

SCR²-ST: Combine Single Cell with Spatial Transcriptomics for Efficient Active Sampling via Reinforcement Learning

Junchao Zhu¹

Ruining Deng²

Junlin Guo¹

Tianyuan Yao¹

Chongyu Qu¹

Juming Xiong¹

Siqi Lu³

Zhengyi Lu¹

Yanfan Zhu¹

Marilyn Lions¹

Yuechen Yang¹

Yalin Zheng⁴

Yu Wang⁵

Shilin Zhao⁵

Haichun Yang⁵

Yuankai Huo¹

JUNCHAO.ZHU@VANDERBILT.EDU

RUD4004@MED.CORNELL.EDU

JUNLIN.GUO@VANDERBILT.EDU

TIANYUAN.YAO@VANDERBILT.EDU

CHONGYU.QU@VANDERBILT.EDU

JUMING.XIONG@VANDERBILT.EDU

SLU09@WM.EDU

ZHENGYI.LU@VANDERBILT.EDU

YANFAN.ZHU@VANDERBILT.EDU

MARILYN.M.LIONS@VANDERBILT.EDU

YUECHEN.YANG@VANDERBILT.EDU

YALIN.ZHENG@LIVERPOOL.AC.UK

YU.WANG.2@VUMC.ORG

SHILIN.ZHAO.1@VUMC.ORG

HAICHUN.YANG@VUMC.ORG

YUANKAI.HUO@VANDERBILT.EDU

¹ Vanderbilt University, TN, USA

² Weill Cornell Medicine, NY, USA

³ The College of William and Mary, VA, USA

⁴ University of Liverpool, UK

⁵ Vanderbilt University Medical Center, TN, USA

Editors: Under Review for MIDL 2026

Abstract

Spatial transcriptomics (ST) is an emerging technology that enables researchers to investigate the molecular relationships underlying tissue morphology. However, acquiring ST data remains prohibitively expensive, and traditional fixed-grid sampling strategies lead to redundant measurements of morphologically similar or biologically uninformative regions, thus resulting in scarce data that constrain current methods. The well-established single-cell sequencing field, however, could provide rich biological data as an effective auxiliary source to mitigate this limitation. To bridge these gaps, we introduce SCR²-ST, a unified framework that leverages single-cell prior knowledge to guide efficient data acquisition and accurate expression prediction. SCR²-ST integrates a single-cell guided reinforcement learning-based (SCRL) active sampling and a hybrid regression-retrieval prediction network SCR²Net. SCRL combines single-cell foundation model embeddings with spatial density information to construct biologically grounded reward signals, enabling selective acquisition of informative tissue regions under constrained sequencing budgets. SCR²Net then leverages the actively sampled data through a hybrid architecture combining regression-based modeling with retrieval-augmented inference, where a majority cell-type filtering mechanism suppresses noisy matches and retrieved expression profiles serve as soft labels for

auxiliary supervision. We evaluated SCR²-ST on three public ST datasets, demonstrating SOTA performance in both sampling efficiency and prediction accuracy, particularly under low-budget scenarios. Code is publicly available at: <https://github.com/hrblab/SCR2ST>.

Keywords: Computational Pathology, Spatial Transcriptomics, Active Learning




Figure 1: **Comparison between traditional ST sampling and our active sampling.** *Left:* Traditional ST methods rely on fixed-grid sampling regardless of biological importance, leading to redundant measurements in similar regions and inefficient use of sequencing budgets. *Right:* Our proposed approach actively selects informative spots by incorporating single-cell prior knowledge, reducing redundancy while preserving biologically diverse regions.

1. Introduction

Spatial transcriptomics (ST) provides a new perspective for studying the relationship between pathological tissue structures and their spatial gene expression patterns (Burgess, 2019; Asp et al., 2019, 2020). However, acquiring ST data remains relatively expensive (Choe et al., 2023), which together pose challenges for large-scale data collection in practice (He et al., 2020; Zhu et al., 2025d).

Histology features exhibit strong correlations with gene expression patterns (Badea and Stănescu, 2020), providing a foundation for image-based gene expression prediction (He et al., 2020; Zhu et al., 2025a). Deep learning methods have begun leveraging histology images to infer ST expression profiles of each tissue slide (Xie et al., 2024; Yang et al., 2023; Zhu et al., 2025b,c). Representative approaches include regression-based ST-Net (He et al., 2020), HisToGene (Pang et al., 2021), and EGN (Yang et al., 2023), which directly predict expression values from local image appearance; and retrieval-based vision-omics contrastive learning methods, such as BLEEP (Xie et al., 2024) and mlxSTExp (Min et al., 2024).

However, traditional fixed-grid sampling inevitably acquires many spatially adjacent regions with highly similar morphology, leading to substantial molecular redundancy and reduced biological diversity. Its non-selective nature also results in the inclusion of biologically uninformative areas (Schroeder et al., 2025; Grases and Porta-Pardo, 2025). Consequently, the effective information density of the dataset is low, causing a mismatch between

sequencing cost and informative yield, thus constraining the performance and scalability of image-based ST prediction methods.

The limited availability of ST data motivates the integration of external biological knowledge to compensate for inherent constraints in coverage and data quality. In particular, the single-cell sequencing field provides substantially richer priors, supported by large-scale datasets (Regev et al., 2017) and powerful foundation models (Cui et al., 2024), with sample sizes typically exceeding ST by more than an order of magnitude (Svensson et al., 2018). Single-cell profiles resolve cellular types, states, and regulatory programs (Cao et al., 2019; Stuart et al., 2019), offering mechanistic insight into gene expression variation across tissues. Incorporating such fine-grained priors into ST analysis introduces valuable structural guidance and biological constraints, helping mitigate challenges related to limited sampling.

To achieve this, we introduce SCR²-ST, a unified framework that leverages single-cell prior knowledge to guide both efficient data acquisition and expression prediction. Our framework comprises two components. First, we develop a single-cell guided reinforcement learning-based (SCRL) active sampling strategy that jointly leverages single-cell priors and spatial tissue cues to construct a biological reward function, which enables the policy network to adaptively prioritize informative regions while avoiding redundant measurements, maximizing the utility of each sequenced spot under constrained budgets. Within this framework, we further propose a hybrid regression-retrieval prediction network SCR²Net, which integrates regression modeling with retrieval-augmented inference. The retrieval branch aggregates signals from morphologically similar spots, while a majority cell-type filtering mechanism suppresses unreliable matches, balancing global structural learning with context-aware expression transfer. Our contributions can be summarized as fourfold:

- We introduce SCR²-ST, a pioneering and generalizable framework that leverages single-cell prior as an auxiliary source to overcome the scarcity of ST data. It jointly enables efficient data acquisition and accurate expression prediction for ST profile.
- Within this framework, we propose a reinforcement learning-based (SCRL) active sampling strategy that prioritizes informative regions under constrained sequencing budgets through biologically grounded reward signals.
- Building upon this, we develop SCR²Net, a hybrid prediction network that integrates direct regression with retrieved soft label supervision, with a majority cell-type filtering module that suppresses unreliable matches in heterogeneous tissues.
- We provide a systematic benchmark across three public ST datasets under varied sampling budgets, with code and tools released to support reproducible research.

2. Method

2.1. Overall Framework

We introduce SCR²-ST, a unified framework that leverages single-cell prior knowledge to guide efficient data acquisition and accurate expression prediction under limited data budgets, as illustrated in Figure 2. Within SCR²-ST, our single-cell-guided reinforcement

learning-based (SCRL) active sampling strategy integrates single-cell priors with spatial features, using a multi-objective reward function to iteratively drive the policy toward selecting the most informative spots. To fully exploit the abundant single-cell priors in prediction, we design a hybrid regression-retrieval prediction network SCR²Net that fuses direct regression with retrieval-augmented soft label supervision on the actively sampled set.




Figure 2: **Overview of our proposed SCR²-ST framework.** *Left:* Single-cell-guided reinforcement learning-based (SCRL) active sampling strategy that integrates vision features and external single-cell priors to iteratively select informative spots, ensuring efficient data acquisition under limited data budgets. *Right:* Our SCR²Net that infers gene expression from histology images with retrieval-augmented reference. Majority cell-type filtering guided by single-cell priors is applied to suppress unreliable matches in heterogeneous regions.

2.2. Active Sampling via Single-cell Guided Reinforcement Learning

2.2.1. POLICY NETWORK FOR ACTIVE SAMPLING

Prior to active sampling, we perform dense visual feature extraction on tissue sections. Specifically, we uniformly partition WSIs patches and employ the pre-trained UNI (Chen et al., 2024) to extract visual embeddings $\{e_i\}_{i=1}^N$ for each patch, along with their corresponding spatial coordinates $\{(x_i, y_i, w_i)\}_{i=1}^N$, where w_i denotes the slide identifier. Based on these embeddings, we construct a lightweight policy network $\pi_\theta(\cdot)$ that outputs a sampling priority score for each candidate location as

$$\pi_\theta(e_i) = W_2 \cdot \text{ReLU}(W_1 e_i), \quad (1)$$

where $W_1 \in \mathbb{R}^{128 \times d}$ and $W_2 \in \mathbb{R}^{1 \times 128}$ are learnable parameters. The scores are then normalized into a probability distribution via softmax as $p_i = \frac{\exp(\pi_\theta(e_i))}{\sum_{j=1}^N \exp(\pi_\theta(e_j))}$. At iteration

t , the policy network samples k new locations S_t from the unsampled candidate set \mathcal{U}_t according to this probability distribution, and adds them to sampling pool $\mathcal{S} = \bigcup_{\tau \leq t} S_\tau$. The sampling process terminates when the number of samples reaches the total budget B .

2.2.2. MULTI-OBJECTIVE REWARD DESIGN

After obtaining sample set S_t at round t , we evaluate sampling quality and construct multi-objective reward signals to update the policy network. We extract ST expression embeddings $\{\mathbf{z}_i\}_{i \in S_t}$ by pretrained scGPT (Cui et al., 2023) for sampled locations and reference embeddings $\{\mathbf{q}_j\}_{j=1}^M$ from external single-cell data. The reward function measures sampling quality from two complementary perspectives: biological diversity and spatial uniformity.

Single-Cell Prior-Guided Biological Diversity Reward. To quantify how well the sample set explores the single-cell state space, we first apply PCA to reduce the single-cell embeddings $\{\mathbf{q}_j\}$ to 50 dimensions, then cluster them into C latent cell state clusters using MiniBatchKMeans, obtaining the cluster center set $\{\boldsymbol{\mu}_c\}_{c=1}^C$. The coverage reward measures the fraction of clusters reached by the sample set:

$$R_{\text{sc}}(S_t) = \frac{|\{\arg \min_c \|\mathbf{z}_i - \boldsymbol{\mu}_c\|_2 : i \in S_t\}|}{C}, \quad (2)$$

where $|\cdot|$ denotes set cardinality. Each sampled point \mathbf{z}_i is assigned to its nearest cluster, and coverage is computed as the ratio of unique clusters covered to total clusters C .

We then match each ST embedding \mathbf{z}_i to the most similar single-cell embedding via cosine similarity and retrieve the corresponding cell type label to evaluate the cell type diversity of selected samples as:

$$j^*(i) = \arg \max_j \frac{\mathbf{z}_i^\top \mathbf{q}_j}{\|\mathbf{z}_i\| \|\mathbf{q}_j\|}, \quad (3)$$

We compute the cell type distribution in the sample set as $P(k) = |\{i \in S_t : \text{type}(j^*(i)) = k\}|/|S_t|$, and define the diversity reward based on normalized entropy:

$$R_{\text{type}}(S_t) = \frac{-\sum_k P(k) \log(P(k) + \varepsilon)}{\log(K + \varepsilon)}, \quad (4)$$

where K is the number of cell types observed in the sample set and ε is a small constant for numerical stability. More uniform type distributions yield higher rewards, encouraging preferential sampling of regions with greater cellular heterogeneity.

Spatial Distribution Diversity Reward. Spatial distribution of sampled points also affects information density. An ideal sampling strategy should balance two objectives: (1) spatial dispersion to avoid over-clustering in local regions; (2) uniform coverage to ensure that unsampled locations have nearby reference points. Therefore, we define dispersion D_{disp} as the average pairwise distance among sampled points, where larger values indicate better dispersion. We define coverage D_{cover} as the average distance from all candidate locations to their nearest sampled point, where smaller values indicate more uniform coverage:

$$D_{\text{disp}}(S_t) = \frac{1}{|S_t|^2} \sum_{i,j \in S_t} \|(x_i, y_i) - (x_j, y_j)\|_2, \quad D_{\text{cover}}(S_t) = \frac{1}{N} \sum_{i=1}^N \min_{j \in S_t} \|(x_i, y_i) - (x_j, y_j)\|_2. \quad (5)$$

The spatial distribution diversity reward combines both metrics:

$$R_{\text{spa}}(S_t) = \frac{D_{\text{disp}}(S_t) + D_{\text{cover}}(S_t)}{2}. \quad (6)$$

2.2.3. COMBINED REWARD AND POLICY OPTIMIZATION

We linearly combine the three reward components into a composite signal:

$$R(S_t) = w_{\text{sc}} \cdot R_{\text{sc}}(S_t) + w_{\text{type}} \cdot R_{\text{type}}(S_t) + w_{\text{spa}} \cdot R_{\text{spa}}(S_t), \quad (7)$$

where w_{sc} , w_{type} , and w_{spa} control the relative contributions of single-cell manifold coverage, cell type diversity, and spatial distribution diversity, respectively. We then update the policy network parameters using the composite reward:

$$\nabla_{\theta} \mathcal{J} = \mathbb{E}_{S_t \sim \pi_{\theta}} [R(S_t) \cdot \nabla_{\theta} \log \pi_{\theta}(S_t)], \quad (8)$$

where \mathcal{J} is the expected cumulative reward. Through gradient ascent optimization, the policy network progressively learns to balance biological diversity and spatial uniformity, steering the sampling strategy toward more informative tissue regions.

2.3. SCR²Net: Single-Cell Guided Regression-Retrieval Network

To further leverage single-cell prior knowledge, we design SCR²Net with two complementary paths, including a direct regression path for image-to-expression mapping, and a retrieval-augmented path that provides soft supervision by retrieving similar samples from the training set as an external knowledge base.

2.3.1. SINGLE-CELL GUIDED RETRIEVAL MODULE

Cross-Modality Alignment. Direct regression alone struggles to capture complex expression patterns under limited training samples. To address this, we introduce a retrieval-augmented module that treats the training set as an external memory bank encoding single-cell knowledge, providing soft supervision for the regression pathway.

We design two projection heads with identical architecture to map image features f_{img} from the visual encoder and gene expression embeddings into a shared semantic space. An InfoNCE loss \mathcal{L}_{con} is applied to align vision-omics representations and update the projection head. We then compute cosine similarity between the query image and reference samples:

$$\text{sim}(f_{\text{img}}, y_j) = \frac{\phi_{\text{img}}(f_{\text{img}})^{\top} \phi_{\text{expr}}(y_j)}{\|\phi_{\text{img}}(f_{\text{img}})\| \|\phi_{\text{expr}}(y_j)\|}, \quad (9)$$

where ϕ_{img} and ϕ_{expr} denote the image and expression projection heads, respectively. We select the top- K most similar samples to construct the reference set.

Cell-Type-Aware Filtering and Knowledge Distillation. Expression patterns vary significantly across cell types in ST data, while vision representation could resemble, thus directly aggregating all retrieved samples may introduce noise. To ensure biological consistency, we introduce a majority cell-type filtering mechanism. We count cell type distribution among the top- K samples and retain only those belonging to the T most frequent cell types.

The mean expression of filtered samples serves as the retrieved soft label $\hat{y}_{\text{ret}} = \frac{1}{|\mathcal{R}|} \sum_{j \in \mathcal{R}} y_j$, where \mathcal{R} denotes the filtered retrieval set.

To account for retrieval quality, we introduce a similarity-based confidence mask m , where higher retrieval similarity leads to greater weight on the distillation loss. The retrieved prediction \hat{y}_{ret} then guides the regression path through knowledge distillation, with loss function \mathcal{L}_{ret} weighted by a hyperparameter λ_{KD} denoted as:

$$\mathcal{L}_{\text{ret}} = \lambda_{\text{KD}} \cdot m \cdot \|\hat{y} - \hat{y}_{\text{ret}}\|^2 \quad (10)$$

2.3.2. REGRESSION PATH AND TRAINING OBJECTIVE

We adopt DenseNet-121 (Huang et al., 2017) pre-trained on ImageNet as the visual encoder to capture histomorphological patterns. Given an input patch, the encoder produces a compact feature vector through global average pooling. A two-layer MLP then decodes the visual features into gene expression predictions \hat{y} , forming the direct regression path.

To supervise the regression prediction, we employ two complementary losses. The MSE loss $\mathcal{L}_{\text{reg}} = \|y - \hat{y}\|^2$ directly minimizes the difference between predictions and ground truth, while a Pearson Correlation Coefficient (PCC) loss $\mathcal{L}_{\text{pcc}} = 1 - \text{PCC}(y, \hat{y})$ to capture the correlation structure across genes, which is important for preserving gene-spatial relationships. The total loss integrates direct supervision from both regression losses and soft supervision from the retrieval-based distillation with hyperparameters λ_r and λ_p , denoted as:

$$\mathcal{L} = \lambda_r \cdot \mathcal{L}_{\text{reg}} + \lambda_p \cdot \mathcal{L}_{\text{pcc}} + \mathcal{L}_{\text{ret}}, \quad (11)$$

3. Data and Experiments

Datasets and Preprocessing. We evaluate all methods using three public ST datasets, including HER2 (Andersson et al., 2021), Breast Cancer (He et al., 2020), and Kidney (Lake et al., 2023). For each spot, we cropped a 224×224 pixel patch centered on spatial coordinates as model input. We selected top 300 genes with highest average variance as prediction targets. Following BLEEP (Xie et al., 2024), we applied a $\log(1+x)$ transformation on raw readouts. For the external single-cell datasets, we use two million cells from (Lake et al., 2025) as the reference for the Kidney dataset, and around three million cells from (Chen et al., 2025; Reed et al., 2024; Klughammer et al., 2024) as the reference for the Breast Cancer and HER2 datasets. A detailed profile of datasets is provided in Appendix A.

Baseline. We compared our model against SOTA methods, including regression-based models ST-Net (He et al., 2020), EGN (Yang et al., 2023), HisToGene (Pang et al., 2021), His2ST (Zeng et al., 2022), and TRIPLEX (Chung et al., 2024), and retrieval-based models BLEEP (Xie et al., 2024) and mclSTExp (Min et al., 2024). All methods were trained and evaluated under consistent experimental settings to ensure fair comparison. To validate our sampling strategy, we select Monte Carlo random sampling, uncertainty-based sampling (Safaei et al., 2024), and diversity-driven sampling (Zhdanov, 2019) for comparison.

Evaluation Metrics. We employed Pearson correlation coefficient (PCC), mean squared error (MSE), and mean absolute error (MAE) to comprehensively assess model performance in gene expression prediction from both spatial correlation and error perspectives.

Implementation Details. All experiments were conducted on a single NVIDIA RTX A6000 GPU. We employed SGD optimizer with momentum of 0.9 and weight decay of

10^{-4} . The initial learning rate was set to $lr_0 = 10^{-4}$, with a cosine annealing schedule that gradually decays the learning rate to 10^{-6} . The training batch size was set to 256. Details of experimental implementation and hyperparameter settings are listed in Appendix B.




Figure 3: **Comparison of sampling strategies across multiple methods with metrics reported (A) MSE and (B) PCC.** We evaluate performance under data ratios of 10%–75% to simulate low-resource scenarios, with each curve denoting a sampling strategy. Our SCRL sampling consistently achieves better performance across datasets and models, particularly in low-budget (10%–25%) regimes.

4. Results

4.1. Active Sampling under Budget Constraints

To validate the effectiveness of our active sampling strategy via single-cell guided reinforcement learning (SCRL), we conducted a systematic evaluation across methods. As shown in Figure 3 and Figure 6 in Appendix C, we compared four sampling strategies under training data ratios ranging from 10% to 75%. Experimental results demonstrate that SCRL sampling achieves optimal performance across all datasets and model combinations, with particularly advantages in low-budget scenarios (10%–25%). For Breast Cancer dataset at a 10% sampling ratio, SCRL sampling reduces the MSE from approximately 0.85 to 0.75 and improves the PCC from 0.04 to 0.14 on ST-Net compared to random sampling. This trend holds consistently across other datasets and diverse methods.

Notably, we observed differential sensitivity to sampling strategies across model types. Retrieval-based models exhibit greater sensitivity to data quality compared to end-to-end

regression models, resulting in larger performance gaps between different sampling strategies. This can be attributed to the nature of contrastive learning, which is highly dependent on the quality and diversity of training samples. Our SCRL sampling strategy balances biological quality and diversity. Specifically, single-cell references ensure that sampled spots cover critical cell subpopulations, while spatial density information guides the sampling process to preserve morphological diversity. This dual constraint enables SCRL to achieve stable and consistent performance improvements across both training paradigms.

Table 1: **Performance comparison on gene expression prediction task.** The best performance is highlighted in **orange** and second highest in **blue**, where we can observe that SCR²Net outperforms the SOTAs across most metrics on most datasets.

Model	Breast Cancer			HER2			Kidney		
	MSE ↓	MAE ↓	PCC ↑	MSE ↓	MAE ↓	PCC ↑	MSE ↓	MAE ↓	PCC ↑
ST-Net	0.6318	0.6377	0.1592	0.9237	0.7559	0.2709	0.7460	0.6811	0.1851
His2ST	0.6999	0.6682	0.0612	0.9928	0.8034	0.1045	0.7912	0.7080	0.0571
HisToGene	0.6521	0.6486	0.1149	0.9702	0.8050	0.1392	0.8540	0.7373	0.1134
EGN	0.6662	0.6558	0.1462	0.8916	0.7640	0.2524	0.7574	0.6864	0.1632
TRIPLEX	0.6672	0.6590	0.1093	0.9356	0.7752	0.2167	0.7168	0.6692	0.0930
BLEEP	0.6266	0.6044	0.2041	0.9507	0.7613	0.2834	0.8246	0.7167	0.2020
mc1STExp	0.6472	0.6202	0.1645	0.8882	0.7367	0.2651	0.7438	0.6759	0.1580
SCR ² Net (Ours)	0.5848	0.5725	0.1940	0.9139	0.7042	0.3028	0.7038	0.6611	0.2391

4.2. Empirical Validation on Gene Expression Prediction

We conducted four-fold cross-validation at the sample level to validate SCR²Net against SOTAs. Table 1 summarizes quantitative comparisons across different cohorts. Our SCR²Net outperforms existing methods in almost all metrics, achieving the lowest MSE and MAE ↓ as well as the highest PCC on most datasets. For example, on the Kidney dataset, SCR²Net improves PCC by a clear margin to 0.2391, compared with prior baselines of 0.2020. Furthermore, as illustrated in Figure 5 in Appendix C, SCR²Net maintains strong predictive robustness under varying sampling budgets with our SCRL sampling strategy. The performance gap is pronounced under low sampling ratios (10–25%), where other approaches suffer degradation due to sparse tissue coverage. In contrast, SCR²Net mitigates this by acquiring informative tissue regions and leveraging retrieval-based auxiliary priors, resulting in a more reliable predictive performance with reduced sequencing costs.

4.3. Ablation Study

Reward Function in SCRL sampling. We conducted an ablation analysis on the Biological Prior Reward and the Spatial Density Reward. As shown in Figure 7 in Appendix C, when using only Biological Reward, model performs well at low sampling ratios (10% and 25%); however, as the ratio increases, redundant samples limit further performance gains. Conversely, using only Spatial Reward yields results similar to random sampling, as it cannot directly assess the informativeness of samples. Combining both rewards ensures both biological quality and diversity and allows the model to achieve optimal performance with a performance curve exhibiting a stable upward trend as the sampling ratio increases.

Table 2: **Ablation study and hyperparameter sensitivity analysis in SCR²Net**, where SCR²Net achieves the optimal results with all blocks, and a moderate hyperparameter setting provides the best balance between noise and information.

Functional Block & Setting		Breast Cancer			Kidney		
		MSE ↓	MAE ↓	PCC ↑	MSE ↓	MAE ↓	PCC ↑
w.o. Retrieval Reference Module		0.6318	0.6377	0.1592	0.7460	0.6811	0.1851
w.o. Cell Type Filtering		0.6032	0.5992	0.1786	0.6952	0.6531	0.2235
w. All functional blocks		0.5848	0.5725	0.1940	0.7038	0.6611	0.2391
Retrieval Module	$K = 10, T = 3$	0.6079	0.6198	0.1680	0.7236	0.6814	0.2187
	$K = 20, T = 5$	0.5912	0.6059	0.1711	0.6886	0.6643	0.2225
	$K = 50, T = 10$	0.5848	0.5725	0.1940	0.7038	0.6611	0.2391
	$K = 100, T = 20$	0.5848	0.5925	0.1839	0.7239	0.6712	0.1977
Confidence Mask	$m = 0.05$	0.6052	0.6035	0.1743	0.7135	0.6659	0.2102
	$m = 0.15$	0.5848	0.5725	0.1940	0.7038	0.6611	0.2391
	$m = 0.35$	0.6233	0.6281	0.1590	0.7465	0.7005	0.1920

Functional Blocks in SCR²Net. As shown in Table 2, removing the Retrieval Reference Module leads to an increase in MSE ↓ from 0.7038 to 0.7460 and a decrease in PCC ↑ from 0.2391 to 0.1851 on the Kidney dataset, which indicates its effectiveness in providing reference priors by incorporating similar spots. Meanwhile, the majority cell type filtering mechanism suppresses interference from noisy references by excluding low-quality retrieved spots, thereby enhancing the biological relevance and quality of retrieval references.

Sensitivity Analysis of Hyperparameters. We tested different combinations of candidate pool size K , retained cell types T , and confidence mask threshold m for the retrieval module. Results in Table 2 indicate that overly small values of K and T (e.g., $K = 10$, $T = 3$) or a higher threshold m limit the richness of reference information and reduce the number of effective reference spots. Conversely, overly large settings of K and T or a lower threshold fail to effectively filter noisy matches, leading to performance degradation. Therefore, moderate hyperparameter settings achieve the optimal trade-off between suppressing noisy matches and preserving informative retrieval references.

5. Conclusion

We propose SCR²-ST, a unified framework that bridges single-cell prior knowledge with ST to enable efficient data acquisition and accurate expression prediction. Moving beyond traditional fixed-grid sampling, SCRL constructs biologically grounded reward signals by integrating single-cell prior knowledge with spatial density cues, guiding the policy network to prioritize informative regions while avoiding redundant measurements. SCR²Net then fuses direct regression with retrieval-augmented inference, where cell-type-aware filtering suppresses noise and retrieved soft labels regularize the learning process. Extensive experiments on public datasets validate our framework’s superior performance across diverse prediction architectures, with notable gains under low-budget constraints. By unifying active sampling and hybrid prediction within a single-cell-empowered paradigm, SCR²-ST offers a scalable solution for efficient spatial transcriptomics modeling and establishes a foundation for future research in budget-aware biomedical data acquisition.

Acknowledgments

This research was supported by NIH R01DK135597 (Huo), DoD HT9425-23-1-0003 (HCY), NSF 2434229 (Huo), and KPMP Glue Grant. This work was also supported by Vanderbilt Seed Success Grant, Vanderbilt Discovery Grant, and VISE Seed Grant. This project was supported by The Leona M. and Harry B. Helmsley Charitable Trust grant G-1903-03793 and G-2103-05128. This research was also supported by NIH grants R01EB033385, R01DK132338, REB017230, R01MH125931, and NSF 2040462. We extend gratitude to NVIDIA for their support by means of the NVIDIA hardware grant. This work was also supported by NSF NAIRR Pilot Award NAIRR240055.

References

- Alma Andersson, Ludvig Larsson, Linnea Stenbeck, Fredrik Salmén, Anna Ehinger, Sunny Z Wu, Ghamdan Al-Eryani, Daniel Roden, Alex Swarbrick, Åke Borg, et al. Spatial deconvolution of her2-positive breast cancer delineates tumor-associated cell type interactions. *Nature communications*, 12(1):6012, 2021.
- Michaela Asp, Stefania Giacomello, Ludvig Larsson, Chenglin Wu, Daniel Fürth, Xiaoyan Qian, Eva Wärdell, Joaquin Custodio, Johan Reimegård, Fredrik Salmén, et al. A spatiotemporal organ-wide gene expression and cell atlas of the developing human heart. *Cell*, 179(7):1647–1660, 2019.
- Michaela Asp, Joseph Bergensträhle, and Joakim Lundeberg. Spatially resolved transcriptomes—next generation tools for tissue exploration. *Bioessays*, 42(10):1900221, 2020.
- Liviu Badea and Emil Stănescu. Identifying transcriptomic correlates of histology using deep learning. *PloS one*, 15(11):e0242858, 2020.
- Darren J Burgess. Spatial transcriptomics coming of age. *Nature Reviews Genetics*, 20(6):317–317, 2019.
- Junyue Cao, Malte Spielmann, Xiaojie Qiu, Xingfan Huang, Daniel M Ibrahim, Andrew J Hill, Fan Zhang, Stefan Mundlos, Lena Christiansen, Frank J Steemers, et al. The single-cell transcriptional landscape of mammalian organogenesis. *Nature*, 566(7745):496–502, 2019.
- Andrew Chen, Lina Kroehling, Christina S Ennis, Gerald V Denis, and Stefano Monti. A highly resolved integrated transcriptomic atlas of human breast cancers. *bioRxiv*, pages 2025–03, 2025.
- Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Bowen Chen, Andrew Zhang, Daniel Shao, Andrew H Song, Muhammad Shaban, et al. Towards a general-purpose foundation model for computational pathology. *Nature Medicine*, 2024.
- Kyongho Choe, Unil Pak, Yu Pang, Wanjun Hao, and Xiuqin Yang. Advances and challenges in spatial transcriptomics for developmental biology. *Biomolecules*, 13(1):156, 2023.

- Youngmin Chung, Ji Hun Ha, Kyeong Chan Im, and Joo Sang Lee. Accurate spatial gene expression prediction by integrating multi-resolution features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11591–11600, 2024.
- Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, and Bo Wang. scgpt: Towards building a foundation model for single-cell multi-omics using generative ai. *bioRxiv*, 2023.
- Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. scgpt: toward building a foundation model for single-cell multi-omics using generative ai. *Nature methods*, 21(8):1470–1480, 2024.
- Daniela Grases and Eduard Porta-Pardo. A practical guide to spatial transcriptomics: lessons from over 1000 samples. *Trends in Biotechnology*, 2025.
- Bryan He, Ludvig Bergenstr hle, Linnea Stenbeck, Abubakar Abid, Alma Andersson,  ke Borg, Jonas Maaskola, Joakim Lundberg, and James Zou. Integrating spatial gene expression and breast tumour morphology via deep learning. *Nature biomedical engineering*, 4(8):827–834, 2020.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- Johanna Klughammer, Daniel L Abravanel,  sa Segerstolpe, Timothy R Blosser, Yury Goltsev, Yi Cui, Daniel R Goodwin, Anubhav Sinha, Orr Ashenberg, Michal Slyper, et al. A multi-modal single-cell and spatial expression map of metastatic breast cancer biopsies across clinicopathological features. *Nature Medicine*, 30(11):3236–3249, 2024.
- Blue B Lake, Rajasree Menon, Seth Winfree, Qiwen Hu, Ricardo Melo Ferreira, Kian Kalhor, Daria Barwinska, Edgar A Otto, Michael Ferkowicz, Dinh Diep, et al. An atlas of healthy and injured cell states and niches in the human kidney. *Nature*, 619(7970):585–594, 2023.
- Blue B Lake, Ricardo Melo Ferreira, Jens Hansen, Rajasree Menon, Jeannine Basta, Heather Thiessen Philbrook, Stephanie Reinert, Robin Fallegger, Asmita K Lagwankar, Xi Chen, et al. Cellular and spatial drivers of unresolved injury and functional decline in the human kidney. *BioRxiv*, pages 2025–09, 2025.
- Wenwen Min, Zhiceng Shi, Jun Zhang, Jun Wan, and Changmiao Wang. Multimodal contrastive learning for spatial gene expression prediction using histology images. *Briefings in Bioinformatics*, 25(6):bbae551, 2024.
- Minxing Pang, Kenong Su, and Mingyao Li. Leveraging information in spatial transcriptomics to predict super-resolution gene expression from histology images in tumors. *BioRxiv*, pages 2021–11, 2021.

- Austin D Reed, Sara Pensa, Adi Steif, Jack Stenning, Daniel J Kunz, Linsey J Porter, Kui Hua, Peng He, Alecia-Jane Twigger, Abigail JQ Siu, et al. A single-cell atlas enables mapping of homeostatic cellular shifts in the adult human breast. *Nature genetics*, 56(4): 652–662, 2024.
- Aviv Regev, Sarah A Teichmann, Eric S Lander, Ido Amit, Christophe Benoist, Ewan Birney, Bernd Bodenmiller, Peter Campbell, Piero Carninci, Menna Clatworthy, et al. The human cell atlas. *elife*, 6:e27041, 2017.
- Bardia Safaei, VS Vibashan, Celso M De Melo, and Vishal M Patel. Entropic open-set active learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 4686–4694, 2024.
- Amelia Schroeder, Melanie L Loth, Chunyu Luo, Sicong Yao, Hanying Yan, Daiwei Zhang, Sarbottam Piya, Edward Plowey, Wenxing Hu, Jean R Clemenceau, et al. Scaling up spatial transcriptomics for large-sized tissues: uncovering cellular-level tissue architecture beyond conventional platforms with iscale. *Nature methods*, pages 1–12, 2025.
- Tim Stuart, Andrew Butler, Paul Hoffman, Christoph Hafemeister, Efthymia Papalexi, William M Mauck, Yuhan Hao, Marlon Stoeckius, Peter Smibert, and Rahul Satija. Comprehensive integration of single-cell data. *cell*, 177(7):1888–1902, 2019.
- Valentine Svensson, Roser Vento-Tormo, and Sarah A Teichmann. Exponential scaling of single-cell rna-seq in the past decade. *Nature protocols*, 13(4):599–604, 2018.
- Ronald Xie, Kuan Pang, Sai Chung, Catia Perciani, Sonya MacParland, Bo Wang, and Gary Bader. Spatially resolved gene expression prediction from histology images via bi-modal contrastive learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Yan Yang, Md Zakir Hossain, Eric A Stone, and Shafin Rahman. Exemplar guided deep neural network for spatial transcriptomics analysis of gene expression prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5039–5048, 2023.
- Yuansong Zeng, Zhuoyi Wei, Weijiang Yu, Rui Yin, Yuchen Yuan, Bingling Li, Zhonghui Tang, Yutong Lu, and Yuedong Yang. Spatial transcriptomics prediction from histology jointly through transformer and graph neural networks. *Briefings in Bioinformatics*, 23(5):bbac297, 2022.
- Fedor Zhdanov. Diverse mini-batch active learning. *arXiv preprint arXiv:1901.05954*, 2019.
- Junchao Zhu, Ruining Deng, Junlin Guo, Tianyuan Yao, Siqu Lu, Chongyu Qu, Juming Xiong, Yanfan Zhu, Zhengyi Lu, Yuechen Yang, et al. Computer vision methods for spatial transcriptomics: A survey. *bioRxiv*, pages 2025–10, 2025a.
- Junchao Zhu, Ruining Deng, Junlin Guo, Tianyuan Yao, Juming Xiong, Chongyu Qu, Mengmeng Yin, Yu Wang, Shilin Zhao, Haichun Yang, et al. Img2st-net: efficient high-resolution spatial omics prediction from whole-slide histology images via fully convolutional image-to-image learning. *Journal of Medical Imaging*, 12(6):061410–061410, 2025b.

Junchao Zhu, Ruining Deng, Tianyuan Yao, Juming Xiong, Chongyu Qu, Junlin Guo, Siqu Lu, Yucheng Tang, Daguang Xu, Mengmeng Yin, et al. Magnet: Multi-level attention graph network for predicting high-resolution spatial transcriptomics. *arXiv preprint arXiv:2502.21011*, 2025c.

Junchao Zhu, Ruining Deng, Tianyuan Yao, Juming Xiong, Chongyu Qu, Junlin Guo, Siqu Lu, Mengmeng Yin, Yu Wang, Shilin Zhao, et al. Assign: an anatomy-aware spatial imputation graph network for 3d spatial transcriptomics. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 30829–30838, 2025d.

Appendix A. Additional Details of Dataset

We evaluate all methods on three public spatial transcriptomics (ST) datasets: HER2 (Andersson et al., 2021), Breast Cancer (He et al., 2020), and Kidney (Lake et al., 2023). The HER2 dataset comprises 8 patient samples with 36 WSIs and 13,620 spatial spots in total. The Breast Cancer dataset consists of 23 samples with 68 WSIs and 30,066 spots. The Kidney dataset includes 22 samples with 23 WSIs and 25,944 spots. The spot diameter for HER2 and Breast Cancer is 100 μm , whereas the Kidney dataset adopts a smaller 55 μm spot diameter. For the external single-cell datasets, we use two million cells from (Lake et al., 2025) as the reference for the Kidney dataset, and around three million cells from (Chen et al., 2025; Reed et al., 2024; Klughammer et al., 2024) as the reference for the Breast Cancer and HER2 datasets.

For each spot location, we extracted a 224×224 pixel histology patch centered on its spatial coordinate as model input. To construct the prediction targets, we selected the top 300 genes with the highest variance in expression within each dataset. Following BLEEP (Xie et al., 2024), we applied a $\log(1+x)$ transformation to the raw count matrices to alleviate the heavy-tailed distribution characteristic of ST expression data (He et al., 2020). The dataset-specific selected genes are visualized in Appendix Figure 4.

Dataset	Selected Genes
HER2	'IGKC', 'IGHG3', 'IGLC2', 'MUC1', 'IGHM', 'IGHA1', 'ERBB2', 'IGLC3', 'CALML5', 'IGHG4', 'PTPRF', 'FNL', 'ACTG1', 'PSMB3', 'MIEN1', 'TPT1', 'S100A9', 'CISD3', 'BEST1', 'COL1A2', 'IGHG1', 'FASN', 'COL1A1', 'ACTB', 'B2M', 'TFE3', 'GRB7', 'APOE', 'FTH1', 'PCGF2', 'MGP', 'PABPC1', 'SERF2', 'PSMD3', 'COL3A1', 'S100A8', 'DDX5', 'NDUFB9', 'HLA-B', 'S100A6', 'GNAS', 'TMSB10', 'EEF2', 'ADAM15', 'MLL16', 'KRT18', 'AZGP1', 'TRAF4', 'TUBA1B', 'CD24', 'LASP1', 'S100A14', 'SCD', 'KRT7', 'STARD3', 'TAGLN2', 'PERP', 'AEBP1', 'CD74', 'PTMS', 'HLA-DRA', 'CRIP2', 'IFI27', 'PPP1R1B', 'VMP1', 'PNT1', 'FTL', 'COX7C', 'GNB2', 'EEF1D', 'MYL6', 'F4H8', 'FNB1L', 'GAPDH', 'FAU', 'SPINT2', 'PRDX1', 'ATPIA1', 'SYNGR2', 'XBP1', 'CCT3', 'MMP14', 'HNRNP2B1', 'COX5C', 'ORMDL3', 'EIF4G1', 'CALR', 'PSAP', 'KRT18', 'PLXNB2', 'CTSD', 'SPARC', 'SSR4', 'HLA-C', 'CLDN3', 'GRINA', 'KRT19', 'APOC1', 'SLC25A6', 'HLA-A', 'EIF4G2', 'RACK1', 'HSP90AB1', 'TAGLN', 'FLNA', 'C1QA', 'ATG10', 'HSP90AA1', 'S100A11', 'BGN', 'HSPB1', 'HLA-E', 'C3', 'CNN3', 'ATP6VOB', 'PRRC2A', 'SPDEF', 'UQCRCQ', 'MDK', 'COMP', 'SDCT', 'LMNA', 'POSTN', 'CD63', 'JTB', 'CHCHD2', 'HSPA8', 'IGFBP2', 'OAZ1', 'TSPO', 'IDH2', 'MUC1', 'LLGL', 'PGA3', 'CD99', 'PRSS8', 'MIDN', 'CIB1', 'TUBB', 'A2M', 'ALDOA', 'SCAND1', 'COPPS9', 'MZT2B', 'LUM', 'LYGE', 'LGALS3', 'CLDN4', 'ELOVL1', 'CST3', 'CFI', 'CALM', 'NACA', 'COL6A2', 'PSMB4', 'DBI', 'LAPTM4A', 'GPX4', 'PIP4K2B', 'ERGIC1', 'SUPT6H', 'SCUB5', 'DHCR24', 'TXNIP', 'MYL9', 'PEBP1', 'LAPTM5', 'UBA52', 'HSP90B1', 'SEC61A1', 'MMACHC', 'MAP3K12', 'MAPKAPK2', 'PNMT', 'CTSB', 'NBL1', 'DDIT4', 'FUCAC2', 'JUP', 'PLD3', 'BSG', 'NUPR1', 'LGALS1', 'LGALS3BP', 'UBC', 'SNRBP', 'PTMA', 'BST2', 'PCSK7', 'ZYG5', 'KDELRT', 'IGFBP7', 'COX6B1', 'UBE2M', 'COX4I1', 'PPPCA', 'COL18A1', 'COPE', 'CYBA', 'PTGES3', 'MYH9', 'TAPBP', 'CD81', 'EDF1', 'SSR2', 'CLTC', 'HEBP', 'RALY', 'PPDPF', 'ZYZ', 'HINT1', 'INTS1', 'PCBP2', 'TMBIM6', 'ZBTB7B', 'COX5B', 'ACTN4', 'ZFP641', 'FADS2', 'CTTN', 'TIMP1', 'LMAN2', 'STARD10', 'VCP', 'CHPF', 'C12ORF57', 'UBI5', 'GUK1', 'DCN', 'PFDN5', 'SLC2A4RG', 'C1QB', 'S100A10', 'ENO1', 'TMSB4X', 'KRT8', 'SREBF1', 'CLDN', 'PKM', 'SRRM2', 'ARHGDI1', 'IFI6', 'NDUFB7', 'ABCD3', 'PPP1R14B', 'ROMO1', 'FKBP2', 'PFKL', 'KDELRT', 'NDUFA3', 'MRPL12', 'LSM7', 'CIORF122', 'CPD', 'SMARCD2', 'NDUFA4', 'LDHB', 'RABAC1', 'ATP6AP1', 'ITM2B', 'SEC61G', 'SERINC2', 'RRBP1', 'EPCAM', 'PGK 1', 'HLA-DPA1', 'HMI3', 'TMED9', 'UQCRI1', 'CYB561', 'TSTD1', 'PTBPT', 'LPCAT1', 'TRIM28', 'UBB', 'AP2S1', 'SH3BGR13', 'EIF3B', 'AKT1', 'PSMD8', 'SDF4', 'CCDC152', 'MCL1', 'PSMB1', 'UQCRI0', 'LSM4', 'EIF4EBP1', 'CD46', 'PDIA6', 'SOD1', 'APP', 'TYMP', 'APIS1', 'CSDE1', 'ARPC1B', 'CD9', 'POLR1', 'CCND1'
Kidney	'IGKC', 'UMOD', 'MT-CO1', 'GPX3', 'SPPI', 'MT-CO2', 'DEFB1', 'IGHG3', 'IGHG4', 'ALDOB', 'ATP6VOC', 'IGLC2', 'IGHG1', 'IGHA1', 'RPS23', 'MT-CO3', 'RPS12', 'MT-ND3', 'RPLP1', 'RPS27', 'RPS18', 'RPL34', 'RPS15A', 'MT-ND1', 'RPL37A', 'MT-CYB', 'SLC12A1', 'MTIG', 'MT-ATP6', 'RPS27A', 'RPL37', 'RPL41', 'IGFBP7', 'APOE', 'RPS15', 'RPL12', 'RPL17', 'EEF1A', 'RPL26', 'RPS21', 'RPS24', 'ATPIA1', 'RPS28', 'TMSB4X', 'WFDC2', 'RPL32', 'RPL30', 'PCK1', 'RPS4X', 'RPS8', 'RPS14', 'RPL11', 'RPS13', 'C7', 'MT-ND2', 'RPL19', 'RPS7', 'A2M', 'RPS29', 'RPS25', 'MGP', 'RPS16', 'RPS6', 'RPS2', 'MT-ND4', 'RPL9', 'ATPIB1', 'RPL7A', 'FAU', 'RPL35A', 'IGLC3', 'PDZK1IP1', 'RPL28', 'TPT1', 'RPL10', 'PIGR', 'RPLP2', 'RPL15', 'MT-ND5', 'FTH1', 'COX7C', 'RPS3A', 'RPL23A', 'RPS19', 'HSD11B2', 'TMSB10', 'GATM', 'RPL5', 'RPL38', 'ATP5F1E', 'RPS26', 'RPL31', 'SERF2', 'UQCRCQ', 'KNG1', 'AQP2', 'SOD2', 'TAGLN', 'RPL24', 'RPL35', 'RPL13A', 'PTHIR', 'RPL21', 'PP1A', 'S100A6', 'RPS20', 'RPS9', 'RPL3', 'RPLP0', 'RPL18', 'RPS10', 'AQP1', 'RPL36A', 'NDUFA4', 'COX4I1', 'SPARC', 'HLA-DRA', 'CST3', 'VIM', 'UQCRI1', 'RPS3', 'RPL10A', 'DCN', 'CD24', 'COX5B', 'CD74', 'RPL29', 'RPL6', 'SERPINI1', 'RPS17', 'CKB', 'GAPDH', 'SLC12A3', 'COX6AF', 'IGLC1', 'FTL', 'B2M', 'RPL8', 'COX6B1', 'MALAT1', 'CTSB', 'RPL18A', 'RPL22', 'GSTP1', 'COX8A', 'RPS11', 'ATPSMG', 'UQCRCB', 'HLA-A', 'CHCHD10', 'RPL7', 'HLA-B', 'ITM2B', 'ACTG1', 'IGFBP5', 'LDHB', 'COL1A2', 'CDH16', 'RPL23', 'FXFD2', 'PODXL', 'RPL13', 'RPL14', 'CD81', 'NDUFB2', 'S100A11', 'CALB1', 'IFTM3', 'SELENOP', 'EEF1G', 'RPL27', 'RPL36', 'COX6C', 'ATP5MC3', 'ADIRF', 'MUC1', 'PTGDS', 'GNAS', 'CLCNKB', 'COX7B', 'CTSD', 'MYL9', 'MT2A', 'MAL', 'RACK1', 'ACTB', 'MT-ND4L', 'CA12', 'NAT8', 'CRYAB', 'ADGRG1', 'COX7A2', 'PEBP1', 'EIF1', 'PSAP', 'UBA52', 'REN', 'ATP5F1B', 'NDUFA1', 'ATP5ME', 'RPSA', 'TIMP3', 'NEAT1', 'ATP5F1A', 'UBI5', 'MTIE', 'NACA', 'MMP7', 'COL1A1', 'CD63', 'TIMP1', 'HINT1', 'ENO1', 'ELOB', 'UQCRIH', 'UBB', 'RPL27A', 'RPS5', 'MIF', 'MYL6', 'AEBP1', 'TOMM7', 'ATP5P', 'ACTA2', 'SLC25A6', 'CLU', 'SLC25A5', 'SRP14', 'GPX4', 'S100A10', 'MIOX', 'GSSG', 'NDUFB1', 'CTSH', 'PFDN5', 'SLC1A1', 'S100A2', 'PKM', 'TSPAN1', 'LUM', 'RNASE1', 'ASB1', 'APP', 'BGN', 'HSPB1', 'RPL4', 'COL3A1', 'TMA7', 'SLC5A12', 'SERPINAS', 'TFTM2', 'SERPING1', 'SPINK1', 'UQCRI0', 'PFN1', 'CHCHD2', 'UBC', 'CXCL14', 'TMBIM6', 'SAT1', 'LRP2', 'APLP2', 'CFI1', 'SLC13A3', 'EEF2', 'CD59', 'HLA-DRA', 'CALR', 'PPP1R1A', 'NME2', 'TPT1', 'ATP6VOB', 'GSTM3', 'COXSA', 'ANXA2', 'TXNIP', 'PRDX5', 'NDUFB7', 'ATP5F1D', 'OST4', 'PRDX1', 'ATP5MC2', 'SLC25A3', 'CYBSA', 'MYL12B', 'OAZ1', 'NDUFB9', 'ATP5PO', 'CYSTMI', 'TUBA1B', 'ANPEP', 'HSP90AA1', 'FABP1', 'PPDPF', 'IGFBP4', 'DSTN', 'EPAS1', 'RNASEK', 'PTMA'
Breast Cancer	'TFE3', 'IGLL5', 'ERBB2', 'MUC1', 'MGP', 'RPS3', 'S100A9', 'KRT19', 'HAJ', 'NHERF1', 'COX6C', 'AZGP1', 'TMSB10', 'TF127', 'RPL18A', 'B2M', 'CAB', 'FASN', 'RPL19', 'APOE', 'HLA-B', 'HSPB1', 'ACTB', 'PABPC1', 'RPS15', 'RPS12', 'FNL', 'IGFBP2', 'RPS8', 'HLA-C', 'FTH1', 'CCND1', 'S100A8', 'CTSD', 'RPL31', 'RPL37A', 'EEF2', 'KRT18', 'XBP1', 'TAGLN', 'RPS17', 'RPS21', 'COL1A1', 'BEST1', 'SNHG25', 'ACTG1', 'RPL18', 'RPS18', 'RPS2', 'CST3', 'HLA-A', 'PSMD3', 'UBA52', 'SSR4', 'RPL29', 'RPL13', 'A2M', 'RPS6', 'AEBP1', 'GAPDH', 'RPS4X', 'RPL30', 'TTL', 'RPS11', 'RPLP1', 'RPL8', 'RPS24', 'RPS28', 'RPL28', 'RPS27A', 'RPL10', 'RPS19', 'BST2', 'RPL37', 'RPS27', 'S100A11', 'CD74', 'CLDN3', 'EVL', 'KRT8', 'RPLA', 'RPS14', 'RPL35', 'PPPCA', 'F4H8', 'SYNGR2', 'HSP90AB1', 'HLA-DRA', 'SLC39A6', 'RPLP0', 'RPL27A', 'RPL23', 'RPS29', 'SERF2', 'RPL12', 'LYGE', 'RPS8', 'RACK1', 'GNAS', 'LGALS3BP', 'S100A14', 'MZT2B', 'PFN1', 'IFTM3', 'RPL11', 'HLA-E', 'RPL35A', 'H3-3B', 'RPL9', 'RPL31A', 'MIEN1', 'PRDX2', 'GPX4', 'RPL15', 'APOC1', 'BGN', 'CLU', 'PSAP', 'ATP6VOB', 'C1QA', 'ISG15', 'TPT1', 'RPL36', 'IFI6', 'RPL23A', 'RPL27', 'RPS9', 'FXFD3', 'C3', 'GATA3', 'KRT7', 'BSG', 'MZT2A', 'PTMA', 'PLXNB2', 'EIF4A1', 'STARD10', 'IGFBP', 'RPL3', 'TAPBP', 'RPS16', 'SPDEF', 'NDUFA13', 'ELF3', 'RPS3A', 'ALDOA', 'TAGLN2', 'CD81', 'TMSB4X', 'ATP6AP1', 'COL1A2', 'BCAP31', 'MALAT1', 'DDX', 'TUBA1B', 'CFI1', 'MYL9', 'ATPIA1', 'PTPRF', 'RPS23', 'RPL', 'ANPEP', 'CALR', 'RPL32', 'NBEAL1', 'FAU', 'GNB2', 'RPS8', 'RPLP2', 'S100A6', 'RPS13', 'SDCL', 'RPS7', 'UQCRCQ', 'SLC25A6', 'PRDX1', 'VIM', 'RPL14', 'GASS', 'RPL4', 'TUFM', 'POLR1', 'NR2F6', 'RPL34', 'TMEM123', 'GRINA', 'UBC', 'H1-10', 'TER3', 'OAZ1', 'METRN', 'IGFBP4', 'RPS15A', 'ENOL', 'CISD3', 'PSMB3', 'POSTN', 'RPL7', 'TIMP1', 'FLNA', 'CD24', 'RPS25', 'CRIP2', 'LAPTM5', 'NDUFB9', 'ARHGDI1', 'JUP', 'GSTP1', 'CLDN4', 'COLA1', 'CYBA', 'EDF1', 'RPL24', 'EIF3C', 'NDUFA11', 'COX8A', 'NPDC1', 'SPARC', 'RPS20', 'CD63', 'MFE8', 'GLUL', 'PPDPF', 'SREBF1', 'MAGED2', 'DEGS2', 'C12orf57', 'RPSA', 'SPINT2', 'RPL41', 'RHOC', 'SLC44A2', 'SERINC', 'PTMS', 'COX6B1', 'CTSB', 'UBB', 'WDR83OS', 'HNRNP2B1', 'ATP5F1E', 'C1QB', 'GRN', 'COL6A2', 'SEC61A1', 'ATG10', 'EIFSA', 'CRABP2', 'ELOB', 'TT', 'COX4I1', 'SELENOW', 'TXNIP', 'UQCRI', 'RPL6', 'MMP14', 'MDK', 'GUK1', 'TUBB', 'H1-2', 'LMNA', 'RPL22', 'MYL6', 'CHCHD2', 'CTTN', 'RPL10A', 'EEF1A', 'LSM4', 'DHCR24', 'CAPNS1', 'LAPTM4A', 'SSR2', 'UBE2M', 'TSPO', 'LMAN2', 'RPS10', 'PFDN5', 'NENF', 'TSKU', 'SERPINH1', 'IGLL1', 'LGALS1', 'CHPF', 'EEF1D', 'MIDN', 'ATP5MC2', 'CLDN7', 'STARD3', 'LGLN2', 'CCT3', 'EIF3B', 'ZNF90', 'SH3BGR13', 'MZB1', 'CDC37', 'PTBP1', 'AP2M1', 'IGFBP7', 'ENSA', 'EIF4G1'

Figure 4: Selected high variance genes for gene expression prediction task across datasets.

Appendix B. Additional Implementation Details

To evaluate the effectiveness of our sampling strategy, we compare it against three representative baselines: Monte Carlo random sampling, uncertainty-based sampling (Safaei et al., 2024), and diversity-driven sampling (Zhdanov, 2019).

Uncertainty-based sampling. This method estimates prediction uncertainty via Monte Carlo Dropout. Specifically, we insert a Dropout layer (drop rate = 0.1) after the feature extraction block of the vision encoder, and the model outputs the expression values of 300 genes. During uncertainty estimation, the Dropout layer remains activated, and we perform $T = 20$ stochastic forward passes for each patch. We compute the variance across these predictions and use its mean as the entropy score. Higher entropy indicates greater model uncertainty, and patches with high entropy are prioritized during sampling.




Figure 5: **Comparison of different sampling ratios (10-75%) across datasets.** Our SCR²Net consistently achieves lower MSE and higher PCC, especially under low sampling budgets.

Diversity-driven sampling. This method encourages sample diversity based on feature similarity. We first extract 1024-dimensional visual features using the vision encoder during training. These features are standardized and reduced to 128 dimensions via PCA, followed by clustering using DBSCAN. To avoid insufficient cluster granularity, we incorporate a dynamic adjustment mechanism: the minimum cluster count is adaptively set to $\sqrt{N}/5$, where N denotes the total number of samples. If DBSCAN yields too few clusters, we automatically switch to KMeans to enforce the desired number of clusters. During sampling, patches are drawn uniformly from each cluster to ensure diverse coverage of tissue regions.

Our SCRL sampling strategy. For the multi-objective reward function, we set $w_{sc} = 20$, $w_{type} = 5$, and $w_{spa} = 0.05$ to balance manifold coverage, cell-type diversity, and spatial density constraints. The loss weights λ_r , λ_p , and λ_{KD} are set to 1.0, 0.25, and 0.25, respectively. We adopt a similarity confidence mask with threshold $m = 0.15$. Active sampling proceeds for 20 rounds, with the sampled training set updated every 5 epochs. The initial round employs random sampling for warm-up. A fixed seed of 42 is used for reproducibility. During retrieval, we select the top-50 most similar expression profiles and retain only the top-10 dominant cell types for robust reference aggregation.




Figure 6: **Comparison of sampling strategies across multiple methods on different datasets.** We evaluate performance under training data ratios of 10%–75% to simulate low-resource scenarios, with each curve denoting a sampling strategy. Our proposed SCRL sampling consistently achieves better performance across datasets and models, particularly in low-budget (10%–25%) regimes.

Appendix C. Additional Experimental Results

Due to space limitations in the main manuscript, we provide additional experimental results in this appendix, including the comparison of different sampling ratios across datasets (Figure 5), the comparison of sampling strategies evaluated on mlSTExp and TRIPLEX benchmarks (Figure 6), and an ablation study on the reward function design (Figure 7).




Figure 7: **Ablation study of reward function in our proposed SCRL sampling.**