
Reasoning Across Space: Tiny Recursive Models for Spatial Omics

Anonymous Authors¹

Abstract

Spatial omics has enabled the study of gene expression within intact tissue, but existing models are limited by local processing, scarce annotations, and poor generalization across tissue sections. Current approaches perform spatial reasoning in a single pass over local regions, restricting the integration of long-range dependencies that underlie many biological signals. In this work, we investigate whether recursive computation provides a natural inductive bias for spatial omics, enabling global reasoning over entire tissue samples without test-time training. We apply Tiny Recursive Models (TRM) for spatial omics tasks and show that it consistently outperforms parameter-matched supervised baselines across robust validation splits. TRM also surpasses several slice-specific unsupervised methods in domain identification and achieves state-of-the-art performance among supervised methods in gene expression prediction. Further analysis reveals that recursion enables test-time scaling and adaptive computation, and that the model is partially generalizable across tissue types for domain identification. Overall, our results demonstrate that recursive computation provides a scalable, data-efficient framework for spatial omics, overcoming the locality constraints of existing approaches without sacrificing generalization.

1. Introduction

Spatial omics has transformed biology, enabling the study of gene expression directly within intact tissue and allowing the elucidation of molecular programs with spatial resolution (Stahl et al., 2016; Chen et al., 2020). Spatial omics presents a challenging learning regime. First, annotations are scarce and expensive, making fully supervised learning difficult at scale. Second, models must explicitly capture spatial structure and dependencies across cells and tissue regions, rather than treating observations as independent samples.

This has motivated a range of models that operate on spatial graphs or localized regions of tissue. For example, spatial domain identification methods typically train individual

graph-based models per tissue section, propagating information through local neighborhoods (Dai et al., 2026; Long et al., 2023; Dong & Zhang, 2022; Hu et al., 2021). Similarly, recent spatial gene expression prediction methods often partition data into independent patches and integrate spatial neighborhoods and histological features (Jia et al., 2024). While effective, most of these approaches remain fundamentally local: spatial reasoning is effectively performed in a single pass over local neighborhoods, without global iterative refinement. This restricts the ability to integrate long-range tissue interactions, that are essential for many biological mechanisms. As a result, most methods often rely on fitting directly to the test sample itself, struggling to generalize across tissue slices or tissue types. Ultimately, current models face a trade-off between scalability and the inductive bias required for global spatial reasoning.

Recent work has shown that recurrent forward passes provide an alternative axis of scaling deep neural networks by increasing the effective depth of the model without increasing parameter count (Geiping et al., 2025; Goyal et al., 2026): rather than learning deeper networks, the same block is repeatedly applied to iteratively refine intermediate representations. Tiny Recursive Models (TRM) (Jolicoeur-Martineau, 2025) additionally demonstrate that such depth recurrence across multiple levels offers additional benefits in terms of reasoning ability in data-scarce regimes, where a small shared network repeatedly refines its own predictions through iterative self-correction.

In this work, we investigate whether recursive computation provides a natural inductive bias for spatial omics. We focus on two representative tasks: spatial-domain identification from gene expression and histology-to-gene-expression prediction. We examine whether TRMs can reason jointly over entire spatial samples using a single shared recurrent model, rather than operating on local neighborhoods or independent patches. We evaluate both tasks under a strict leave-one-slice-out (LOSO) setting, where models must generalize to unseen tissue sections and patient biology without test-time training. Recursive reasoning substantially improves robustness in this regime, outperforming parameter-matched supervised baselines across all LOSO folds and exceeding most established per-slice unsupervised methods in domain identification and all supervised methods in gene expression prediction. We further analyze how recursion affects

performance through test-time scaling and adaptive computation (ACT), show that the reasoning latent space learned by TRM encourages generalization across tissue types in domain identification, and propose practical ACT strategies for continuous prediction tasks.

2. Method

Tasks: We consider two representative spatial transcriptomics (ST) tasks that are both widely used benchmarks, making them suitable for direct comparison with prior work:

1. **Spatial Domain Identification (SDI) task:** Given spot-level gene expression measurements from a 2D Visium section, cluster spots into biologically meaningful spatial domains.
2. **Histology-to-Gene expression prediction (H2G) task:** Given an H&E-stained histology image, predict the spatially resolved gene expression profile at each spot conditioned on paired H&E and ST measurements.

Both tasks provide a natural testbed for evaluating recursive models, as they require learning global spatial dependencies from high-dimensional, sample-limited data.

Architecture: We use the Tiny Recursive Model (TRM) as our core architecture, as it is designed to improve reasoning in data-scarce settings through iterative refinement rather than increased parameter count. Instead of producing the final answer in one forward pass, the TRM recursively refines its predictions using an auxiliary latent reasoning state. Intuitively, the latent state acts as an internal scratchpad that accumulates task-relevant information, whereas the prediction state represents the model’s current best answer—for our tasks, spatial domain assignments (SDI task) or gene expression matrices (H2G task). Starting from randomly initialized noise, the same network updates both states, gradually transforming an initially uncertain prediction into a more accurate output. Across recursion cycles, the same shared network is reused, allowing the model to increase effective depth without increasing the number of parameters. This provides a strong inductive bias for our setting, where learning robust global structure is more important than scaling model size alone. The TRM’s backbone is a 2-layer transformer with 896K trainable parameters, $N_{Z_L}=6$ latent steps, $N_{Z_H}=3$ deep steps, and $N_{outer}=8$ outer steps. In a single forward step, our TRM thus makes $(N_{Z_L} + 1) \cdot N_{Z_H} \cdot 2 = 42$ forward passes. The maximum TRM rollout is $42 \cdot N_{outer} = 336$. Exact pseudocode is provided in Appendix A (Figure 7).

Adaptive computation time (ACT) strategy:

Classical ACT (Graves, 2016) was designed for tasks with a clear “solved/not-solved” signal (e.g., digit addition). Spatial gene-expression prediction is a continuous task: the

output is a high-dimensional regression target, and quality improves *continuously* with depth. We study four halt recipes on TRM using ($N_{z_H}=1$, $N_{z_L}=3$, $K_{max}=8$). Three are slide-level halt-target recipes; **PCC-halt**, **RVD-halt**, and **MSE-halt** that supervise the halt head against the panel-wide PCC / RVD / MSE on a 20% per-step held-out subset of the slide’s supervised spot features. The fourth, **per-class PCC-halt**, supervises a separate halt logit per gene with the halt-prob threshold calibrated to sit *inside* the per-input distribution of converged halt probabilities.

Zero-shot evaluation protocol on a different tissue. To probe out-of-distribution generalization of recursive reasoning, we evaluate the DLPFC-LOSO TRM checkpoints on the 10x Visium human breast cancer dataset of Andersson et al. (2021) without any retraining. This is a structurally distinct tissue type (breast carcinoma rather than cortex), but it shares the same human gene space: 2,872 of the 3,000 DLPFC HVGs are present in the breast cancer dataset, and the remaining 128 genes are zero-filled. We project breast cancer expression onto the DLPFC shared PCA space and pass the resulting 50-dimensional features directly to all 12 DLPFC LOSO checkpoints. Cluster assignments are obtained by k -means with $k=20$ (matching the dataset’s 20 fine-grained histological annotation classes spanning IDC, DCIS/LCIS, tumor edge, and healthy tissue) on the z_L latent, which consistently outperforms z_H in our DLPFC experiments; we report the best ARI over the 12 LOSO checkpoints, and additionally test how much spatial information further influences clustering by concatenating the spatial coordinates with z_L before k -means.

3. Experiments and Results

3.1. Baselines

We compare TRM with two supervised baselines. Our first baseline, **No Recursion 2-Layer**, is a 2-layer non-recursive transformer, equivalent to setting $N_{Z_L}=1$, $N_{Z_H}=1$, and $N_{outer}=1$. This controls for per-step model capacity while removing iterative refinement entirely. Our second baseline is a **Universal Transformer** (Dehghani et al., 2019) without ACT: a 2-layer shared-weight transformer repeated 21 times, thus bringing its effective depth to 42 (equivalent to the effective depth of a single TRM forward pass). This provides a depth- and weight-sharing-matched control for repeated computation, allowing us to test whether TRM’s gains stem from its structured separation of latent-state refinement and solution refinement, rather than from recurrence depth alone. A fully depth-matched baseline, that is, training a model with depth equivalent to the maximum TRM rollout, consistently overfitted and failed to generalize under the same training data, and hence we omit it from the comparison.

For both tasks, we additionally compare against several

Table 1. Per-slice ARI for all 12 LOSO SDI folds. TRM outperforms both baselines on all 12 slices. †Published per-slice results from Dai et al. (2026) source data (unsupervised methods that train on the test slice).

	151507	151508	151509	151510	151669	151670	151671	151672	151673	151674	151675	151676	Mean
No Rec. 2-Layer (LOSO)	.412	.366	.466	.337	.566	.580	.578	.546	.447	.406	.419	.413	.461
Univ. Transformer (LOSO)	.462	.407	.458	.408	.636	.612	.541	.547	.437	.442	.410	.439	.483
TRM (LOSO)	.519	.492	.555	.525	.695	.736	.665	.591	.520	.520	.466	.426	.559
SpaGCN†	.414	.369	.354	.336	.423	.406	.380	.363	.340	.352	.292	.293	.360
STAGATE†	.485	.471	.518	.483	.570	.535	.573	.510	.478	.483	.404	.376	.491
GraphST†	.504	.527	.589	.547	.621	.588	.571	.547	.547	.480	.417	.440	.532
3d-OT†	.643	.639	.685	.622	.684	.633	.849	.798	.703	.722	.646	.631	.688




Figure 1. DLPFC test-time recursion ablation (mean \pm std, 12 SDI LOSO folds). ARI increases with total processing blocks, saturating around 3. Latent recursion steps further improve performance.




Figure 2. H&E overlay of ground truth and TRM predictions with 1, 3, and 6 latent recursion steps (slice 151508, LOSO). Increasing the number of latent recursion steps sharpens cortical layer boundaries.

established state-of-the-art per-slice methods: for the SDI task, SpaGCN (Hu et al., 2021), STAGATE (Dong & Zhang, 2022), GraphST (Long et al., 2023), and 3d-OT (Dai et al., 2026), whereas for the H2G task, STEM (Zhu et al., 2025), BLEEP (Xie et al., 2023), TRIPLEX (Chung et al., 2024), and HisToGene (Pang et al., 2021).

3.2. Spatial Domain Identification

We use the dorsolateral prefrontal cortex (DLPFC) dataset (Maynard et al., 2021), which includes 12 human DLPFC Visium sections with expert annotations for seven cortical layers. We map each sample to a 2D grid and featurize gene expression with PCA over 3,000 highly variable genes (HVGs). We adopt a strict **leave-one-slice-out (LOSO)** protocol: each fold trains on 11 slices and evaluates zero-shot on the held-out 12th slice. Importantly, the 12th sample is excluded from PCA fitting and is projected into the training PCA space only at test time. Full implementation details are provided in Appendix B.

Comparison to baselines. We evaluate all methods using the Adjusted Rand Index (ARI) between predicted cluster assignments and the ground-truth cortical layer annotations. Table 1 shows that TRM substantially outperforms both supervised baselines on all test slices, demonstrating that recursive refinement provides benefits beyond increased effective depth alone. Compared to established, per-slice methods, TRM achieves a mean LOSO ARI of 0.559, outperforming SpaGCN (0.360) and STAGATE (0.491), and GraphST (0.532), despite *never observing the test slice dur-*

ing training. We note that this comparison is not strictly apples-to-apples: these methods are unsupervised and are fit explicitly on the test slice itself, while the TRM is trained on other slices. Nevertheless, the fact that a cross-slice model matches or surpasses per-slice methods underscores its generalization. Interestingly, the only method that substantially outperforms TRM-LOSO is 3d-OT (0.688), which combines a PointNet++ backbone with an optimal transport alignment step. 3d-OT encodes spatial structure through a fixed hierarchical decomposition, recursively over nested spatial partitions to build multi-scale representations. This provides a strong built-in inductive bias for spatial data, further strengthened by explicit access to the test slice through optimal transport alignment. Conversely, the TRM operates under a stricter setting where it does not impose an explicit spatial hierarchy, but rather learns to construct it dynamically through repeatedly updating its predictions via recursive computation.

Do recursions help performance? To isolate the role of the three abstractions of recursion, we perform inference-time ablations while keeping the model frozen. We observe that the model needs ~ 3 total processing blocks to reach peak ARI, where each block comprises multiple latent recursion steps followed by one deep recursion update (Figure 1(a)). These blocks can arise from any combination of deep supervision steps and deep recursion cycles—the two axes are interchangeable at test time because the outer step boundary performs only read-only operations (Appendix C). Next, we compute the ARI vs N_{Z_L} for different values of N_{Z_H} . At $N_{Z_H} = 1$, improvement is capped at 3 N_{Z_L} steps, indi-




Figure 3. Zero-shot DLPFC \rightarrow breast cancer transfer. The TRM z_L latent benefits substantially from concatenated spatial coordinates (+0.15 ARI), more than the non-recursive baseline does.

cating that further deep recursion is necessary to improve latent recursion performance, which is validated by improvement for more N_{Z_L} steps at $N_{Z_H} \in \{2, 3\}$. Furthermore, at full N_{outer} and N_{Z_H} steps, we see significant improvement when increasing N_{Z_L} (Figure 1(c)), reducing the number of latent recursions from 6 to 1 drops mean ARI by up to 0.25. Figure 2 qualitatively visualizes the combined effect of recursive refinement on one example slice: with only 1 latent recursion step, the model produces near-random predictions (ARI = 0.003); at 3 steps it recovers the gross laminar structure (ARI = 0.417); and at 6 steps it achieves the sharpest layer boundaries (ARI = 0.489).

Zero-shot transfer to human breast cancer. Following the protocol described in Section 2, we evaluate the DLPFC-LOSO TRM checkpoints on the breast cancer dataset of Andersson et al. (2021) without any retraining and report the best ARI against the 20 fine-grained histological annotation classes (Table 2, Figure 3). TRM achieves 0.388 ARI on the z_L latent alone, already exceeding the No-Recursion 2-Layer baseline (0.343), the Universal Transformer (0.206), and matching the unsupervised SpaGCN baseline (0.390). Concatenating the spatial coordinates onto z_L before clustering raises ARI to **0.538**; the same spatial-coordinate addition also helps the non-recursive baseline (0.343 \rightarrow 0.458), suggesting that the recursive latent encodes a representation that is complementary to, rather than redundant with, explicit spatial context. Conversely, increasing N_{outer} at evaluation degrades zero-shot ARI, indicating that deeper refinement on this OOD task overcommits to the in-distribution prior. Overall, we find that TRM moderately generalizes across tissue types in the SDI task without any test-time training, and recursion is a transferable architectural feature in the sense that it produces a latent that benefits from added spatial context more than its non-recursive counterpart does.

3.3. Spatial gene expression prediction

We use the HER2ST dataset (Andersson et al., 2021), which contains 36 samples with spot-level gene expression data and paired H&E images. We map each sample to a 2D spot grid and featurize each spot with a frozen Gigapath (Xu et al., 2024) embedding (1536-d) extracted from a

Table 2. Human breast cancer spatial clustering (max ARI, 20 fine classes). Zero-shot methods are trained on LOSO DLPFC data only. [†]Values taken from Dai et al. (2026). [§]L2-normalised z_L .

Method	Max ARI
PCA k -means (DLPFC HVG projection)	0.289
No Rec. 2-Layer	0.343
No Rec. 2-Layer + spatial	0.458
Univ. Transformer	0.206
TRM (z_L)	0.388
TRM (z_L + spatial[§])	0.538
SpaGCN	0.390
STAGATE	0.572
GraphST	0.577
3d-OT [†]	0.683

Table 3. HER2ST B1 holdout under the Zhu et al. (2025) protocol (300 HMMVG, log-transformed counts). [†]Published source data from Zhu et al. (2025).

	PCC-10	PCC-300	RVD
No Rec. 2-Layer	0.865	0.762	0.257
Univ. Transformer	0.830	0.745	1.088
TRM ($N_{outer}=8$)	0.883	0.785	0.354
HisToGene [†]	0.6812	0.5250	10.34
BLEEP [†]	0.7727	0.5652	0.6025
TRIPLEX [†]	0.7907	0.5766	0.6428
STEM [†]	0.8298	0.5984	0.0693

224 \times 224 H&E patch centered on the spot; per-sample grids are padded to a square 35 \times 35 union shape and accompanied by a binary tissue mask. Following (Zhu et al., 2025), targets are log-transformed values of the top 300 high mean and high variance (HMMVG) genes selected on the union of train-slide highly-variable genes per the protocol of Zhu et al. (2025). All models are trained under a single-slide holdout protocol, matching Zhu et al. (2025): B1 is the test slide chosen for its difficulty, B6 is the validation slide for checkpoint selection, and the remaining 34 slides form the training set.

Comparison to baselines. We evaluate with top- k mean Pearson correlation in log space (PCC- k) and relative variation distance (RVD) (Table 3). We see that TRM (fixed-K ACT) outperforms all methods on PCC across the whole panel. STEM has a diffusion objective and leads in RVD, but for regression (our task) we find that this RVD gap can be mitigated by using ACT-halting (Table 8). More detailed metrics are provided in Appendix B.

Do recursions help performance? We repeat the test-time recursion ablation of the SDI block-equivalence procedure (Figure 1) on the HER2ST B1 holdout using the converged TRM ($k=8$, 300-HMMVG) checkpoint, sweeping N_{Z_L} , N_{Z_H} , and N_{outer} at inference while keeping the

model frozen (Figure 5); following Zhu et al. (2025), we omit error bars because this is a single-slide holdout. The qualitative shape mirrors DLPFC: roughly three processing blocks suffice to reach peak PCC-top-50 (rising from 0.49 at one block to 0.72 by three blocks and saturating, with RVD dropping from 0.85 to a minimum of 0.26 at three blocks before stabilising near 0.35), and outer and deep recursion are interchangeable at fixed $N_{Z_H} \cdot N_{outer}$. The deep-recursion update is again the load-bearing component of the latent→deep→latent message-passing loop: with a single deep recursion ($N_{Z_H}=1$), latent depth alone caps at PCC-mean ≈ 0.51 even at $N_{Z_L}=8$; $N_{Z_H}=2$ reaches 0.69 and $N_{Z_H}=3$ reaches 0.72 (Figure 5b). Unlike DLPFC, however, increasing N_{Z_L} at the default config ($N_{outer}=8$, $N_{Z_H}=3$) does not improve PCC: all $N_{Z_L} \in \{1, \dots, 8\}$ give identical PCC-mean = 0.7221, PCC-50 = 0.8414, and RVD = 0.354 to four decimal places (Figure 5c), indicating that the trained model has fully saturated and additional latent applications make no measurable difference. We attribute this to the gene-expression target being a smoother, lower-frequency function of the foundation-model features than the cortical-layer label: once the deep-recursion loop has refined the spatial features, additional latent passes have nothing left to improve. This finding also indicates that with the default TRM config, ACT would not work because outer recursions add no useful improvement to performance, motivating the smaller-recursion architecture used in Section 4 that re-exposes a learnable outer-recursion gradient.

Alignment with expert annotations. Following the downstream evaluation protocol used by STEM (Zhu et al., 2025), we also assess whether the predicted expression profiles preserve biologically meaningful spatial structure. We cluster TRM predictions and compare the resulting clusters with the expert annotations of the tissue regions. Figure 4 shows that the predicted expression from TRM yields clusters that align well with the major annotated regions. From literature, we select a set of cancer-related genes and visualize model predictions for these in Appendix E.

4. ACT in the Non-Discrete Setting

Classical ACT (Graves, 2016) was designed for tasks with a clear “solved/not-solved” signal per input (e.g., digit addition). Spatial gene-expression prediction is a continuous task: the output is a high-dimensional regression target, quality improves *continuously* with depth, and there is no binary correctness signal. We adapt ACT to this regime on the H2G task using a shallow-inner architecture ($N_{z_H}=1$, $N_{z_L}=3$, $K_{max}=8$), trained on the STEM 2025 HER2ST protocol (Zhu et al., 2025) with 302 HMHV genes. The rationale for scaling N_{z_H} , N_{z_L} down from the defaults — namely that the default architecture saturates N_{outer} at step 1 while the shallow architecture exposes a learnable test-time

Table 4. ACT comparison on HER2ST B1 holdout (shallow TRM, $K_{max}=8$). \bar{n}_{halt} is the empirical mean halt step over the 36 slides; values $< K_{max}$ indicate halting fires.

Strategy	PCC-300	RVD	\bar{n}_{halt}
TRM fixed $N_{outer}=8$	0.785	0.354	8.0
PCC-halt (slide-level)	0.723	0.163	8.0
RVD-halt (slide-level)	0.723	0.164	8.0
MSE-halt (slide-level)	0.779	0.109	2.0
Per-class PCC-halt	0.780	0.158	4.9
STEM (Zhu et al., 2025)	0.598	0.069	—

scaling gradient — is given in Appendix D.2.

Halt-head supervision. We compare four halt-target families. Three supervise a single halt logit against a slide-level quality estimate computed on a 20% per-step held-out subset of the supervised spot features (the held-out subset is re-sampled every iteration so that each spot supervises both objectives across an epoch but never within the same forward pass). At step k :

$$\text{PCC-halt: } q^* = \sigma((\overline{\text{PCC}}_k - \tau) s), \quad \tau=0.65, s=20,$$

$$\text{RVD-halt: } q^* = \sigma((\tau - \text{RVD}_k) s), \quad \tau=0.4, s=4,$$

$$\text{MSE-halt: } q^* = \sigma((\tau - \text{MSE}_k) s), \quad \tau=0.9, s=10,$$

all trained against the halt logit via BCE. The fourth recipe, **per-class PCC-halt**, supervises one halt logit per gene against its own held-out PCC trajectory ($q_{halt,g}^* = \sigma((\text{PCC}_{g,k} - \tau) s)$); the slide-level halt decision $q_{halt} = \frac{1}{G} \sum_g q_{halt,g}$ is unchanged but the head receives a richer per-gene gradient. Here, q^* denotes the target halt probability, $\sigma(\cdot)$ is the logistic sigmoid function, τ is the metric-specific quality threshold, s controls the sharpness of the transition around that threshold, k indexes the refinement step, g indexes genes, G is the total number of genes, $\overline{\text{PCC}}_k$ is the mean held-out Pearson correlation across genes at step k , and RVD $_k$ and MSE $_k$ are the corresponding held-out relative volume difference and mean squared error estimates.

Quality-compute results. Table 4 reports panel-wide PCC, RVD, and the empirical mean halt step \bar{n}_{halt} across the 36 slides. The two key takeaways: (i) every halt-loss recipe cuts RVD by $\geq 2\times$ over fixed-depth TRM at matched PCC, with **MSE-halt** achieving the lowest slide-level RVD (0.109, $3.2\times$ better than fixed-depth, closing $\sim 75\%$ of the gap to STEM’s diffusion-sampled baseline); (ii) only **per-class PCC-halt** produces non-trivial input-adaptive halting ($\bar{n}_{halt}=4.9$, between the floor and ceiling), giving $\sim 42\%$ mean compute reduction at matched PCC.

Slide-level halt heads collapse but regularize. PCC-halt and RVD-halt do not fire: their converged mean halt probability (0.27 and 0.07) is uniformly below the default




Figure 4. (a) The H&E held out slide, (b) Manual pathologist annotations, (c) TRM-predicted clusters align well with expert annotations




Figure 5. Test-time recursion ablation on the HER2ST B1 holdout using PCC-top-50 and RVD.

threshold $\theta=0.7$, so the model defaults to K_{\max} on every slide. The cause is calibration mismatch — the BCE loss is dominated by early steps where the target sits near 0, pulling the converged halt logit down (full mechanism in Appendix D.2). Even without firing, both recipes still cut RVD $\sim 2\times$ versus fixed-depth TRM: the halt-loss BCE acts as a variance-aware regularizer on the prediction head whether or not the controller halts. MSE-halt sits at the opposite extreme — its target trajectory has the largest dynamic range, the halt head saturates above θ on every input, and the model halts at the configured floor. Crucially, MSE is also the prediction-head loss, so MSE-halt adds a sample-level view of the same quality signal z_H already optimizes, biasing the prediction head toward preserving across-spot variance — exactly what RVD penalizes — and explaining the headline RVD gain.

Per-class PCC-halt yields input-adaptive halting. Two changes recover real adaptive halting: (i) lower the halt-prob threshold to $\theta=0.4$, moving it *inside* the per-input distribution of converged halt probabilities, and (ii) supervise per gene rather than per slide, so each of the 302 halt logits sees its own PCC trajectory and the slide-level decision averages 302 informative gradients. Figure 6 shows the result: across the 36 HER2ST slides the empirical halt step is bimodal at $\{2, 8\}$, with 17 slides halting at the floor (halt-prob ≥ 0.69), 17 collapsing below threshold and running to K_{\max} , and the test slide B1 landing between the two clusters at step 5. Halt step is inversely correlated with per-slide PCC — *high-quality slides halt early; low-quality slides keep computing*

Table 5. Cluster sizes and test-slide outcomes across four HER2ST holdout folds (shallow TRM + per-class PCC-halt, threshold 0.4). “Early” = halt step ≤ 4 ; “Late” = halt step = $K_{\max}=8$.

Fold	Test slide	Test PCC	RVD	Halt step	Early	Late
0	B1	0.7253	0.158	5 (boundary)	18/36	17/36
1	B6	0.7196	0.144	2 (early)	14/36	22/36
2	C1	0.6257	0.186	2 (early)	19/36	17/36
3	D1	0.5381	0.286	8 (late)	13/36	23/36

— which is exactly the input-adaptive structure classical ACT was designed to produce, recovered here for a continuous structured-output task once the threshold is calibrated to the per-input halt-prob distribution.

Cross-fold validation. To rule out a single-fold artifact, we retrain the same recipe on three additional held-out test slides drawn from different patients (B6, C1, D1). The bimodality replicates on every fold (Table 5) with mild variation in cluster sizes (early cluster 13–19 of 36; late cluster 17–23 of 36) but unchanged qualitative structure. Across all four folds the held-out test slide lands in the cluster consistent with its own converged PCC: high-PCC test slides (B1, B6, C1) halt early, the low-PCC holdout (D1, PCC 0.54) sits in the late cluster and runs to K_{\max} , and only B1 happens to land at the boundary at step 5. Per-fold halt-step distributions are shown in Appendix Figure 11.

Compute-quality trade-off. Aggregating \bar{n}_{halt} across the 36 slides on the per-class checkpoint gives a $\sim 42\%$ mean compute reduction relative to fixed $N_{\text{outer}}=8$ at matched PCC and $2.2\times$ lower RVD (0.158 vs. 0.354). At the more aggressive operating point $\text{min_halt_steps}=1$ the saving rises to 50.7% at a 0.4% PCC drop; lowering θ further to 0.3 regresses on PCC because the under-trained halt head fires before the model has converged. A complementary analog of these recipes for the discrete DLPFC SDI task — where the halt head is supervised by a soft target derived from held-out cross-entropy — is provided

PCC-halt with halt-threshold 0.4 on shallow TRM (HER2ST, all 36 slides)




Figure 6. Input-adaptive halting under per-class PCC-halt with halt threshold 0.4, evaluated per-slide on all 36 HER2ST slides. (a) Halt-step histogram is bimodal at $\{2, 8\}$; the test slide B1 sits at step 5 (mean 4.92, std 2.96). (b) Halt step vs. per-slide PCC: easy slides ($\text{PCC} \gtrsim 0.70$) halt at the `min_halt_steps` floor, hard slides ($\text{PCC} \lesssim 0.55$) never cross threshold and run to K_{max} . (c) Halt step vs. converged halt probability: the 0.4 threshold (red dashed) cleanly partitions the halt-prob distribution into the two halt-step clusters.

in Appendix D.1; the qualitative conclusion mirrors the H2G case (halt-loss-as-regularizer dominates the quality–compute frontier), with the caveat that on a discrete output the natural calibration of the halt head is coarser.

5. Discussion and Limitations

This work asks a simple question: can iterative reasoning, rather than a single forward pass, improve how models represent tissue? Across both tasks, our results suggest that the answer is yes. On spatial domain identification, recursion is clearly beneficial: TRM consistently outperforms parameter-matched non-recursive baselines, generalizes under leave-one-slice-out evaluation, and remains competitive with several per-slice methods despite never fitting on the test slice itself. On histology-to-gene-expression prediction, the picture is more nuanced. TRM again improves over supervised baselines, but the recursion ablations show that the benefit concentrates in the deep latent–prediction interaction, while repeated latent-only refinement quickly saturates once strong histology foundation-model features are available. Additionally, we see strong generalization benefits, where TRM can zero-shot transfer cross-tissue in domain identification. Furthermore, our experiments delineate what recursive computation contributes to performance. The gains do not appear to arise from effective depth alone: the Universal Transformer baseline, which also reuses weights across many layers, does not match TRM. Instead, the advantage seems to come from the particular factorization of computation into latent-state refinement, deep prediction refinement, and outer-loop iteration, which is specific to TRM. This structured recurrence gives the model a mechanism for progressively refining local evidence with global tissue organization. The test-time ablations further reinforce this view by showing that compute can be traded for quality up to a point, after which the model saturates. That behavior is appealing for spatial biology, where samples vary widely

in complexity and where a fixed computation budget may be unnecessarily rigid. Taken together, these results suggest that recursive reasoning provides strong benefits and improves generalization in spatial omics.

Limitations: We consider only two representative tasks in spatial omics, so the impact of recursive reasoning at the scale of foundation models remains unexplored. Additionally, we propose heuristic solutions to repurpose ACT for continuous domains; inventing smarter approaches for ACT is the next step as our proposed solutions require delicate dataset-specific hyperparameter tuning. The histology-to-expression experiments also use a single foundation-model backbone, leaving robustness across pretrained feature spaces unclear. Finally, the models are deterministic and do not estimate uncertainty, limiting their use in settings where calibrated confidence is important.

References

- Albergaria, A., Paredes, J., Sousa, B., Milanezi, F., Carneiro, V., Bastos, J., Costa, S., Vieira, D., Lopes, N., Lam, E. W., Lunet, N., and Schmitt, F. Expression of FOXA1 and GATA-3 in breast cancer: the prognostic significance in hormone receptor-negative tumours. *Breast Cancer Research*, 11(3):R40, 2009. doi: 10.1186/bcr2327.
- Andersson, A., Larsson, L., Stenbeck, L., Salmén, F., Ehinger, A., Wu, S. Z., Al-Eryani, G., Roden, D., Swarbrick, A., Borg, Å., et al. Spatial deconvolution of her2-positive breast cancer delineates tumor-associated cell type interactions. *Nature communications*, 12(1):6012, 2021.
- Bassi, C., Fortin, J., Snow, B. E., Wakeham, A., Ho, J., Haight, J., You-Ten, A., Cianci, E., Buckler, L., Gorini, C., Stambolic, V., and Mak, T. W. The PTEN and ATM axis controls the G1/S cell cycle checkpoint and tumorigenesis in HER2-positive breast cancer. *Cell Death & Differentiation*, 28(11):3036–3051, 2021. doi: 10.1038/s41418-021-00799-8.
- Chen, W.-T., Lu, A., Craessaerts, K., Pavie, B., Frigerio, C. S., Corthout, N., Qian, X., Laláková, J., Kühnemund, M., Voytyuk, I., et al. Spatial transcriptomics and in situ sequencing to study alzheimer’s disease. *Cell*, 182(4):976–991, 2020.
- Chung, Y., Ha, J. H., Im, K. C., and Lee, J. S. Accurate spatial gene expression prediction by integrating multi-resolution features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11591–11600, 2024.
- Dai, B., Yi, L., Wang, P., Li, H., Hu, P., Song, Y., Xing, J., Feng, Z., Yuan, Z., and Zuo, Y. 3d-ot: a deep geometry-aware framework for heterogeneous slices alignment of spatial multi-omics. *Nature Methods*, pp. 1–12, 2026.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, Ł. Universal transformers. *International Conference on Learning Representations*, 2019.
- Dong, K. and Zhang, S. Deciphering spatial domains from spatially resolved transcriptomics with an adaptive graph attention auto-encoder. *Nature communications*, 13(1):1739, 2022.
- Geiping, J., McLeish, S., Jain, N., Kirchenbauer, J., Singh, S., Bartoldson, B. R., Kailkhura, B., Bhatele, A., and Goldstein, T. Scaling up test-time compute with latent reasoning: A recurrent depth approach. *arXiv preprint arXiv:2502.05171*, 2025.
- Goyal, S., Agrawal, S., Anil, G. G., Jain, P., Paul, S., and Kusupati, A. Elt: Elastic looped transformers for visual generation. *arXiv preprint arXiv:2604.09168*, 2026.
- Graves, A. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.
- Hu, J., Li, X., Coleman, K., Schroeder, A., Ma, N., Irwin, D. J., Lee, E. B., Shinohara, R. T., and Li, M. 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.
- Jia, Y., Liu, J., Chen, L., Zhao, T., and Wang, Y. Thiotogene: a deep learning method for predicting spatial transcriptomics from histological images. *Briefings in Bioinformatics*, 25(1):bbad464, 2024.
- Jolicoeur-Martineau, A. Less is more: Recursive reasoning with tiny networks, 2025. URL <https://arxiv.org/abs/2510.04871>.
- Long, Y., Ang, K. S., Li, M., Chong, K. L. K., Sethi, R., Zhong, C., Xu, H., Ong, Z., Sachaphibulkij, K., Chen, A., et al. Spatially informed clustering, integration, and deconvolution of spatial transcriptomics with graphst. *Nature communications*, 14(1):1155, 2023.
- Maynard, K. R., Collado-Torres, L., Weber, L. M., Uyttingco, C., Barry, B. K., Williams, S. R., Catallini, J. L., Tran, M. N., Besich, Z., Tippani, M., et al. Transcriptome-scale spatial gene expression in the human dorsolateral prefrontal cortex. *Nature neuroscience*, 24(3):425–436, 2021.
- Nielsen, T. O., Leung, S. C. Y., Rimm, D. L., Dodson, A., Acs, B., Badve, S., Denkert, C., Ellis, M. J., Fineberg, S., Flowers, M., Kreipe, H. H., Laenkholm, A.-V., Pan, H., Penault-Llorca, F. M., Polley, M.-Y., Salgado, R., Smith, I. E., Sugie, T., Bartlett, J. M. S., McShane, L. M., Dowsett, M., and Hayes, D. F. Assessment of Ki67 in breast cancer: Updated recommendations from the international Ki67 in breast cancer working group. *Journal of the National Cancer Institute*, 113(7):808–819, 2021. doi: 10.1093/jnci/djaa201.
- Nieto, M. A., Huang, R. Y.-J., Jackson, R. A., and Thiery, J. P. EMT: 2016. *Cell*, 166(1):21–45, 2016. doi: 10.1016/j.cell.2016.06.028.
- Pang, M., Su, K., and Li, M. Leveraging information in spatial transcriptomics to predict super-resolution gene expression from histology images in tumors. *BioRxiv*, pp. 2021–11, 2021.
- Perou, C. M., Sørli, T., Eisen, M. B., van de Rijn, M., Jeffrey, S. S., Rees, C. A., Pollack, J. R., Ross, D. T., Johnsen, H., Akslen, L. A., Fluge, Ø., Pergamenschikov, A., Williams, C., Zhu, S. X., Lønning, P. E., Børresen-Dale, A.-L., Brown, P. O., and Botstein, D. Molecular

- 440 portraits of human breast tumours. *Nature*, 406(6797):
441 747–752, 2000. doi: 10.1038/35021093.
- 442
443 Sahlberg, K. K., Hongisto, V., Edgren, H., Mäkelä, R., Hell-
444 ström, K., Due, E. U., Vollan, H. K. M., Sahlberg, N.,
445 Wolf, M., Børresen-Dale, A.-L., Perälä, M., and Kallion-
446 iemi, O. The HER2 amplicon includes several genes
447 required for the growth and survival of HER2 positive
448 breast cancer cells. *Molecular Oncology*, 7(3):392–401,
449 2013. doi: 10.1016/j.molonc.2012.10.012.
- 450
451 Shiloh, Y. ATM and related protein kinases: safeguarding
452 genome integrity. *Nature Reviews Cancer*, 3(3):155–168,
453 2003. doi: 10.1038/nrc1011.
- 454
455 Sørli, T., Perou, C. M., Tibshirani, R., Aas, T., Geisler, S.,
456 Johnsen, H., Hastie, T., Eisen, M. B., van de Rijn, M.,
457 Jeffrey, S. S., Thorsen, T., Quist, H., Matese, J. C., Brown,
458 P. O., Botstein, D., Lønning, P. E., and Børresen-Dale,
459 A.-L. Gene expression patterns of breast carcinomas
460 distinguish tumor subclasses with clinical implications.
461 *Proceedings of the National Academy of Sciences*, 98(19):
462 10869–10874, 2001. doi: 10.1073/pnas.191367098.
- 463
464 Ståhl, P. L., Salmén, F., Vickovic, S., Lundmark, A.,
465 Navarro, J. F., Magnusson, J., Giacomello, S., Asp, M.,
466 Westholm, J. O., Huss, M., et al. Visualization and analy-
467 sis of gene expression in tissue sections by spatial tran-
468 scriptomics. *Science*, 353(6294):78–82, 2016.
- 469
470 Stein, D., Wu, J., Fuqua, S. A., Roonprapunt, C., Yajnik,
471 V., D’Eustachio, P., Moskow, J. J., Buchberg, A. M., Os-
472 borne, C. K., and Margolis, B. The SH2 domain protein
473 GRB-7 is co-amplified, overexpressed and in a tight com-
474 plex with HER2 in breast cancer. *The EMBO Journal*,
475 13(6):1331–1340, 1994. doi: 10.1002/j.1460-2075.1994.
tb06386.x.
- 476
477 Theodorou, V., Stark, R., Menon, S., and Carroll, J. S.
478 GATA3 acts upstream of FOXA1 in mediating ESR1
479 binding by shaping enhancer accessibility. *Genome Re-
480 search*, 23(1):12–22, 2013. doi: 10.1101/gr.139469.112.
- 481
482 Tobin, N. P. and Bergh, J. Analysis of cyclin D1 in breast
483 cancer: A call to arms. *Current Breast Cancer Reports*, 4
484 (3):171–173, 2012. doi: 10.1007/s12609-012-0083-7.
- 485
486 Xie, R., Pang, K., Chung, S., Perciani, C., MacParland,
487 S., Wang, B., and Bader, G. Spatially resolved gene
488 expression prediction from histology images via bi-modal
489 contrastive learning. *Advances in Neural Information
490 Processing Systems*, 36:70626–70637, 2023.
- 491
492 Xu, H., Usuyama, N., Bagga, J., Zhang, S., Rao, R., Nau-
493 mann, T., Wong, C., Gero, Z., González, J., Gu, Y., et al.
494 A whole-slide foundation model for digital pathology
from real-world data. *Nature*, 630(8015):181–188, 2024.
- Zhu, S., Zhu, Y., Tao, M., and Qiu, P. Diffusion gener-
ative modeling for spatially resolved gene expression
inference from histology images. In *The Thirteenth
International Conference on Learning Representations*,
2025. URL [https://openreview.net/forum?
id=FtjLUHyZAO](https://openreview.net/forum?id=FtjLUHyZAO).

A. The TRM Architecture

The pseudocode follows Figure 3 of Jolicoeur-Martineau (2025) exactly; \cdot we only replace the discrete ACT target ($\hat{y}=y_{\text{true}}$) with a soft loss-proxy target q^* .

```

def latent_recursion(x, y, z, n = 6):
    for i in range(n):
        z = net(x, y, z)
        y = net(y, z)
    return y, z
    latent reasoning
    refine output answer

def deep_recursion(x, y, z, n = 6, T = 3):
    recursing T - 1 times to improve y and z
    no gradients needed
    with torch.no_grad():
        for j in range(T - 1):
            y, z = latent_recursion(x, y, z, n)
            recursing once to improve y and z
            y, z = latent_recursion(x, y, z, n)
        return (y.detach(), z.detach()),
            output_head(y), Q_head(y)

Deep Supervision
for x_input, y_true in train_data_loader:
    y, z = y_init, z_init
    for step in range(N_supervision):
        x = input_embedding(x_input)
        (y, z), y_hat, q_hat = deep_recursion(x, y, z)
        loss = softmax_cross_entropy(y_hat, y_true)
        q_star = sigma((tau - CE_holdout) * s)
        loss += binary_cross_entropy(q_hat, q_star)
        loss.backward()
        opt.step()
        opt.zero_grad()
        if q_hat > 0:
            break
    early-stopping
    
```

Figure 7. TRM training pseudocode. We follow the original TRM recursion structure and replace the discrete ACT target with a soft target q^* .

B. Task Formulation and Data Overview

Datasets

DLPCF (Maynard et al., 2021). This dataset contains 12 dorsolateral prefrontal cortex sections from 3 neurotypical donors, profiled on the $10\times$ Visium platform. Each slide carries $\sim 3,600$ valid under-tissue spots, and each spot is manually annotated by neuropathologists into one of seven cortical regions (white matter plus layers 1–6); spots without a confident layer call are masked out. The dataset also distributes the H&E image accompanying each section.

HER2⁺ ST (Andersson et al., 2021). This dataset contains thirty-six legacy Spatial Transcriptomics sections from eight HER2⁺ breast cancer patients (lettered A–H, with up to six serial sections per patient). Each section is captured on the original ST 1k-spot array at a $100\ \mu\text{m}$ spot

size, yielding ~ 176 – 712 valid spots per slide. Each spot is paired with a corresponding H&E patch from the section’s whole-slide image, which is released alongside the count data. Slide B1 (and every patient’s first section) carries free-form pathologist annotations of breast-cancer histology classes (invasive cancer, breast glands, connective tissue, adipose, undetermined).

Task Definitions

We study two complementary tasks on spatial transcriptomics (ST) data. In both settings, a slide is represented on a discrete spot grid with binary tissue mask $M_s \in \{0, 1\}^{H_s \times W_s}$, and we write $\Omega_s = \{(i, j) : M_s(i, j) = 1\}$ for the valid tissue positions.

Task 1: Cross-slice Spatial Domain I (SDI) identification

In the DLPCF data, each slide provides gene-expression features $X_s \in \mathbb{R}^{H_s \times W_s \times d_{\text{in}}}$ together with layer annotations $Y_s \in \{1, \dots, K\}^{H_s \times W_s}$ on Ω_s , where in our setup $d_{\text{in}} = 50$ PCA components computed from the top 3,000 batch-aware HVGs after $\log_2(\text{counts}+1)$ and $K = 7$ cortical layers. The goal is to learn a mapping

$$f_\theta : (X, M) \mapsto \hat{Y} \in \{1, \dots, K\}^{H \times W}$$

that predicts a domain label at every valid spot. Training minimizes masked per-spot cross-entropy,

$$\mathcal{L}_{\text{cls}}(\theta) = \frac{1}{|\Omega|} \sum_{(i,j) \in \Omega} \text{CE}(f_\theta(X, M)_{i,j}, Y_{i,j}),$$

The input grid is placed on a shared 78×128 canvas, linearly projected to hidden dimension $d=192$, RMS-normalized, and augmented with 2D sinusoidal positional encodings plus a learnable context token. During training, each slice is presented as 20 augmented copies per epoch, with joint transformations applied to the feature grid, label grid, and mask: random 180° rotation, horizontal and vertical flips, Gaussian noise $\mathcal{N}(0, 0.1^2)$ on valid spots, and feature dropout with probability 0.1 per PCA dimension. Validation and test evaluation use the unaugmented grid, and evaluation follows leave-one-slice-out (LOSO): for each held-out slide s^* , the model is trained on all other slides and evaluated zero-shot on s^* .

Task 2: Histology-to-gene-expression prediction (H2G)

For the histology-to-expression task, the input at inference time is only the H&E image. For each valid spot $(i, j) \in \Omega_s$, we extract a 224×224 image patch centered on that spot and embed it with a frozen pathology foundation model ϕ (Gigapath, 1536 dimensions), producing a feature tensor $V_s \in \mathbb{R}^{H_s \times W_s \times d_{\text{fm}}}$. The prediction target is a per-spot gene-expression vector $Y_s(i, j) \in \mathbb{R}^G$ of log-normalized count

Table 6. Full HER2ST B1 holdout metrics under the Zhu et al. (2025) evaluation protocol (300 HMHVG, $\log_2(\text{counts}+1)$, single test slide). [†]Published source data from Zhu et al. (2025).

	PCC-10	PCC-50	PCC-300	MAE	MSE	RVD
No Rec. 2-Layer	0.865	0.820	0.762	0.690	0.872	0.257
Univ. Transformer	0.830	0.792	0.745	0.708	0.822	1.088
TRM ($N_{\text{outer}}=4$)	0.869	0.824	0.771	0.689	0.836	0.383
TRM ($N_{\text{outer}}=8$)	0.883	0.841	0.785	0.678	0.810	0.354
HisToGene [†]	0.6812	0.6345	0.5250	0.9367	1.3468	10.34
BLEEP [†]	0.7727	0.7141	0.5652	0.8328	1.2428	0.6025
TRIPLEX [†]	0.7907	0.7394	0.5766	0.9311	1.3456	0.6428
STEM [†]	0.8298	0.7726	0.5984	0.7547	1.0742	0.0693

values over a fixed gene panel; for HER2ST, we use the top-300 HMHVG panel selected from the training slides, following Zhu et al. (2025). The model learns

$$g_\theta : (V, M) \mapsto \hat{Y} \in \mathbb{R}^{H \times W \times G},$$

trained with masked mean squared error,

$$\mathcal{L}_{\text{reg}}(\theta) = \frac{1}{|\Omega|G} \sum_{(i,j) \in \Omega} \sum_{c=1}^G (g_\theta(V, M)_{i,j,c} - Y_{i,j,c})^2.$$

The spot embeddings are padded to a shared 35×35 grid together with the tissue mask. Following Zhu et al. (2025), we precompute Gigapath embeddings for all eight elements of the dihedral group D_4 and sample one coherent view per slide during training; we additionally use 20 augmented copies per sample per epoch, with Gaussian noise ($\sigma=0.05$) and feature dropout ($p=0.05$) applied to the embedding tensor during training only. Validation and test evaluation use the unaugmented upright view. Evaluation uses a single-slide holdout: one slide is reserved for testing (HER2ST: B1), the model is trained on the remaining slides, and predictions are scored on the held-out valid spots. For each gene c , we compute the across-spot Pearson correlation ρ_c between prediction and target, then report top- k mean Pearson correlation PCC- k for $k \in \{10, 50, 200\}$, along with MSE, MAE, and the relative variation distance (RVD),

$$\text{RVD} = \frac{1}{G} \sum_{c=1}^G \frac{(\sigma_c^{2,\text{pred}} - \sigma_c^{2,\text{gt}})^2}{(\sigma_c^{2,\text{gt}})^2},$$

which measures how well predicted across-spot variance matches the true biological heterogeneity.

Observations about performance on HER2ST. Table 6 shows a detailed comparison between TRM and baselines for the HER2ST dataset. We see that the No-recursion variant already beats the baselines in terms of convergence error and PCC. However, on RVD, Stem’s per-spot diffusion sampling reaches 0.069 on HER2ST B1, an order of magnitude below our best regression result. Two things follow.

(i) *Joint-prediction over the whole slide significantly increases performance over baseline work.* Even without recursion, predicting the entire spot grid jointly under a strong regulariser already exceeds prior regression baselines, which recursion further improves on.

(ii) *The RVD gap is structural, not a tuning issue.* By the law of total variance, $\text{Var}_x(y) = \mathbb{E}_x[\text{Var}(y | x)] + \text{Var}_x(\mathbb{E}[y | x])$. Any MSE-trained deterministic regressor $\hat{y}(x) = \mathbb{E}[y | x]$ emits only the second term and discards the spot-conditional residual $\mathbb{E}_x[\text{Var}(y | x)]$, which dominates on the 302-HMHVG panels and is exactly what RVD scores. PCC- k is invariant to global rescaling, so the same conditional-mean collapse can preserve top- k ranking while flattening variance—hence we beat Stem on PCC- k and lose by an order of magnitude on RVD. Zhu et al. (2025) introduce RVD precisely because “PCC in log-transformed space would be surprisingly high if the prediction is simply the mean expression across all genes,” and motivate diffusion on the prior grounds that the image→expression map is intrinsically one-to-many. Drawing $S \approx 20$ samples per spot from $p_\theta(y | x)$ and averaging restores both variance terms—spot-to-spot diversity preserves $\text{Var}_x(\mathbb{E}[y | x])$, sample-to-sample stochasticity reinjects $\mathbb{E}_x[\text{Var}(y | x)]$ —collapsing the RVD ratio while keeping PCC- k competitive. However, using different ACT halting strategies along with scaling down N_{z_L} and N_{z_H} , we observed that different ACT objectives added variance-stabilizing properties to the final predictions.

C. Ablation Methodology

Procedure. All recursion ablations use models trained with default hyperparameters (deep recursion cycles = 3, latent recursion steps = 6, `halt_max_steps` = 8) in the LOSO (SDI) or held-out-B1 slice (H2G) setting. At test time, we override one or more parameters and evaluate on the held-out slice. For each of the 12 LOSO folds, we evaluate all parameter combinations and report mean \pm std across folds.

Block equivalence. A TRM forward pass is parameterized by three counters, all referring to repeated applications of the shared transformer \mathcal{F}_θ :

- N_{z_L} (*latent recursion, also denoted ℓ*): how many times the latent state z_L is refreshed in a row before the next deep update.
- N_{z_H} (*deep recursion, also denoted h*): how many such latent groups are stacked, each followed by one deep update of z_H .
- N_{outer} (*outer / deep-supervision steps, also denoted n*): how many times the model traverses the entire deep-recursion stack, with a (read-only) outer-step boundary

in between.

A single *processing block* is the basic unit of Figure 1a: N_{z_L} latent updates of z_L (each a call to \mathcal{F}_θ with context z_H+x), followed by one deep update of z_H from z_L . One inner pass executes N_{z_H} such blocks back-to-back; one full forward pass executes N_{outer} inner passes. At test time, the outer-step boundary performs only read-only operations (logit decoding, halt-head evaluation) and does *not* reset or modify the carried (z_H, z_L) state. The sequence of \mathcal{F}_θ calls—and therefore the hidden-state trajectory—is controlled by N_{outer} and N_{z_H} only through their *product*: any two configurations with the same $N_{\text{outer}} \cdot N_{z_H}$ produce a bit-identical test-time forward pass. We verified this empirically: across all 12 folds, configurations with equal $N_{\text{outer}} \cdot N_{z_H}$ produce identical ARI to six decimal places (Table 7).

Table 7. Block equivalence verification (slice 151670). Configurations with the same total blocks ($n \times h$) yield identical ARI.

Config	Total blocks	ARI
$n=1, h=2$	2	0.6769
$n=2, h=1$	2	0.6769
$n=1, h=3$	3	0.7117
$n=3, h=1$	3	0.7117
$n=2, h=3$	6	0.7169
$n=3, h=2$	6	0.7169
$n=6, h=1$	6	0.7169

Grid ablation. For Figure 1a, we evaluate all (n, h) combinations with $n \in \{1, \dots, 8\}$ and $h \in \{1, 2, 3\}$ at default latent recursion steps $\ell=6$, then collapse results by total blocks $n \times h$. For Figure 1b, we fix $n=1$ and vary deep recursion cycles $h \in \{1, 2, 3\}$ with latent recursion steps $\ell \in \{1, 2, 3, 4, 6, 8\}$. For Figure 1c, we use the standard ablation at default $n=3, h=3$ with $\ell \in \{1, \dots, 8\}$.

D. ACT details

Here, we provide details on the ACT recipes, calibration choices, formulae, and results referenced in Section 4, separated by dataset. Section D.1 covers the DLPFC LOSO setting (the original ACT-on-domain-identification analysis); Section D.2 covers the HER2ST gene-expression setting (the regression analog and the input-adaptive halt result). Across all ACT recipes, the halt head is supervised via a per-step, randomized within-sample spot holdout, disjoint from the regression-loss spots. Note that these spots (FM features) are seen by the model during training but are cut off from backpropagation via the regression loss, instead being supervised by the halt head.

Table 8. ACT strategy comparison on LOSO fold 151673 (DLPFC).

Strategy	Test ARI	Mean steps	Best epoch
Loss-Proxy ACT	0.520	8.0	6
Weighted BCE	0.465	2.0	13
Random-K [2,8]	0.457	4.0	100
Random-K [2,16]	0.462	8.0	100

D.1. DLPFC LOSO ACT

Classical ACT (Graves, 2016) was designed for tasks with a clear “solved/not-solved” signal per input, e.g., digit addition, where the model emits one scalar answer it can self-verify. Spatial-domain identification is technically still discrete — each spot is assigned to one of K cortical layers, but the output is a *joint multi-class distribution over all* ~ 4000 spots in the sample, so partial correctness is the rule and a binary correct/not-correct signal per input is essentially never achievable. We compare three practical recipes on slice 151673 LOSO — training on all slices except 151673 and evaluating on it — with results in Table 8.

Loss-Proxy ACT. The halt head is supervised by a soft target derived from held-out cross-entropy. We want the ACT controller to encourage generalization:

$$q_{\text{halt}}^* = \sigma((\tau - \text{CE}_{\text{holdout}}) \cdot s), \quad (1)$$

where $\tau=0.7$ is a quality threshold and $s=4$ controls sharpness. This is also the formulation summarized in the pseudo code in Figure 7.

Weighted BCE. This is a calibrated variant of the same halt-head supervision, with the halt threshold lowered to 0.75 and early-stopping patience relaxed to 8 epochs. We also bias the target to account for class imbalance, thereby promoting good performance across all classes. These changes bias training toward earlier stopping and produce the 2.0 mean-step solution reported in Table 8 and Figure 9.

Random-K. Each training sample is assigned a random target depth $K \sim \text{Uniform}[2, K_{\text{max}}]$, where the model must always take at least two exploratory steps, and the halt head is trained to stop exactly at step K . We sweep $K_{\text{max}} \in \{2, \dots, 8\}$ and also with $K_{\text{max}}=16$ (to make it a fair comparison to the other methods with mean steps 8) to study how the width of the training distribution affects both test ARI and the learned halting point (learned using the random-K values as supervision) (Figure 8).

Loss-Proxy ACT achieves the best ARI (0.520) on a held-out set by optimizing for halt head *emph* answer quality, but it almost always uses the full 8-step budget. Weighted




Figure 8. Random-K sweep over training range $[2, K_{\max}]$ on fold 151673. (a) Test ARI increases with K_{\max} up to ≈ 4 , then plateaus below the Loss-Proxy ACT baseline (dashed). (b) The learned halt point is pinned to K_{\max} for $K_{\max} \leq 4$; for $K_{\max} \geq 5$ the model abandons early halting and runs all 8 steps.




Figure 9. ACT strategy comparison on fold 151673. (a) Test ARI: Loss-Proxy achieves the highest quality. (b) Compute cost: Weighted BCE and Random-K learn to halt early. (c) Convergence: Loss-Proxy is fastest; Random-K is considerably slower.

BCE instead biases the controller toward *good-enough* solutions that account for class imbalance and is the only setting that yields real anytime behavior: mean steps drop to 2.0 (a 75% compute reduction) while retaining moderate quality (ARI 0.465). Random-K, where K corresponds to the choice of N_{outer} , trains over a random depth budget between $[2, K_{\max}]$ and acts as a controller-free regularized alternative; in practice, however, it does not match Loss-Proxy ACT. The key takeaway (Table 8) is that these recipes occupy different points on the quality–compute frontier: Loss-Proxy ACT is best for accuracy and encourages generalization, Weighted BCE offers a practical, compute-saving tradeoff, and Random-K is useful for budget diversification but tricky to train.

D.2. HER2ST ACT: Extended Details

This appendix provides additional mechanism details and the multi-fold halt-step distributions referenced from Section 4. The headline ACT comparison table, halt-recipe definitions, the input-adaptive bimodality result, the cross-fold validation table, and the compute–quality trade-off summary are all in the main paper (Tables 4 and 5, Figure 6). All experiments use the shallow-inner architecture ($N_{z_H}=1$, $N_{z_L}=3$, $K_{\max}=8$) trained on the STEM 2025 HER2ST protocol (Zhu et al., 2025) with 302 HMHVG genes. Across all ACT recipes, the halt head is supervised via a per-step, randomized within-sample spot holdout, disjoint from the regression-loss spots: these spots (FM features) are seen by the model during training but are cut off from backpropaga-




Figure 10. Test-time scaling on the small-recursion architecture trained with random- $K \sim \text{Uniform}[2, 16]$. PCC increases from 0.69 at $N_{\text{outer}}=1$ to 0.72 at $N_{\text{outer}}=4$ before saturating.

tion via the regression loss, instead being supervised by the halt head.

The intuition for scaling down N_{z_H} and N_{z_L} . The default TRM ($N_{z_H}=3$, $N_{z_L}=6$) saturates after a single inner pass: sweeping N_{outer} at test time is flat from step 1. This indicates that the expressivity of the inner recursions saturates performance, meaning more N_{outer} steps cannot improve. We hence scale down the expressivity of the latent ($N_{z_L}=3$) and deep ($N_{z_H}=1$) recursions and, by doing so, observe input-adaptive halting in the continuous setting. To test whether this scaled-down architecture provides a learnable gradient for increasing N_{outer} , we trained a model with a random $K \sim \text{Uniform}[2, 16]$ ACT strategy. This model exhibits a real test-time scaling increase (PCC 0.69 at $N_{\text{outer}}=1 \rightarrow 0.72$ at $N_{\text{outer}}=4$, Figure 10), and hence gives ACT a non-trivial outer-recursion gradient to operate on.

Why slide-level PCC-halt and RVD-halt do not fire.

The model halts at step k if $\sigma(q_{\text{halt}}) > \theta$ for some halt-prob threshold θ (default $\theta=0.7$). PCC-halt’s converged mean halt probability is 0.27 across all 36 HER2ST slides; RVD-halt’s is 0.07. In both cases, the halt head’s output collapses to a value well below θ *uniformly across inputs*, so the controller never crosses threshold, and the model defaults to $K_{\max}=8$ steps on every slide. The reason is a calibration mismatch: PCC-halt’s halt target $\sigma((\text{PCC} - 0.65) \cdot 20)$ saturates near 0.83 at the converged-PCC ceiling ($\overline{\text{PCC}} \approx 0.72$), but the BCE loss is dominated by the early steps where the target is near 0, pulling the halt logit’s converged value down. RVD-halt has the same dynamic with even more skew, since RVD has a wider range and the target sits closer to 1 at convergence than to 0. The remedy is either (i) recalibrate the threshold below the per-input converged distribution (Figure 6c), or (ii) modify the supervision target so the halt head converges higher — both of which we exploit in the main text. Despite never firing, both recipes still cut RVD by $\sim 2\times$ over fixed-depth TRM (Table 4): the halt-loss BCE term acts as a strong variance-aware regularizer on the prediction head whether or not the controller eventually

halts.

Why MSE-halt halts immediately yet wins on RVD.

MSE-halt’s sigmoid $\sigma((\tau - \text{MSE}_k) s)$ at $\tau = 0.9$, $s = 10$ places the midpoint at the converged held-out MSE (~ 0.7 – 1.0), so the supervision target rises sharply across training-time steps from ~ 0 at $k = 1$ to ~ 0.95 at saturation. The halt head therefore converges to mean halt probability 0.997, comfortably above the threshold $\theta = 0.7$ on every input, and halts at the earliest step the configuration allows.

Despite halting at the configured minimum, MSE-halt achieves the lowest slide-level RVD across all four recipes (0.109, $3.2\times$ better than fixed-depth TRM at matched PCC). Two factors account for this. First, MSE-halt’s halt-target trajectory has the largest dynamic range across steps of the three slide-level recipes: PCC saturates at ≈ 0.8 ; RVD-halt’s target saturates at only ≈ 0.5 , since τ sits right at the converged RVD; and MSE-halt’s target saturates near 0.95. Its BCE halt loss is therefore the strongest signal pushing the gradient back through z_H during training.

Second, and more importantly, MSE is the exact same loss used to supervise the prediction head, so the halt loss adds a sample-level view of the same quality signal the prediction head receives at the spot level. This double supervision forces z_H to encode a per-sample coherence signal in addition to per-spot fit, which empirically biases the prediction head toward preserving across-spot variance — exactly what RVD penalizes. PCC-halt and RVD-halt provide weaker BCE gradients, with saturated targets of ≈ 0.8 and ≈ 0.5 , respectively, and use quality estimators that are not the prediction loss. As a result, the same regularization channel is correspondingly weaker.

E. Marker-Gene Panel on HER2ST B1

E.1. HER2ST: Qualitative Marker-Gene Validation

The marker-gene grid in Figure 12 provides a qualitative sanity check on whether the model’s predictions are biologically faithful, rather than merely statistically well-correlated. The HER2-amplicon row supplies the strongest such check: ERBB2 and GRB7 should localize to the same spots within the pathologist-annotated invasive cancer regions, because GRB7 is frequently co-amplified *and co-expressed* with ERBB2/HER2 in breast cancer (Stein et al., 1994; Sahlberg et al., 2013). Both ground-truth and predicted patterns concentrate inside those annotations. The cell-cycle (CCND1, MKI67) row extends this to graded, non-amplicon-driven proliferation signals, with CCND1 and Ki67/MKI67 has been established as a proliferation-associated marker in breast cancer (Tobin & Bergh, 2012; Nielsen et al., 2021). The luminal row (ESR1, GATA3, FOXA1, MUC1) probes more diffuse expression patterns,

since ESR1, GATA3, FOXA1, and the associated epithelial/luminal programs define luminal breast cancer identity (Perou et al., 2000; Sørli et al., 2001; Theodorou et al., 2013; Albergaria et al., 2009). The tumor-suppressor row (ATM, PTEN), grouped because both have established roles in the DNA-damage response in breast cancer (Shiloh, 2003; Bassi et al., 2021), and the EMT row (VIM) (Nieto et al., 2016) probes mesenchymal/EMT-associated signal, which is biologically distinct from the sharply localized HER2-amplicon markers and may be influenced by stromal or mixed-cell composition. We emphasize that the figure is not cherry-picked. The twelve genes were specified a priori from the breast cancer biology before any predictions were inspected, and the checkpoint ($k = 8$) was likewise fixed before this analysis.

F. Limitations

We consider only two representative tasks in spatial omics, so the impact of recursive reasoning at the scale of foundation models remains unexplored. Additionally, we propose heuristic solutions to repurpose ACT for continuous domains; inventing smarter approaches for ACT is the next step as our proposed solutions require delicate dataset-specific hyperparameter tuning. Additionally, the lack of absolute uncertainty or prediction confidence limits the application of the recursive reasoning paradigm to real-world clinical tasks.

770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824

Per-slide halt-step distribution under PCC-halt (thresh = 0.4, shallow TRM) replicates across 4 folds




Figure 11. Per-slide halt-step distributions under per-class PCC-halt (threshold 0.4, shallow TRM) across four independently retrained checkpoints with different held-out test slides. **Top row:** histogram of empirical halt step over all 36 slides, split by train/val/test. **Bottom row:** halt step versus per-slide PCC; the test slide is shown as a red star. Bimodality at $\{\min_halt_steps=2, K_{max}=8\}$ replicates on every fold, and the test slide always lands in the cluster predicted by its difficulty.

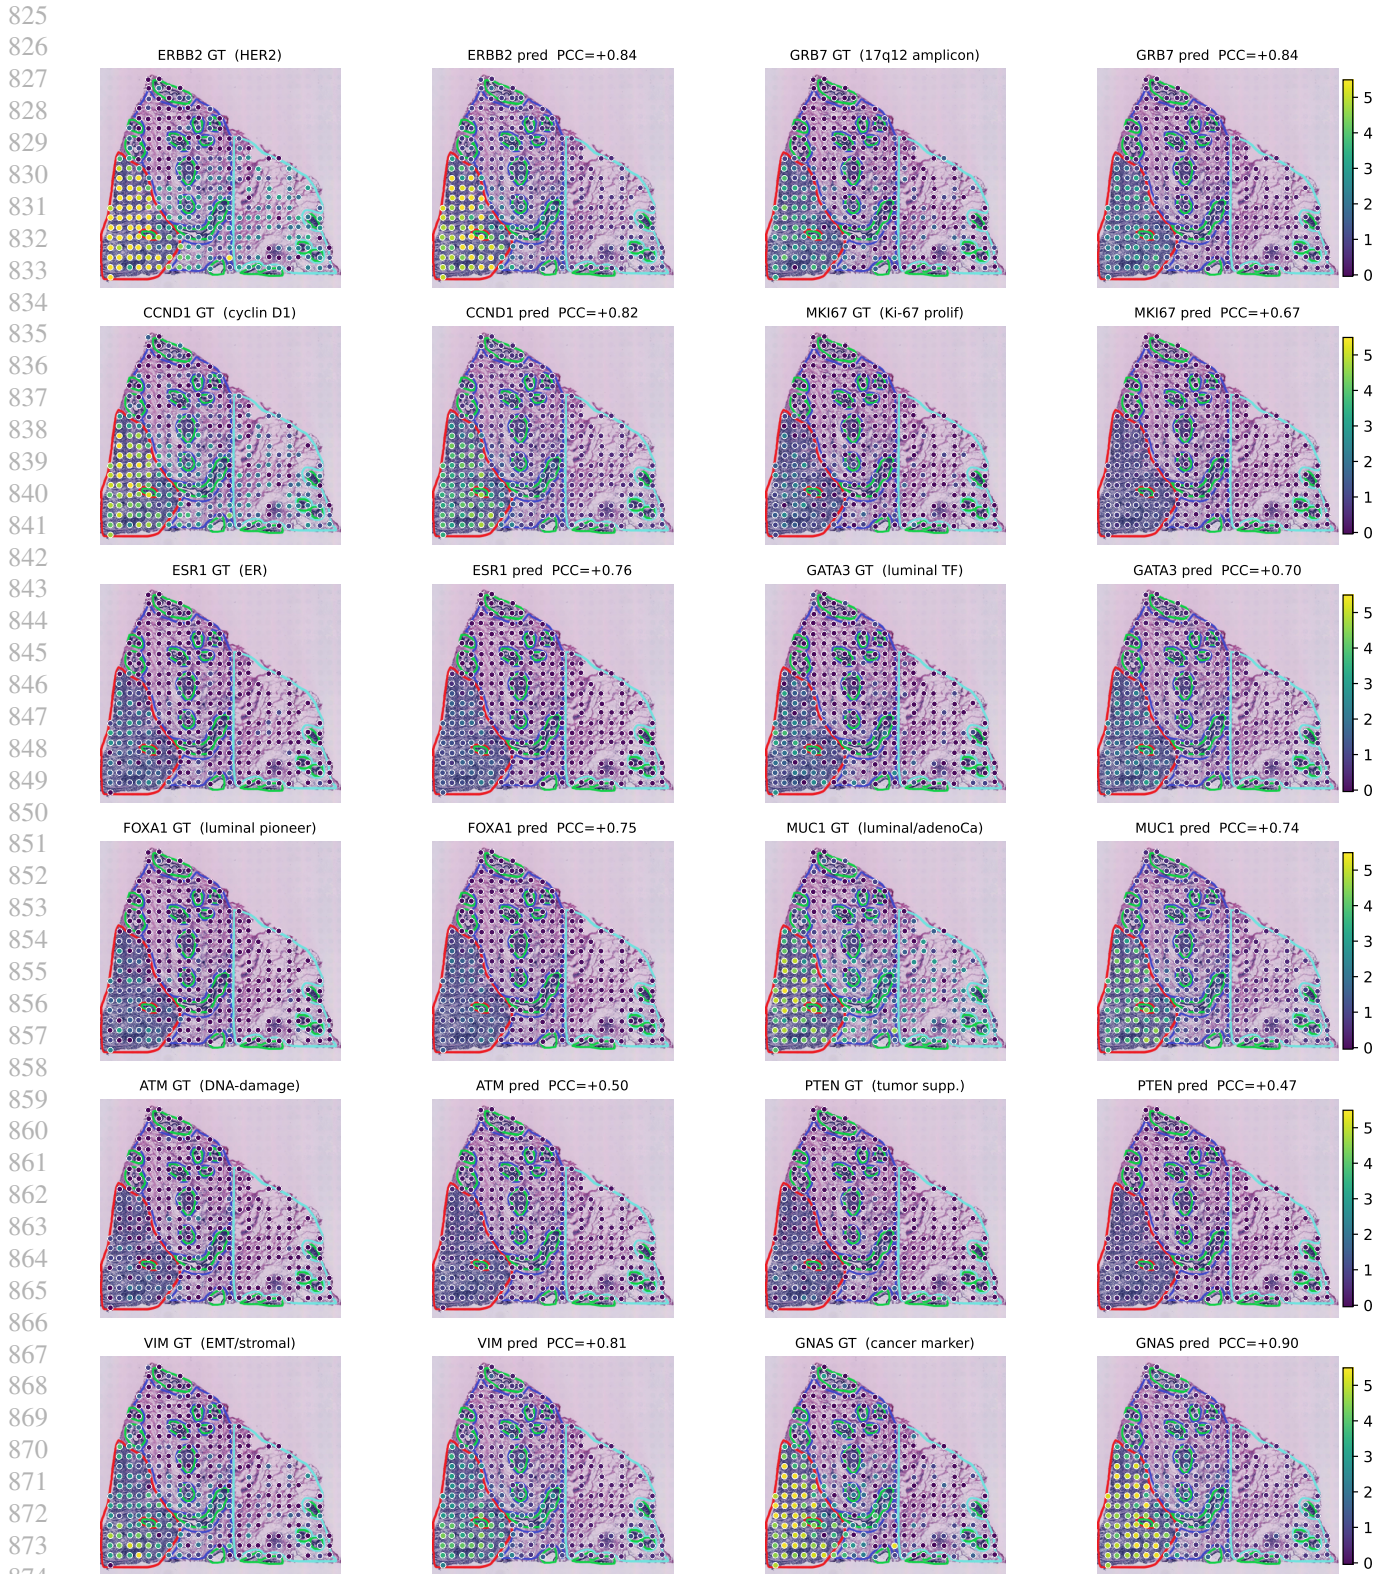


Figure 12. HER2ST B1 marker-gene grid. All expression values are normalized together to allow for direct comparison. Rows are grouped by biological axis: HER2 amplicon (*ERBB2*, *GRB7*); cell cycle (*CCND1*, *MKI67*); hormone receptor / luminal (*ESR1*, *GATA3*); luminal / adenocarcinoma (*FOXA1*, *MUC1*); DNA-damage response/tumor suppressor (*ATM*, *PTEN*); EMT / cancer marker (*VIM*, *GNAS*).