e) quantitative evaluation on the held-out S100 section; and (f) representative striatum, hippocampus, and cerebellum regions.

(Zhang et al., 2023) by expanding sparse measured sections into a dense three-dimensional molecular atlas. Next, we examine temporal interpolation and forecasting on planarian regeneration data to determine whether HyperBridge can recover unobserved biological stages and predict future tissue states. Finally, an ablation study on the Allen MERFISH subset isolates the contributions of hypergraph representation, dynamic hypergraph updating, and unbalanced transport. In all experiments, target sections or stages are withheld from representation learning and bridge fitting, and predictions are evaluated against the corresponding measurements using coordinate-based nearest-neighbor matching. Dataset-specific splits, preprocessing procedures, and metric definitions are provided in Supplementary Section A.

3.1 Cross-spatial-omics platform benchmarking

Figure 2 tests a common missing-section task across four distinct measurements: spot-level RNA and H&E in human DLPFC Visium, cell-level RNA in human cSCC Deep-STARmap, cell-level protein markers in breast-cancer imaging mass cytometry (IMC), and targeted RNA in mouse-brain Xenium.The IMC experiment is a spatial proteomics task: each segmented cell carries a 25-marker protein-abundance vector, not a gene-expression profile. HyperBridge reconstructs the missing section's spatial distribution and marker patterns across depth, including the displayed epithelial/tumour, vascular, and stromal signals. The other panels show cortical organization in DLPFC, an internal three-dimensional cSCC slab, and the held-out Xenium middle section. For Xenium, Cell PCC is 0.522 for HyperBridge, 0.496 for SpatialZ, and 0.312 for isoST.

3.2 Large-scale 3D atlas reconstruction

We next apply the bridge between 129 measured Allen MERFISH sections to form a 1,311-layer atlas, including 1,182 generated layers (Fig. 3a–c). Generated points are virtual molecular samples, not individually identified new biological cells. Orthogonal sections, marker fields, and regional views of the striatum, hippocampus, and cerebellum assess anatomical coherence (Fig. 3d,f). On

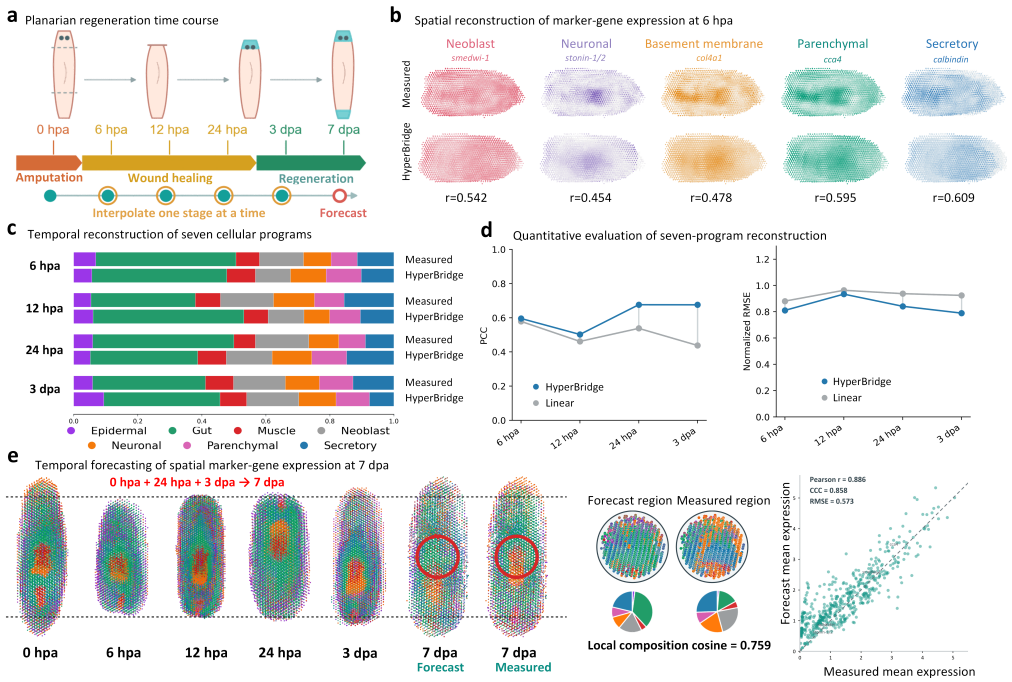

Figure 4: **Temporal reconstruction and forecasting.** (a) Planarian regeneration stages; (b) measured and reconstructed five-marker patterns; (c,d) seven cellular-program composition and metric comparisons with linear interpolation; and (e) spatial marker-gene forecasting from earlier stages to 7 dpa.

the held-out S100 section, HyperBridge achieves a global expression MSE of 0.148 and class-level ACC of 0.357, compared with MSE 0.220–0.247 and ACC 0.231–0.281 for the displayed baselines (Fig. 3e). Its relative mass error is 5.22%, versus 8.96–9.03% for those baselines.

3.3 Temporal reconstruction and forecasting

The spatial bridge is also tested across regeneration stages rather than anatomical depth (Fig. 4a). At 6 hpa, HyperBridge reconstructs five marker-gene patterns with spatial PCC values of 0.454–0.609 (Fig. 4b). Across held-out intermediate stages, seven predefined cellular-program proportions follow the measured stage-wise changes and improve on linear interpolation in the displayed PCC and normalized RMSE comparisons (Fig. 4c,d). For forecasting, observations at 0 hpa, 24 hpa, and 3 dpa are used to predict 7 dpa. The mean-expression comparison gives Pearson $r = 0.886$, CCC $= 0.858$, and RMSE $= 0.573$; local program composition has cosine similarity 0.759 (Fig. 4e). Supplementary Section B provides the five-marker maps across all six measured stages and the complete program readouts. These are population-level reconstructions, not cell-lineage trajectories.

3.4 Component ablation

We ablate one mechanism at a time on the Allen S100 held-out task (Table 1). Replacing hyperedge aggregation with a pairwise GNN reduces Cell PCC from 0.4937 to 0.3951. Freezing the bridge hypergraph reduces Gene PCC from 0.1874 to 0.1575. Enforcing balanced transport raises relative mass error from 0.0522 to 0.6975. The full model is best on the reported Cell PCC, Gene PCC, MSE, ACC, NMI, and mass error; its ARI advantage over the fixed-hypergraph variant is small. Thus the evidence supports the three mechanisms most clearly through complementary molecular and mass measurements, without attributing every score change to one component. Variant definitions and common training/evaluation controls are given in Supplementary Section A.

Table 1: Allen S100 component ablation. All variants share the same held-out target and evaluation protocol. Mass error is the relative total-mass discrepancy; arrows indicate the preferred direction.

Variant	Cell PCC ↑	Gene PCC ↑	MSE ↓	ACC ↑	NMI ↑	ARI ↑	Mass err. ↓
Full HyperBridge	**0.4937**	**0.1874**	**0.1484**	**0.3571**	**0.1872**	0.1017	**0.0522**
w/o Hypergraph (use GNN)	0.3951	0.1667	0.1986	0.2752	0.1374	0.0933	0.6156
w/o Dynamic Hypergraph Update	0.4238	0.1575	0.1684	0.2947	0.1200	0.1029	0.6499
w/o Unbalanced Transport	0.4191	0.1379	0.1783	0.3167	0.1011	0.0744	0.6975

4 Conclusion and Limitation

HyperBridge reconstructs unobserved tissue states from sparse spatial-omics measurements by combining a three-dimensional hypergraph representation, state-dependent structural updates, and an unbalanced bridge. Across spatial transcriptomic and proteomic assays, it recovers held-out sections and large spatial gaps while preserving molecular and regional organization; the MERFISH and planarian experiments further demonstrate dense three-dimensional atlas construction and reconstruction across biological stages. These results support population-level inference across anatomical depth and time, but the reconstructed paths do not identify individual cell lineages or biological birth and death rates. Accuracy depends on endpoint coverage, registration, and assay calibration; quantifying uncertainty and validating inferred transitions with longitudinal measurements remain important next steps.

AI use statement

In this work, we used generative AI tools to assist with translation, searching for relevant literature, and improving the language, clarity, and readability of the manuscript. We did not use generative AI to develop the methodology, design experiments, generate or clean datasets, formulate mathematical claims or proofs, interpret results, or perform qualitative data analysis. All AI-assisted text was carefully reviewed and revised by the authors, who take full responsibility for the accuracy, originality, and final content of this work.

Code availability

The source code for HyperBridge is publicly available at `https://anonymous.4open.science/r/HyperBridge`.

References

10x Genomics. Fresh frozen mouse brain replicates. Xenium In Situ Gene Expression dataset, 2023. URL `https://www.10xgenomics.com/datasets/fresh-frozen-mouse-brain-replicates-1-standard`.

Dario Bressan, Giorgia Battistoni, and Gregory J. Hannon. The dawn of spatial omics. *Science*, 381 (6657):eabq4964, 2023.

Guanshen Cui, Kangning Dong, Jia-Yi Zhou, et al. Spatiotemporal transcriptomic atlas reveals the dynamic characteristics and key regulators of planarian regeneration. *Nature Communications*, 14:3205, 2023.

Ajinkya Deshpande, Zhilei Bei, Jian Ma, and Spencer Krieger. MIMYR: Generative modeling of missing tissue in spatial transcriptomics. *bioRxiv*, 2025. Preprint.

Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3558–3565, 2019.

Charlotte Giesen, Hao A. O. Wang, Denis Schapiro, Nevena Zivanovic, Alexander Jacobs, Birgit Hattendorf, Philipp Schuetz, Raphael Lamprecht, Justin Usary, Milan Mrksich, et al. Highly multiplexed imaging of tumor tissues with subcellular resolution by mass cytometry. *Nature Methods*, 11:417–422, 2014.

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.

Jianwen Jiang, Yuxuan Wei, Yifan Feng, Jingxuan Cao, and Yue Gao. Dynamic hypergraph neural networks. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pp. 2635–2641, 2019.

Bohan Li, Feng Bao, Yuhang Sun, Tianyun Wang, Fengji Li, Hongjue Li, Qionghai Dai, and Yue Deng. Three-dimensional spatial transcriptomics at isotropic resolution enabled by generative deep learning. *bioRxiv*, 2025. doi: 10.1101/2025.08.15.670472. Preprint.

Senlin Lin, Zhikang Wang, Yan Cui, Qi Zou, Chuangyi Han, Rui Yan, Zhidong Yang, Wei Zhang, Rui Gao, Jiangning Song, et al. Bridging the dimensional gap from planar spatial transcriptomics to 3D cell atlases. *Nature Methods*, 23(2):360–372, 2026.

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.

Lambda Moses and Lior Pachter. Museum of spatial transcriptomics. *Nature Methods*, 19(5):534–546, 2022.

Matteo Pariset, Ya-Ping Hsieh, Charlotte Bunne, Andreas Krause, and Valentin De Bortoli. Unbalanced diffusion Schrödinger bridges. *arXiv preprint arXiv:2306.09099*, 2023. Preprint.

Ajinkya Rao, Dalia Barkley, Gustavo S. Franca, and Itai Yanai. Exploring tissue architecture using spatial transcriptomics. *Nature*, 596(7871):211–220, 2021.

Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.

Xin Sui, Jennifer A. Lo, Shuchen Luo, Yichun He, Zefang Tang, Zuwan Lin, Daniel L. Barabasi, Yiming Zhou, Wendy Xueyi Wang, Jia Liu, and Xiao Wang. Scalable spatial single-cell transcriptomics and translatomics in 3d thick tissue blocks. *Nature Methods*, 22(12):2574–2584, 2025. doi: 10.1038/s41592-025-02867-0.

Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. *arXiv preprint arXiv:2307.03672*, 2023. Preprint.

Gefei Wang, Jia Zhao, Yan Yan, Yang Wang, Angela Ruohao Wu, and Can Yang. Construction of a 3D whole-organism spatial atlas by joint modelling of multiple slices with deep neural networks. *Nature Machine Intelligence*, 5(11):1200–1213, 2023.

Jiangshan Xu, Pengcheng Guo, Shijie Hao, Shuncheng Shangguan, Quan Shi, Giacomo Volpe, Keke Huang, Jing Zuo, Juan An, Yue Yuan, et al. A spatiotemporal atlas of mouse liver homeostasis and regeneration. *Nature Genetics*, 56(5):953–969, 2024.

Ron Zeira, Max Land, Alexander Strzalkowski, and Benjamin J. Raphael. Alignment and integration of spatial transcriptomics data. *Nature Methods*, 19(5):567–575, 2022.

Meng Zhang, Xingjie Pan, Won Jung, Aaron R Halpern, Stephen W Eichhorn, Zhiyun Lei, Limor Cohen, Kimberly A Smith, Bosiljka Tasic, Zizhen Yao, et al. Molecularly defined and spatially resolved cell atlas of the whole mouse brain. *Nature*, 624(7991):343–354, 2023.

APPENDIX

A Experimental Details

A.1 Datasets and Task Definitions

Our experiments cover four section-based spatial reconstruction datasets, one whole-brain MERFISH atlas, and one planarian regeneration time course. The numbers below refer to the processed inputs used by HyperBridge, rather than to the size of the original public repositories. For every task, the target section or stage is removed before representation learning, bridge fitting, checkpoint selection, and output generation.

Visium DLPFC. The Visium DLPFC data provide four registered sections (151673–151676), containing 3,611, 3,635, 3,566, and 3,431 spots, respectively (14,243 spots in total), and a common 33,538-gene feature space together with H&E images. The sections are supplied in the So3D-aligned coordinate system; the recorded inter-section step is 13.7 in the native coordinate metadata. For the held-out reconstruction, the target RNA and target H&E are not exposed to the model: the bridge is conditioned on the visible sections surrounding the target and is evaluated after the prediction file has been written. We report spot-wise expression fidelity and spatial point-cloud agreement for this task.

Human cSCC Deep-STARmap. The human cutaneous squamous cell carcinoma data contain 43,818 spatially resolved cells, 254 genes, and 98 recorded depth planes ($z = 5$–102) from a Deep-STARmap volume. We keep the native XYZ coordinates and use the internal-slab protocol implemented for HyperBridge: cells in the held-out slab are removed from fitting, while the slabs on both sides provide the observed endpoints. The main target is the slab centred at $z = 160$; additional centres are used only for the target-wise stability analysis. RNA counts are converted to per-cell library-normalized values (library size 10,000) and then log-transformed before the hypergraph encoder is trained. The target slab is opened only for the final nearest-neighbour expression and geometry evaluation.

BRCA IMC. The BRCA dataset is an imaging mass cytometry experiment containing 15 tissue sections, 106,706 segmented cells, and 25 protein markers. Section depths range from 0 to 280 μm in 20 μm centre-to-centre increments (2 μm section thickness and an 18 μm physical gap). The supplied marker values are retained in their author-provided continuous scale; no RNA library normalization is applied. HyperBridge uses the visible sections to reconstruct an unobserved cross-z section and evaluates both the reconstructed cell locations and the 25-dimensional protein-marker profiles.

Xenium mouse brain. The Xenium mouse-brain experiment contains three adjacent sections (S1–S3), with a 248-gene targeted panel and 474,734 cells in total: 162,033, 154,654, and 158,047 cells in S1, S2, and S3, respectively. The native centroid coordinates are expressed in micrometres and the nominal serial spacing is 10 μm. S1 and S3 are the only training observations; S2 is fully held out, including its RNA, coordinates, and labels, and is used only for post-generation evaluation. Because the local Xenium files do not provide a curated target cell-type ontology, we report molecular and geometric metrics rather than inventing categorical scores for this task.

Allen mouse brain MERFISH. The Allen MERFISH data contain 129 registered sections, 2,616,328 cells, and a 1,122-gene targeted panel spanning a whole mouse-brain volume. The official expression values are retained in their log2-normalized scale and the registered CCF coordinates are retained in millimetres. HyperBridge is used in two modes: (i) generation of intermediate layers between adjacent measured sections to form a dense atlas, and (ii) a held-out virtual-section benchmark with S100 as the target. For the latter, the target expression, coordinates, cell count, and class labels are opened only after the generated prediction is saved. When class metrics are reported, a random-forest decoder is trained on the observed context labels and is used only as an external evaluator.

Planarian regeneration. The planarian data contain six spatial transcriptomic stages (0 hpa, 6 hpa, 12 hpa, 24 hpa, 3 dpa, and 7 dpa), comprising 53,407 spots and 27,664 shared genes in the processed inputs. Each stage also contains seven fixed cellular-program fields from the cell2location readout (Epidermal, Gut, Muscle, Neoblast, Neuronal, Parenchymal, and Secretory). The main temporal task interpolates a held-out stage between observed stages; the forecasting task uses the observed 0 hpa, 24 hpa, and 3 dpa states to predict 7 dpa. The 7 dpa measurements are not used to fit normalization, the bridge, or the program decoder and are read only for the final marker- and program-level evaluation.

A.2 Platform-Specific Preprocessing

Each dataset was converted to a common record containing a spatial coordinate, a molecular feature vector, a section or stage identifier, and optional labels. For Visium, Xenium, Deep-STARmap, and planarian RNA, the run-specific library-normalization and log transform are applied before feature encoding; the feature set is fixed using observed inputs only. Deep-STARmap uses the per-cell library-size-10,000 transform, while the Xenium and Visium adapters apply the corresponding sparse-count normalization in their run configuration. Targeted panels and MERFISH retain all available panel genes; the Allen MERFISH matrix is used in its official log2-normalized scale. IMC protein features remain separate from RNA normalization and retain their continuous marker values. The seven planarian program fields are normalized within the observed training stages before the temporal bridge is fit. H&E features, when used in a multimodal run, are extracted only from observed input sections and are never read from a held-out target.

All sections were placed in the registered coordinate system supplied by the dataset or by the documented registration procedure. Coordinates were scaled only when required for numerical stability, with their relative geometry and section ordering preserved. Target molecular measurements and target images were excluded from model fitting and checkpoint selection.

A.3 Coordinate Matching and Evaluation Protocol

Generated points generally do not coincide exactly with the coordinates of the ground-truth cells or spots. For expression-level comparison, each ground-truth coordinate r_j is assigned the expression vector of its nearest generated coordinate:

$$i(j) = \arg\min_i \|r_j - \hat{r}_i\|_2, \qquad \hat{x}_j = \hat{x}_{i(j)}. \tag{23}$$

This coordinate-only 1-NN projection is performed after generation and does not use ground-truth expression or labels. Geometric metrics are computed directly from the two point clouds. The same held-out target, coordinate frame, and metric convention are used for all methods within a dataset.

A.4 Evaluation Metrics

To quantify the fidelity of the reconstructed expression, protein, spatial, and cellular-program outputs, we report the following metrics. Each metric is computed only on the held-out target after coordinate matching; the target is never used to fit the model or select a checkpoint.

MSE (mean squared error). Let $X \in \mathbb{R}^{N\times G}$ and $\hat{X} \in \mathbb{R}^{N\times G}$ denote the ground-truth and aligned predicted molecular matrices. The global MSE is computed over all matched locations and all retained features,

$$\mathrm{MSE} = \frac{1}{NG}\sum_{n=1}^{N}\sum_{g=1}^{G}(X_{ng} - \hat{X}_{ng})^2, \tag{24}$$

where N is the number of matched locations and G is the number of retained genes or protein markers. Lower MSE indicates a smaller numerical reconstruction error.

Cell/spot PCC. Cell or spot PCC is computed row-wise across the retained molecular features; it measures whether a reconstructed location has the same molecular profile as its ground-truth counterpart. We report the mean over valid locations.

Gene/marker PCC. Gene or marker PCC is computed column-wise across matched spatial locations; it measures whether the spatial distribution of an individual gene or protein marker is preserved. We report the mean over valid genes or markers. Profiles with zero variance in either argument are excluded from the corresponding PCC summary.

Chamfer distance (CD). For predicted coordinates P and ground-truth coordinates T, the symmetric Chamfer distance is

$$\mathrm{CD}(P,T) = \frac{1}{|P|}\sum_{p\in P}\min_{t\in T}\|p-t\|_2 + \frac{1}{|T|}\sum_{t\in T}\min_{p\in P}\|t-p\|_2. \tag{25}$$

ACC, NMI, and ARI. For the Allen held-out-section evaluation, observed-context labels are decoded by the fixed external random-forest evaluator and compared with the matched target labels. We report accuracy (ACC), normalized mutual information (NMI), and adjusted Rand index (ARI). These metrics are not training losses and are not reported for datasets without a validated target label ontology.

Mass error. Relative mass error is reported for bridge runs that export particle weights or a tissue-mass target:

$$\mathrm{MassError} = \frac{|M_{\mathrm{pred}} - M_{\mathrm{GT}}|}{M_{\mathrm{GT}} + \epsilon}. \tag{26}$$

Program PCC, CCC, normalized RMSE, and composition cosine. For the planarian task, the seven program fields are compared after spatial matching. We report program-profile PCC, concordance correlation coefficient (CCC), normalized RMSE, and local composition cosine similarity. PCC and CCC are higher-is-better; normalized RMSE is lower-is-better. The program metrics are computed on the fixed seven program names and do not replace the gene-level molecular metrics.

A.5 Training, Baselines, and Implementation

HyperBridge was trained using only the observed sections or stages and was evaluated at the designated held-out target after the prediction artifact had been saved. In the section-based benchmark runs, the hypergraph autoencoder uses a 256-dimensional hidden layer and a 96-dimensional latent state. The reconstruction objective is optimized with AdamW (learning rate 10^{-3}, weight decay 10^{-4}), 0.1 input dropout, and gradient-norm clipping at 5; the fixed run seed is 2901. The bridge then operates on the standardized joint state consisting of the registered spatial coordinates and the learned latent state. Dataset-specific training budgets and random seeds are fixed before evaluation and recorded in the corresponding run manifests. The target section is not used for checkpoint selection or conditioning. The planarian bridge is fit on the observed regeneration stages, with the target stage excluded from the normalization and program-readout fit.

The baseline comparison includes linear interpolation, isoST, SpatialZ, and HyperBridge under the same target definition, coordinate frame, molecular feature set, and post-generation matching rule. Random seeds, input–target assignments, selected checkpoints, software hashes, and hardware settings are recorded in the corresponding run manifests. The reported metrics therefore describe the actual HyperBridge outputs and the explicitly matched baselines, not a reimplementation-specific substitute metric. Metrics not used in a corresponding main-text experiment are not added to that comparison.

A.6 Ablation Protocol

The ablation study is performed on the Allen S100 held-out section. S100 is removed before representation learning, bridge fitting, checkpoint selection, and prediction. Every variant uses the same observed context, coordinate system, 1-NN matching rule, random seed, training budget, and final checkpoint policy as the full model. We report Cell PCC, Gene PCC, Global MSE, ACC, NMI, ARI, and Mass Error.

Full HyperBridge. This is the complete model: the static three-dimensional hypergraph encoder, the state-dependent hypergraph update inside the bridge, and the unbalanced transport mechanism with variable active mass.

w/o Hypergraph (use GNN). The hypergraph incidence structure is replaced by an ordinary pairwise graph neighbourhood. Both the representation encoder and the bridge-local message operator use the GNN version, and the decoder is retrained for this ablation; no hyperedge aggregation is retained. The bridge, optimizer budget, and evaluation protocol are otherwise unchanged.

w/o Dynamic Hypergraph Update. The initial hypergraph is constructed in the same way as in the full model, but it is frozen during bridge integration. The bridge therefore cannot rebuild local groups from the current state $q_i(s)$, and the same incidence structure is used at every bridge time. The encoder, decoder, transport mechanism, and training budget remain unchanged.

w/o Unbalanced Transport. The dynamic hypergraph and its state-dependent messages are retained, but the mass-change mechanism is disabled. The bridge is forced to use a balanced transport constraint with fixed total mass,

$$M(s) = M_0 = M_1, \tag{27}$$

and no birth/death or active-mass change is allowed. All molecular, spatial, and optimization settings are kept fixed so that the comparison isolates the effect of unbalanced transport.

B Additional Biological Visualizations

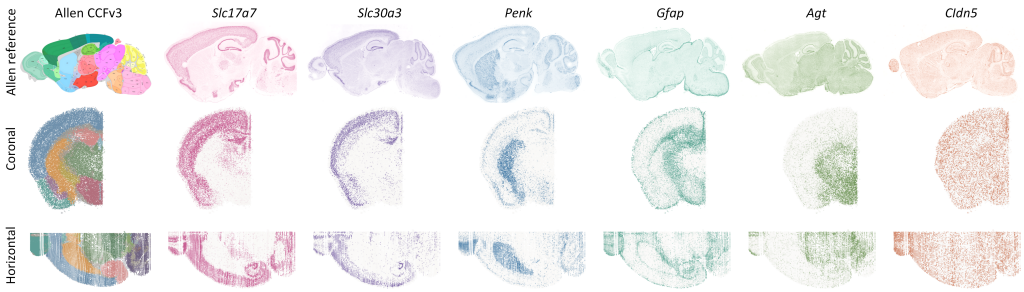

Figure 5: **Additional orthogonal views of the Allen MERFISH reconstruction.** The top row shows the Allen CCFv3 reference and reference marker-gene fields. The middle and bottom rows show HyperBridge-generated virtual sections in the coronal and horizontal planes, respectively, for the same marker genes. The regional annotation colors and marker-specific color scales are kept consistent across views.

This section provides additional biological visualizations that complement the quantitative results and the main-text figures. These views do not introduce new training or evaluation settings; instead, they examine whether the reconstructed molecular patterns preserve recognizable anatomical and biological organization.

Orthogonal Allen MERFISH views. The main text presents the reconstructed Allen atlas from one anatomical orientation. Here, we provide two additional views of the same reconstruction in the coronal and horizontal planes. The Allen CCFv3 reference and marker-gene fields are shown in the top row, while the lower rows display the corresponding HyperBridge-generated virtual views. Consistent regional colors and marker-specific color scales are used across orientations to facilitate the visual comparison of anatomical boundaries and molecular organization.

Planarian marker-gene fields across regeneration. We further visualize marker-gene reconstruction across the planarian regeneration time course, complementing the seven-program analysis in the main text. For each non-anchor stage, the measured marker field is shown alongside the HyperBridge reconstruction. The predicted expression is displayed on the observed target-stage spatial

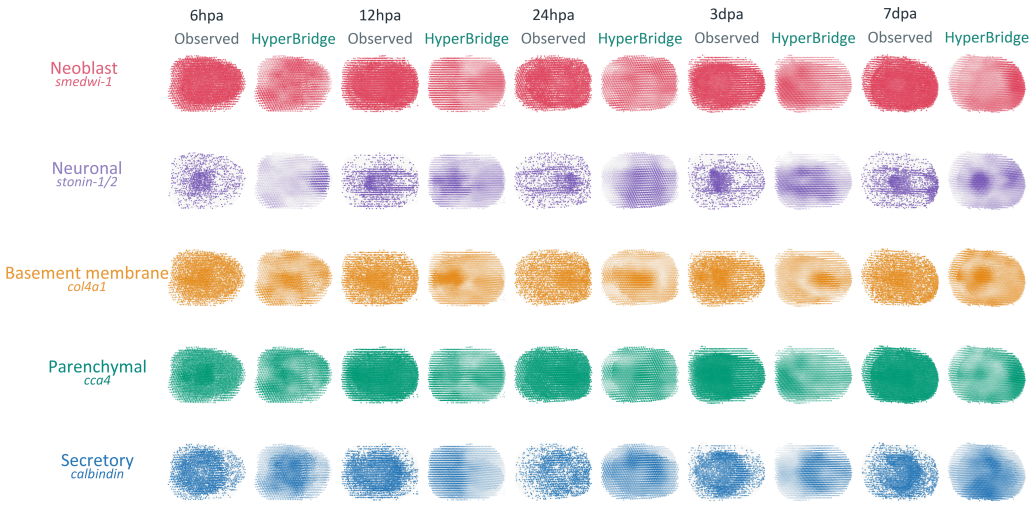

Figure 6: **Marker-gene reconstruction across planarian regeneration.** Each stage is shown as an observed–HyperBridge pair from 6 hpa to 7 dpa. The rows correspond to the neoblast marker *smedwi-1*, neuronal marker *stonin-1/2*, basement-membrane marker *col4a1*, parenchymal marker *cca4*, and secretory marker *calbindin*. HyperBridge predictions are projected onto the observed target-stage geometry for visual comparison.

coordinates, allowing the molecular pattern to be compared without introducing differences caused by alternative point-cloud shapes. The 0 hpa stage serves as the temporal anchor and is therefore not shown as a prediction pair.