

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 HIST: SPATIAL TRANSCRIPTOMICS PREDICTION VIA MULTI-LEVEL HYPERBOLIC REPRESENTATION LEARN- ING

Anonymous authors

Paper under double-blind review

ABSTRACT

Spatial Transcriptomics (ST) merges the benefits of pathology images and gene expression, linking molecular profiles with tissue structure to analyze spot-level function comprehensively. Predicting gene expression from histology images is a cost-effective alternative to expensive ST technologies. However, existing methods mainly focus on spot-level image-to-gene matching but fail to leverage the full hierarchical structure of ST data, especially on the gene expression side, leading to incomplete image-gene alignment. Moreover, a challenge arises from the inherent information asymmetry: gene expression profiles contain more molecular details that may lack salient visual correlates in histological images, demanding a sophisticated representation learning approach to bridge this modality gap. We propose **HiST**, a framework for ST prediction that learns multi-level image-gene representations by modeling the data's inherent hierarchy within hyperbolic space, a natural geometric setting for such structures. First, we design a **Multi-Level Representation Extractor** to capture both spot-level and niche-level representations from each modality, providing context-aware information beyond individual spot-level image-gene pairs. Second, a **Hierarchical Hyperbolic Alignment** module is introduced to unify these representations, performing spatial alignment while hierarchically structuring image and gene embeddings. This alignment strategy enriches the image representations with molecular semantics, significantly improving cross-modal prediction. HiST achieves state-of-the-art performance on three public datasets from different tissues, paving the way for more scalable and accurate spatial transcriptomics prediction.

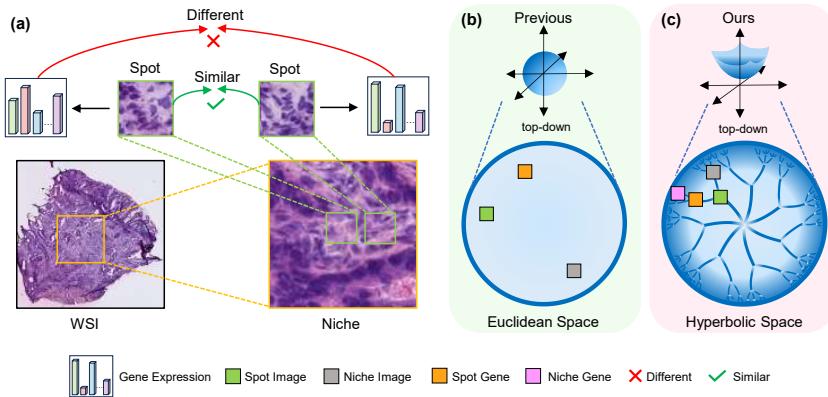


Figure 1: **ST data characteristics.** (a) A WSI contains hierarchical structures and visually similar patterns may correspond to different gene expressions. (b) Other works mainly model ST data in Euclidean Space, which neglects niche-level gene and can lead to biased biological insights. (c) Our hyperbolic approach models hierarchies based on information specificity, where a general concept (image/spot) entails its more specific, information-rich counterpart (gene/niche), enabling more informative representation learning.

054
055
056
057

1 INTRODUCTION

058
059
060
061
062
063
064
065
066
067
068
Pathological images, particularly Hematoxylin and Eosin (H&E) stained Whole Slide Images (WSIs),
provide critical insights into cell morphology and tissue architecture, serving as a cornerstone in
biomedical research and clinical diagnosis (Lu et al., 2024; Chen et al., 2024a; Li et al., 2024). Gene
expression data complement these pathological images by elucidating the molecular mechanisms
underlying observed features, thereby enhancing disease diagnosis and facilitating therapeutic target
identification (Ash et al., 2021). Spatial Transcriptomics (ST) integrates both modalities by capturing
spatially resolved gene expression and cellular morphology simultaneously (Ståhl et al., 2016),
aligning molecular profiles with tissue structure at micrometer resolution (Williams et al., 2022).
Despite its advantages, ST has not achieved widespread clinical adoption due to its high cost and
laborious experiment compared to traditional techniques (Zhang et al., 2022; Choe et al., 2023).
Consequently, there has been an increasing attention in predicting spatially resolved gene expression
directly from pathological images using deep learning approaches (Wang et al., 2025a).069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
Recent studies have explored diverse strategies for this prediction task, including direct inference
from spot-level images (He et al., 2020; Monjo et al., 2022), integration of multi-scale features across
WSIs (Chung et al., 2024; Wang et al., 2025b), and contrastive learning to align spot-level images with
gene expression profiles (Xie et al., 2023). Although these methods have shown promising results,
several critical questions remain underexplored. Our study is motivated by two key questions in ST
prediction: (1) **Can integrating broader pathological and genetic context improve spot-level gene
expression inference?** Previous studies often primarily utilized multi-scale pathological features
for gene expression prediction, neglecting the multi-level structure inherent in gene expression
itself (Jaume et al., 2024b; Chen et al., 2024b), which spans cellular and tissue-level scales. In reality,
both broader pathological context and bulk genetic programs can significantly influence the gene
expression profile at each spot (Chen et al., 2020; Nirmal et al., 2022; Wu et al., 2025; Ye et al., 2024).
(2) **How can we effectively learn more target-modality information (i.e., gene expression) during
training to enhance cross-modal prediction?** Biological heterogeneity frequently results in visually
similar pathology patches exhibiting distinct gene expression patterns (Zhu et al., 2025; Pizurica et al.,
2024; Fujii et al., 2024; Tang et al., 2025), as illustrated in Figure 1 (a). This phenomenon indicates
that standard image encoders may fail to capture the subtle morphological cues for predicting these
molecular variations. Instead of viewing this as an ill-posed one-to-many problem, we contend that
the key is to learn a more powerful and molecularly-informed image representation.086
087
088
089
090
091
092
093
094
095
To address these two questions, we introduce **HiST**, a novel framework for ST prediction by learning
multi-level hyperbolic image-gene representations. HiST tackles these challenges with two core
components. First, our **Multi-Level Representation Extractors** capture hierarchical representa-
tions from both pathology images and their corresponding gene expression profiles. They extract
multimodal information at both spot- level and niche-level, where a niche consists of a central spot
and its surrounding neighbors, enabling the capture of comprehensive morphological and molecular
patterns across spatial scales. Second, our **Hierarchical Hyperbolic Alignment** module acts as
a powerful structural regularizer rather than a generative model. It uses the unique properties of
hyperbolic geometry to impose a meaningful inductive bias on the latent space, guiding the model to
learn molecularly-informed features.096
097
098
099
100
101
102
103
104
105
We define our hierarchical relationships based on **information specificity**. In this view, a concept A
entails a concept B if B is a semantically richer and more specific instance of A. For example, the
concept of a “dog on a beach” is more specific and information-rich than “dog”, and is thus considered
the child concept. Following this principle, we establish two key hierarchies in our framework: (1) A
spot-level representation entails its context-rich niche-level counterpart. (2) A morphological image
entails its corresponding gene expression. This is because the gene profile contains fine-grained
molecular information that offers a much more specific description of the tissue’s state than the
more general pathology image. HiST learns powerful, context-aware representations by modeling
information-based hierarchies in hyperbolic space, which is inherently more suited for capturing such
structures than Euclidean space (Hsu et al., 2021).106
107
We demonstrate HiST’s effectiveness on three public datasets from diverse tissues, where it consis-
tently outperforms state-of-the-art models. Our contributions are summarized as:

- 108 • We propose **HiST**, a novel framework for predicting spatially resolved gene expression
109 from WSIs by learning multi-level hyperbolic representations that capture the intrinsic
110 hierarchical structure of ST data.
- 111 • We design **Multi-Level Representation Extractors** to capture spot- and niche-level repre-
112 sentations from both modalities, providing comprehensive biological insights.
- 113 • We introduce **Hierarchical Hyperbolic Alignment** to structurally regularize the latent
114 space, improving cross-modal feature integration.
- 115 • Extensive experiments on three public datasets demonstrate that HiST consistently out-
116 performs existing approaches, underscoring its robust efficacy in spatial gene expression
117 prediction.

119 2 RELATED WORK

120 2.1 PREDICTION OF GENE EXPRESSION FROM HISTOLOGY IMAGES

121 Recent methodologies for predicting spatially resolved gene expression from histology images have
122 advanced through diverse computational paradigms, including ST-Net (He et al., 2020), BLEEP (Xie
123 et al., 2023), TRIPLEX (Chung et al., 2024), and Stem (Zhu et al., 2025). Local image-to-expression
124 regression models like ST-Net employ ResNet50 (He et al., 2016) to directly map H&E image patches
125 to gene expressions. While effective in deterministic prediction, these methods assume injective
126 mappings between morphology and transcription, overlooking biological heterogeneity. Multi-scale
127 integration approaches like TRIPLEX extract and fuse multi-resolution features from WSIs using
128 attention mechanisms. Although these methods capture multi-resolution visual patterns, they lack
129 explicit constraints to preserve the essential biological hierarchy. Generative models like Stem address
130 the uncertainty in expression prediction by generating probabilistic gene expression profiles. While
131 these paradigms better preserve transcriptional variability, they neglect the inherent data hierarchy. In
132 contrast to these prior works, HiST explicitly models the intrinsic parent-child relationships between
133 spots and their surrounding niches across both imaging and gene expression modalities.

134 2.2 MULTIMODAL CONTRASTIVE REPRESENTATION LEARNING

135 Contrastive learning is a pivotal technique for cross-modal tasks by aligning representations across
136 different modalities. For example, CLIP (Radford et al., 2021) employs contrastive learning to align
137 paired images and texts in a shared Euclidean embedding space. Inspired by CLIP, BLEEP (Xie
138 et al., 2023) adapts contrastive learning to histology and gene expression, using direct interpolation
139 in the embedding space for efficient, decoder-free predictions. These two models rely on Euclidean
140 embeddings, which limit their ability to capture hierarchical relationships. To overcome these
141 limitations, MERU (Desai et al., 2023) embeds image and text into hyperbolic space, leveraging its
142 geometric properties to build a hierarchical representation space through contrastive and entailment
143 losses. Building on MERU, HyCoCLIP (Pal et al., 2024) introduces intra-modal hierarchical modeling
144 by extracting object boxes from images and their corresponding textual descriptions, establishing
145 hierarchical links between box regions and the full image-text pair. HyCoCLIP’s dependence on
146 pre-trained object detection models to derive these boxes from given captions may result in potential
147 inaccuracies. In contrast, HiST directly leverages the inherent structure of ST data, from spot-level to
148 niche-level contexts, avoiding uncertainties associated with external feature extraction.

149 3 METHOD

150 The overview of HiST is illustrated in Figure 2. First, we briefly describe the preliminaries of the
151 hyperbolic geometry in Section 3.1. Second, we present the Multi-Level Representation Extractors in
152 Section 3.2. Third, we introduce the Hierarchical Hyperbolic Alignment in Section 3.3. Finally, we
153 describe the Gene Decoder and our Overall Objective Function in Section 3.4.

154 3.1 PRELIMINARIES

155 **Hyperbolic Geometry** Hyperbolic Geometry is a fundamental class of non-Euclidean geometry
156 with a constant negative curvature. This distinguishing characteristic results in an exponential growth

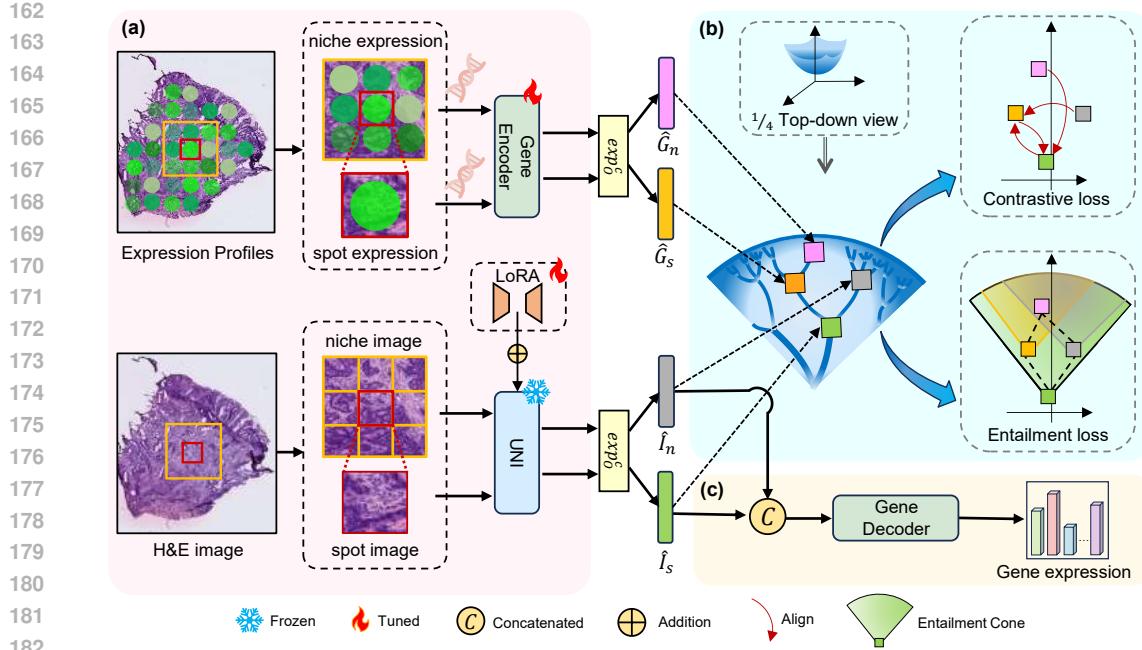


Figure 2: **Overview of HiST.** HiST consists of three components. (a) **Multi-Level Representation Extractors** capture spot- and niche-level features from both images and gene expression. (b) **Hierarchical Hyperbolic Alignment** module projects these features into a shared hyperbolic latent space. It uses contrastive alignment for corresponding image-gene pairs and entailment alignment to structurally regularize the latent space according to information hierarchies. (c) **Gene Decoder** uses the resulting aligned and context-aware image representations to predict spot-level gene expression.

of volume with respect to radius, in stark contrast to Euclidean geometry, which exhibits zero curvature and polynomial volume scaling (Mettes et al., 2024). Consequently, hyperbolic spaces are naturally adept at representing tree-like or hierarchical data structures, where the number of elements increases exponentially with depth (Hsu et al., 2021; Pal et al., 2024). Due to their negative curvature, hyperbolic spaces cannot be isometrically embedded in Euclidean spaces of equivalent dimensionality without compromising distances or angles. To address this issue, several geometric models are employed for their representation and computation, including the Poincaré ball model and the Lorentz model (Cannon et al., 1997; Cho et al., 2022).

Lorentz Model Lorentz model is widely preferred due to its numerical stability and straightforward geodesic calculations (Nickel & Kiela, 2018). The Lorentz model \mathbb{L}_c^n embeds the n -dimensional hyperbolic space as the upper sheet of a two-sheeted hyperboloid in $(n+1)$ -dimensional Minkowski space, with a constant curvature $-c < 0$, it consists of all vectors satisfying:

$$\mathbb{L}_c^n = \{\mathbf{x} \in \mathbb{R}^{n+1} : \langle \mathbf{x}, \mathbf{x} \rangle_{\mathbb{L}} = -\frac{1}{c}, x_{time} = \sqrt{1/c + \|\mathbf{x}_{space}\|^2}, c > 0\}, \quad (1)$$

where points $\mathbf{x} \in \mathbb{R}^{n+1}$ in \mathbb{L}_c^n can be represented as $[x_{time}, \mathbf{x}_{space}]$. $x_{time} \in \mathbb{R}$ and $\mathbf{x}_{space} \in \mathbb{R}^n$ denote the *time component* and the *spatial component* (Desai et al., 2023), respectively. For two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{L}_c^n$, the Lorentzian inner product $\langle \cdot, \cdot \rangle_{\mathbb{L}}$ is defined as $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{L}} = \langle \mathbf{x}_{space}, \mathbf{y}_{space} \rangle_{\mathbb{E}} - x_{time}y_{time}$, where $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{E}}$ represents the Euclidean inner product in \mathbb{R}^n . Besides, the Lorentzian distance $d_{\mathbb{L}}(\mathbf{x}, \mathbf{y})$ measures the length of the shortest path between two points \mathbf{x} and \mathbf{y} , which is formulated as:

$$d_{\mathbb{L}}(\mathbf{x}, \mathbf{y}) = \sqrt{1/c} \cdot \cosh^{-1}(-c\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{L}}). \quad (2)$$

Tangent Space and Exponential Map The tangent space of $\mathbf{x} \in \mathbb{L}_c^n$ is denoted by $\mathcal{T}_{\mathbf{x}}\mathbb{L}_c^n$ which is precisely defined as the set of vectors orthogonal to \mathbf{x} under the Lorentzian inner product:

$$\mathcal{T}_{\mathbf{x}}\mathbb{L}_c^n = \{\mathbf{v} \in \mathbb{R}^{n+1} : \langle \mathbf{x}, \mathbf{v} \rangle_{\mathbb{L}} = 0\}. \quad (3)$$

216 A fundamental mechanism for connecting the tangent space to the hyperbolic manifold is the
 217 exponential map. The exponential map $\exp_x^c : T_x \mathbb{L}_c^n \rightarrow \mathbb{L}_c^n$ projects tangent vector \mathbf{v} onto the \mathbb{L}_c^n
 218 along a geodesic emanating from x in the direction of \mathbf{v} , given by:
 219

$$\exp_x^c(\mathbf{v}) = \cosh(\sqrt{c}\|\mathbf{v}\|_{\mathbb{L}})\mathbf{x} + \frac{\sinh(\sqrt{c}\|\mathbf{v}\|_{\mathbb{L}})}{\sqrt{c}\|\mathbf{v}\|_{\mathbb{L}}}\mathbf{v}, \quad (4)$$

220 where $\|\mathbf{v}\|_{\mathbb{L}} = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle_{\mathbb{L}}}$ is the Lorentzian norm. Moreover, the exponential map serves as a bridge
 221 between Euclidean and hyperbolic geometries. By interpreting Euclidean vectors as tangent vectors
 222 at the origin $\mathbf{O} = [\sqrt{1/c}, 0, \dots, 0] \in \mathbb{R}^{n+1}$ of the hyperbolic space (Mettes et al., 2024; Pal et al.,
 223 2024; Khrulkov et al., 2020), we begin by extending the Euclidean embedding $\mathbf{v}_{euc} \in \mathbb{R}^n$ into \mathbb{R}^{n+1}
 224 by defining a vector $\mathbf{v} = [0, \mathbf{v}_{euc}] \in \mathbb{R}^{n+1}$. This vector \mathbf{v} is situated in the tangent space at the origin
 225 \mathbf{O} of the hyperboloid as $\langle \mathbf{O}, \mathbf{v} \rangle_{\mathbb{L}} = 0$. Thus, \mathbf{v} can be projected onto the hyperboloid \mathbb{L}_c^n employing
 226 the exponential map:
 227

$$\mathbf{x}_{space} = \exp_{\mathbf{O}}^c(\mathbf{v}_{euc}) = \frac{\sinh(\sqrt{c}\|\mathbf{v}_{euc}\|_{\mathbb{E}})}{\sqrt{c}\|\mathbf{v}_{euc}\|_{\mathbb{E}}}\mathbf{v}_{euc}. \quad (5)$$

228 Then we can directly calculate the corresponding time component x_{time} from \mathbf{x}_{space} . The detailed
 229 derivations of the above equations can be found in Appendix A.
 230

231 **Hyperbolic Entailment Loss** The entailment cone \mathcal{R}_y constitutes a region of point y where all
 232 points $x \in \mathcal{R}_y$ represent child concepts of the parent concept y (Ganea et al., 2018; Desai et al.,
 233 2023), defined by the half-aperture:
 234

$$\text{aper}(\mathbf{y}) = \sin^{-1} \left(\frac{2K}{\sqrt{c}\|\mathbf{y}_{space}\|} \right), \quad (6)$$

235 where $K = 0.1$ determines boundary conditions near the origin. To enforce the partial order
 236 relationship where y entails x , the penalty is formulated as:
 237

$$\mathcal{L}_{entail}(\mathbf{y}, \mathbf{x}) = \max(0, \text{ext}(\mathbf{y}, \mathbf{x}) - \text{aper}(\mathbf{y})), \quad (7)$$

238 where $\text{ext}(\mathbf{y}, \mathbf{x})$ denotes the exterior angle defined as $\text{ext}(\mathbf{y}, \mathbf{x}) = \cos^{-1} \left(\frac{x_{time} + y_{time} c \langle \mathbf{y}, \mathbf{x} \rangle_{\mathbb{L}}}{\|\mathbf{y}_{space}\| \sqrt{(c \langle \mathbf{y}, \mathbf{x} \rangle_{\mathbb{L}})^2 - 1}} \right)$.
 239

240 3.2 MULTI-LEVEL REPRESENTATION EXTRACTORS

241 **Multi-Level Pathological Images Extractor** Following the previous works (Xie et al., 2023;
 242 Zhu et al., 2025), a spot-level image patch $X_s \in \mathbb{R}^{3 \times L_s \times L_s}$ of each identified spot is extracted
 243 and preprocessed from a H&E stained image, with the spot positioned at the center of the patch,
 244 as depicted in Figure 2 (a), where L_s represents the size of the spot image. While X_s directly
 245 corresponds to the target spot gene expression, the additional nearby visual information from high-
 246 level pathology patches can significantly contribute to the analysis (Chung et al., 2024; Lin et al.,
 247 2024). Therefore, we introduce the niche-level image patch $X_n \in \mathbb{R}^{3 \times L_n \times L_n}$, which is defined as a
 248 higher-level region composed of the central spot-level patch X_s and its spatially adjacent spot-level
 249 patches. L_n signifies the patch size of the niche image. These neighboring patches are selected based
 250 on spatial proximity using the K-Nearest Neighbors (KNN) algorithm. By cropping the region of
 251 these patches, X_n forms a larger image region that provides a broader field of view and enhanced
 252 contextual information about the surrounding tissue microenvironment.
 253

254 We leverage UNI (Chen et al., 2024a), a pathology foundation model pre-trained on large-scale
 255 histology images, to extract feature embeddings for spot-level and niche-level image patches. As the
 256 original UNI was not well-suited for large-sized niche-level image patches, we resized the images
 257 in our dataset and fine-tuned UNI using the Low-Rank Adaptation (LoRA) technique (Hu et al.,
 258 2022), leading to improved multi-level visual representations. Consider a frozen pre-trained weight
 259 matrix $W_{origin} \in \mathbb{R}^{d \times d}$, where d denotes the dimension. The updated weight matrix is formulated as
 260 $W_{new} = W_{origin} + \Delta W = W_{origin} + BA$, where the update $\Delta W \in \mathbb{R}^{d \times d}$ expressed by the product
 261 of two smaller trainable matrices: $B \in \mathbb{R}^{d \times r}$, $A \in \mathbb{R}^{r \times d}$, and the rank $r \ll d$. This approach enables
 262 us to adapt the model to the characteristics of our data while substantially reducing the computational
 263 resources required for fine-tuning. The multi-level image representations $I_s \in \mathbb{R}^d$ and $I_n \in \mathbb{R}^d$ are
 264 extracted by $I_s, I_n = MIE(X_s, X_n; \theta_{UNI}, \Delta\theta_{lora})$. Here, MIE represent the UNI model adapted
 265 by LoRA, θ_{uni} and $\Delta\theta_{lora}$ denote the frozen parameters of UNI and the trainable parameters of
 266 LoRA modules.
 267

270 **Multi-Level Genomic Profiles Extractor** Let $Y_s \in \mathbb{R}^N$ be the associated spot-level gene expression profile of the spot-level image X_s , where N is the gene set size. In the same vein as the niche-level image patch, we introduce the niche-level gene expression profile $Y_n \in \mathbb{R}^N = \frac{1}{|S|} \sum_{i \in S} Y_s^i$, where
 271 $S = \{Y_s^1, \dots, Y_s^K\}$ denotes the expression profile set of Y_s and its neighbors, $K - 1$ is the number of
 272 selected neighbors. $G_s, G_n = MGE(Y_s, Y_n; \theta_{gene})$, where MGE denotes the multi-level genomic
 273 profiles extractor with trainable parameters implemented by a trainable fully connected network
 274 θ_{gene} , $G_s \in \mathbb{R}^d$ and $G_n \in \mathbb{R}^d$ are the spot-level and niche-level gene embeddings, respectively.
 275
 276

278 3.3 HIERARCHICAL HYPERBOLIC ALIGNMENT

280 To obtain better representations for facilitating the subsequent tasks, the alignment is a pivotal method
 281 which bridges the gap of different modalities (Xie et al., 2023; Zhang et al., 2025; Li et al., 2021).
 282 However, common implementations of alignment, such as BLEEP (Xie et al., 2023), directly close
 283 the different items in Euclidean space, which may not be appropriate for hierarchical data like ST. To
 284 address this problem, we design a **Hierarchical Contrastive Alignment** module, which aligns the
 285 different modalities at different levels in hyperbolic space. Subsequently, we introduce a **Hierarchical**
 286 **Entailment Alignment** module to regularize the partial order in ST data.
 287

288 **Hierarchical Contrastive Alignment (HCA)** Using Equation 5, let $\exp_O^c(\cdot) : \mathbb{R}^d \rightarrow \mathbb{L}_c^d$ map
 289 Euclidean features to hyperbolic space with trainable curvature $-c < 0$ and origin O . This yields
 290 hyperbolic spatial components $\{\hat{I}_s^{space}, \hat{I}_n^{space}, \hat{G}_s^{space}, \hat{G}_n^{space}\} = \exp_O^c(\{I_s, I_n, G_s, G_n\})$, while
 291 the corresponding time components can be calculated by Equation 1. The hyperbolic representations
 292 $\hat{I}_s, \hat{I}_n, \hat{G}_s$ and \hat{G}_n are obtained by concatenating spatial components and time components. To
 293 align the spot-level image embedding to the spot-level embedding, we employ modified infoNCE
 294 loss (Oord et al., 2018), in which the cosine similarity is replaced by the Lorentzian distance $d_{\mathbb{L}}(\cdot, \cdot)$
 295 described in Equation 2. The contrastive loss is defined as follows:
 296

$$296 \mathcal{L}_{align}(\hat{I}_s, \hat{G}_s) = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(d_{\mathbb{L}}(\hat{I}_s^i, \hat{G}_s^i)/\tau)}{\sum_{j=1, j \neq i}^B \exp(d_{\mathbb{L}}(\hat{I}_s^i, \hat{G}_s^j)/\tau)}, \quad (8)$$

297 where B denotes the batch size and τ is the temperature parameter. To better utilize in-batch negatives,
 298 we also align spot-level gene to spot-level image embeddings using $\mathcal{L}_{align}(\hat{G}_s, \hat{I}_s)$. Since spot-level
 299 features represent more general characteristics, a single spot-level feature may correspond to multiple
 300 niche-level features within a batch. To avoid such undesirable negative alignment, we only consider
 301 the alignment from niche-level features to spot-level features, i.e., $\mathcal{L}_{align}(\hat{G}_n, \hat{I}_s)$ and $\mathcal{L}_{align}(\hat{I}_n, \hat{G}_s)$.
 302 The objective function of Hierarchical Contrastive Alignment can be expressed as:
 303

$$305 \mathcal{L}_{HCA} = \frac{1}{4} (\mathcal{L}_{align}(\hat{I}_s, \hat{G}_s) + \mathcal{L}_{align}(\hat{G}_s, \hat{I}_s) + \mathcal{L}_{align}(\hat{G}_n, \hat{I}_s) + \mathcal{L}_{align}(\hat{I}_n, \hat{G}_s)) \quad (9)$$

308 **Hierarchical Entailment Alignment (HEA)** Beyond spot-niche hierarchies, we account for the
 309 non-identical nature of image features and gene features. We recognize that gene features provide
 310 finer-grained molecular insights. Thus, we posit that gene features are the child concept of images in
 311 hyperbolic space. In our ST data, this hierarchy can be summarized as spot-level features entailing
 312 niche-level features, and pathological images entailing their corresponding gene expression profiles.
 313 In order to directly constrain this hierarchical structure, we leverage Hyperbolic Entailment Loss
 314 $\mathcal{L}_{entail}(\cdot, \cdot)$ described in Equation 7. Therefore, the final objective function of this module is
 315 formulated as:
 316

$$316 \mathcal{L}_{HEA} = \frac{1}{4} (\mathcal{L}_{entail}(\hat{I}_s, \hat{I}_n) + \mathcal{L}_{entail}(\hat{G}_s, \hat{G}_n) + \mathcal{L}_{entail}(\hat{I}_s, \hat{G}_s) + \mathcal{L}_{entail}(\hat{I}_n, \hat{G}_n)). \quad (10)$$

319 3.4 GENE DECODER BASED ON ALIGNED REPRESENTATIONS AND OBJECTIVE FUNCTION

321 To predict the spot-level gene expression profiles, we directly concatenate the aligned representations
 322 (I_s and I_n) and feed the result into a gene decoder implemented by Multi-Layer Perceptron
 323 (MLP) (LeCun et al., 2015), which can be expressed by $Y^{pred} = Decoder_{gene}(\text{concat}(I_s, I_n))$.
 324 MSE loss is leveraged to optimize this decoder: $\mathcal{L}_{pred} = \|Y^{pred} - Y_s\|_2^2$.
 325

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

Table 1: Performance comparison on three spatial transcriptomics datasets. Higher values on PCC@10, PCC@50, PCC@200 are better. Lower values on MAE, MSE are better.

Dataset	Model	PCC@10 ↑	PCC@50 ↑	PCC@200 ↑	MSE ↓	MAE ↓
Colorectum	TRIPLEX	0.701±0.128	0.624±0.154	0.462±0.191	1.869±0.803	1.056±0.239
	StNet	0.646±0.134	0.570±0.142	0.419±0.176	1.686±0.373	1.023±0.134
	BLEEP	0.637±0.112	0.556±0.120	0.382±0.160	2.038±0.587	1.096±0.164
	Stem	0.670±0.116	0.573±0.130	0.399±0.166	1.788±0.418	1.032±0.138
	HiST(Ours)	0.721±0.105	0.642±0.128	0.477±0.184	1.498±0.456	0.958±0.158
Skin	TRIPLEX	0.831±0.094	0.799±0.114	0.740±0.142	0.981±0.466	0.685±0.205
	StNet	0.804±0.105	0.779±0.117	0.726±0.140	0.993±0.469	0.689±0.198
	BLEEP	0.788±0.111	0.761±0.123	0.704±0.145	1.117±0.540	0.701±0.221
	Stem	0.782±0.094	0.748±0.113	0.687±0.138	1.276±0.703	0.730±0.261
	HiST(Ours)	0.839±0.086	0.812±0.102	0.758±0.129	0.932±0.418	0.657±0.182
Kidney	TRIPLEX	0.579±0.095	0.485±0.084	0.351±0.066	1.122±0.204	0.855±0.104
	StNet	0.523±0.105	0.435±0.095	0.305±0.064	1.167±0.217	0.847±0.078
	BLEEP	0.518±0.112	0.434±0.102	0.310±0.071	1.233±0.244	0.865±0.085
	Stem	0.535±0.111	0.414±0.084	0.271±0.059	1.380±0.347	0.911±0.115
	HiST(Ours)	0.617±0.094	0.526±0.088	0.390±0.070	1.077±0.155	0.817±0.058

The Training Objective Function The training objective of HiST is twofold: (1) to align pathological image and gene expression profiles across multiple levels by modeling the hierarchical structure of ST data, and (2) to accurately predict gene expression from image features alone. Ultimately, this objective function consists of two components: hierarchical alignment loss and ST prediction loss, defined by:

$$\mathcal{L} = \mathcal{L}_{pred} + \alpha(\mathcal{L}_{HCA} + \beta\mathcal{L}_{HEA}), \quad (11)$$

where α balances the loss components, and β controls the entailment loss effect.

4 EXPERIMENTS AND RESULTS

4.1 EXPERIMENTAL SETTINGS

Dataset To evaluate HiST, we collected three public datasets from the HEST-1K dataset (Jaume et al., 2024a), a high-quality collection of spatial transcriptomics data with standardized processing and rich metadata. (1) Colorectum dataset (Valdeolivas et al., 2024) comprises 14 WSIs ($0.45 \mu\text{m}$ per pixel) with a total of 20,733 spots; (2) Skin (Schäbitz et al., 2022) includes 46 WSIs and over 35,008 spots. The resolution of pathology images is about $0.52 \mu\text{m}/\text{pixel}$; (3) Kidney (Lake et al., 2023) provides 23 WSIs and 25,944 spots at a resolution of approximately $0.76 \mu\text{m}/\text{pixel}$.

ST Preprocessing To account for variations in image resolution across datasets, we adopted a physics-aware patch extraction strategy rather than using fixed pixel dimensions for image cropping. Specifically, we calculated the patch size for each spot based on its physical diameter and crop the corresponding images at their respective resolutions to obtain spot-level image patches. The niche-level patch is created by cropping the region encompassing the central spot and its K-nearest neighbors (determined by spatial coordinates). Subsequently, all extracted patches are resized to a uniform 224×224 pixel resolution. For the gene expression data, we select the top 200 Highly Mean, Highly Variant Genes (HMHVG). Gene expression counts for each spot were subsequently log-transformed.

Evaluation Protocol To ensure robust model evaluation, we performed five independent random splits of the WSI samples for each dataset, allocating 80% for training, 10% for validation, and 10% for testing in each iteration. The exact WSI IDs used for each of the five splits are provided in our code to ensure full reproducibility. Our evaluation metrics include top- k mean Pearson Correlation Coefficient (PCC@ k), mean squared error (MSE), and mean absolute error (MAE), similar to (Zhu et al., 2025; Chung et al., 2024).

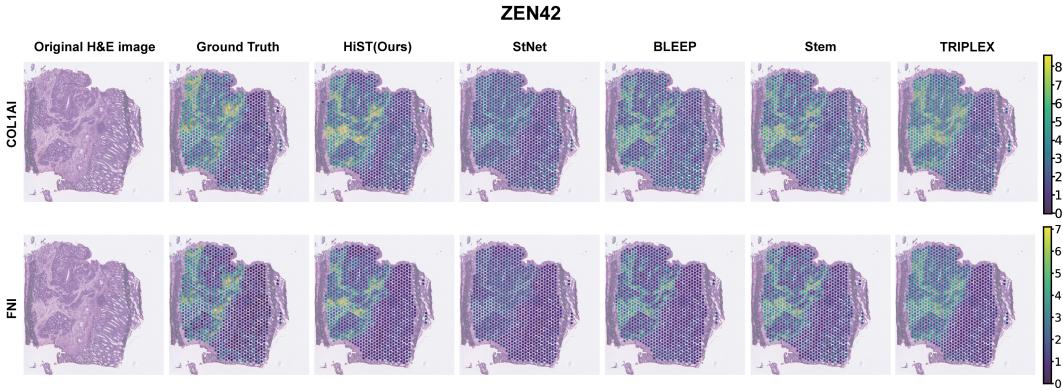


Figure 3: Visualization of the spatial distribution of the COL1A1 (Top) gene and FN1 gene (Bottom) in ZEN42 sample.

4.2 EXPERIMENTAL RESULTS AND VISUALIZATION

Baseline Comparison Table 1 shows that HiST outperforms all existing methods across all datasets (Colorectum, Skin, and Kidney), highlighting HiST’s superior accuracy and robustness in predicting gene expression from pathology images across diverse biological contexts, which were trained under a rigorous, fair comparison protocol (details in Appendix B.2). For instance, compared to the second-best method TRIPLEX, HiST achieves an improvement of 6.33% and 3.24% in PCC@200 on the Kidney and Skin datasets, respectively. Notably, TRIPLEX achieves the next-best performance, underscoring the crucial role of multi-level structure in accurately predicting gene expression. Acknowledging the variance in the results of Table 1, we performed paired t-tests in Appendix C.1 to confirm the statistical significance of HiST’s improvements. To further demonstrate the robustness of our model, we also conducted a series of additional experiments, including an efficiency benchmark (Appendix C.2), cross-laboratory generalization tests (Appendix C.3), robustness analysis on the HVG gene set (Appendix C.4), evaluation with patient-level data splitting (Appendix C.5) and so on.

Biomarker Visualization To further qualitatively assess the model’s behavior, we performed visualization of sample ZEN42 in the Colorectal dataset, focusing on two established colorectal cancer biomarkers, i.e., COL1A1 (Zhang et al., 2018; Pawlak et al., 2025) and FN1 (Sun et al., 2020). Figure 3 demonstrates that HiST more accurately captures the key high-expression regions compared to other methods. More visualizations are available in Section D of Appendix.

Clinical Downstream Task Validation We designed a downstream validation experiment to further validate the clinical utility of the representations learned by HiST. First, we employed our model, pre-trained on the Colorectal dataset, to perform zero-shot inference on H&E slides from an independent external dataset, TCGA-COADREAD (colon and rectal adenocarcinoma). The inferred gene expression profiles were then used to train a Random Forest classifier for predicting microsatellite instability (MSI) status, a critical clinical biomarker for immunotherapy response (Feng et al., 2024). As shown in Table 2, we note that applying some baseline models to this large-scale experiment was computationally infeasible. **TRIPLEX**’s global attention mechanism led to prohibitive memory requirements (>60 GB of VRAM for a single WSI), while **Stem**’s diffusion-based inference was excessively slow (over 230 hours for the entire cohort). The gene profiles predicted by HiST led to significantly better MSI prediction performance compared to other baselines. This result indicates that the representations learned by HiST generalize remarkably well and carry tangible clinical value, underscoring its potential for downstream clinical applications.

Table 2: Performance on MSI Status Classification (AUROC).

Model	AUROC	
	MSI-H	MSS
TRIPLEX	-	-
StNet	0.57 ± 0.08	0.54 ± 0.03
BLEEP	0.55 ± 0.05	0.53 ± 0.06
Stem	-	-
Ours	0.72 ± 0.06	0.60 ± 0.06

432 4.3 ABLATION STUDY
433

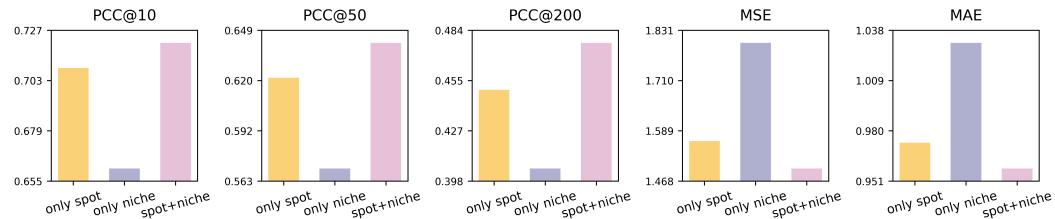
434 We performed an ablation study on the model’s structure and hyperparameters to observe the strategy
435 of alignment, input data for gene decoder and the impact of LoRA. Here, we describe the results in
436 Colorectum dataset. The ablation study on the impact of LoRA, as well as more experiment results
437 on other datasets, can be found in Appendix E.

438
439 **Strategy of Alignment** We compared the different alignment strategies including: a) removing
440 only the gene-image regularization term of the HEA loss (w/o G-I HEA), b) removing the entire HEA
441 loss (w/o HEA), c) removing the entire Hierarchical Hyperbolic Alignment (HHA) module (w/o HEA
442 + HCA), d) replacing the HHA module by a MERU variant in Hyperbolic Space (without multil-level
443 representation learning) (Desai et al., 2023) and e) replacing the HHA module by a CLIP variant in
444 Euclidean Space (Radford et al., 2021), as shown in Table 3. The results imply that our approach
445 can learn better representations by leveraging inherent hierarchies, enhancing the overall model
446 performance. Notably, the performance gap between our full model and its Euclidean counterpart
447 (CLIP) strongly validates our core hypothesis on the superiority of hyperbolic space for this task.
448

449 Table 3: Ablation study of the alignment strategy.

Alignment	PCC@10 \uparrow	PCC@50 \uparrow	PCC@200 \uparrow	MSE \downarrow	MAE \downarrow
w/o G-I HEA	0.708 \pm 0.116	0.633 \pm 0.139	0.470 \pm 0.186	1.523 \pm 0.471	0.972 \pm 0.160
w/o HEA	0.708 \pm 0.118	0.627 \pm 0.142	0.465 \pm 0.186	1.523 \pm 0.444	0.969 \pm 0.154
w/o HEA + HCA	0.697 \pm 0.141	0.615 \pm 0.162	0.456 \pm 0.201	1.675 \pm 0.810	0.997 \pm 0.250
MERU	0.705 \pm 0.101	0.618 \pm 0.159	0.451 \pm 0.142	1.535 \pm 0.341	0.970 \pm 0.125
CLIP	0.693 \pm 0.093	0.605 \pm 0.102	0.441 \pm 0.135	1.523 \pm 0.219	1.011 \pm 0.124
Ours	0.721\pm0.105	0.642\pm0.128	0.477\pm0.184	1.498\pm0.456	0.958\pm0.158

450
451
452
453
454
455
456
457
458
459
460 **Input of Decoder** We evaluated the impact of different input strategies on the decoder’s perfor-
461 mance, including using a) only spot-level image, b) only niche-level images and c) the combination of
462 both. As shown in Figure 4, the results reveal that our combined approach yields the best performance
463 across all metrics. In summary, this ablation study confirms that the integration of both spot-level and
464 niche-level information is critical for achieving optimal performance.
465

473 Figure 4: Ablation study of the input data of decoder.
474475 5 CONCLUSION
476

477 We present HiST, a novel framework that leverages multi-level hyperbolic representations to predict
478 spatial transcriptomics from histology images. By modeling the intrinsic hierarchical structure
479 of ST data within hyperbolic space, HiST learns more comprehensive spatial histological and
480 genetic features. Our comprehensive experimental evaluation demonstrates that HiST consistently
481 outperforms state-of-the-art approaches, underscoring the potential of geometric deep learning in
482 spatial omics analysis. We hope this framework inspires future research toward geometric-aware
483 multimodal learning in the biological domain, harnessing the inherent geometry of biological systems
484 for designing more sophisticated representation models.
485

486 ETHICS STATEMENT
487488 All data utilized in this study, including Whole Slide Images and their corresponding gene expression
489 profiles, were sourced from publicly available and properly cited datasets. No new data involving
490 human subjects was collected for this research, and all data was fully anonymized, containing no
491 personally identifiable information. Our work fully complies with the licensing and terms of use for
492 all original data sources.493
494 REPRODUCIBILITY STATEMENT
495496 We are committed to ensuring the reproducibility of this research. The code, data details, and
497 implementation specifics required to reproduce our experimental results are as follows:
498499 **Code Availability** The complete source code and experiment scripts are available at the following
500 anonymous repository: <https://anonymous.4open.science/r/12116-V2>.
501502 **Dataset and Preprocessing** Our research is based on publicly available datasets. Their detailed
503 descriptions, sources, and the full data preprocessing pipeline (including patch extraction and gene
504 selection) are elaborated in Section 4.1 and Appendix B.1.505 **Implementation Details** Specifics of the model implementation, including the hyperparameter
506 settings for all experiments, and the detailed training procedure, are provided in Appendix B.
507508 We believe this information is sufficient to support the verification and extension of our work.
509510 REFERENCES
511512 Jordan T Ash, Gregory Darnell, Daniel Munro, and Barbara E Engelhardt. Joint analysis of expres-
513 sion levels and histological images identifies genes associated with tissue morphology. *Nature*
514 *communications*, 12(1):1609, 2021.515 Victor Hugo Canela, William S Bowen, Ricardo Melo Ferreira, Farooq Syed, James E Lingeman,
516 Angela R Sabo, Daria Barwinska, Seth Winfree, Blue B Lake, Ying-Hua Cheng, et al. A spatially
517 anchored transcriptomic atlas of the human kidney papilla identifies significant immune injury in
518 patients with stone disease. *Nature communications*, 14(1):4140, 2023.519 James W Cannon, William J Floyd, Richard Kenyon, Walter R Parry, et al. Hyperbolic geometry.
520 *Flavors of geometry*, 31(59-115):2, 1997.
521522 Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Andrew H Song,
523 Bowen Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, et al. Towards a general-purpose
524 foundation model for computational pathology. *Nature Medicine*, 30(3):850–862, 2024a.525 Wei-Ting Chen, Ashley Lu, Katileen Craessaerts, Benjamin Pavie, Carlo Sala Frigerio, Nikky Corthout,
526 Xiaoyan Qian, Jana Laláková, Malte Kühnemund, Iryna Voytyuk, et al. Spatial transcriptomics
527 and in situ sequencing to study alzheimer’s disease. *Cell*, 182(4):976–991, 2020.528 Ying Chen, Jiajing Xie, Yuxiang Lin, Yuhang Song, Wenxian Yang, and Rongshan Yu. Survmbamba:
529 State space model with multi-grained multi-modal interaction for survival prediction. *arXiv*
530 *preprint arXiv:2404.08027*, 2024b.
531532 Seunghyuk Cho, Juyong Lee, Jaesik Park, and Dongwoo Kim. A rotated hyperbolic wrapped normal
533 distribution for hierarchical representation learning. *Advances in Neural Information Processing*
534 *Systems*, 35:17831–17843, 2022.535 Kyongho Choe, Unil Pak, Yu Pang, Wanjun Hao, and Xiuqin Yang. Advances and challenges in
536 spatial transcriptomics for developmental biology. *Biomolecules*, 13(1):156, 2023.
537538 Youngmin Chung, Ji Hun Ha, Kyeong Chan Im, and Joo Sang Lee. Accurate spatial gene expression
539 prediction by integrating multi-resolution features. In *Proceedings of the IEEE/CVF Conference*
on *Computer Vision and Pattern Recognition*, pp. 11591–11600, 2024.

540 Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, and Shanmukha Ramakrishna
 541 Vedantam. Hyperbolic image-text representations. In *International Conference on Machine*
 542 *Learning*, pp. 7694–7731. PMLR, 2023.

543

544 Yu Feng, Wenjuan Ma, Yupeng Zang, Yanying Guo, Young Li, Yixuan Zhang, Xuan Dong, Yi Liu,
 545 Xiaojuan Zhan, Zhizhong Pan, et al. Spatially organized tumor-stroma boundary determines the
 546 efficacy of immunotherapy in colorectal cancer patients. *Nature Communications*, 15(1):10259,
 547 2024.

548 Masayuki Fujii, Shigeki Sekine, and Toshiro Sato. Decoding the basis of histological variation in
 549 human cancer. *Nature Reviews Cancer*, 24(2):141–158, 2024.

550

551 Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic entailment cones for learning
 552 hierarchical embeddings. In *International conference on machine learning*, pp. 1646–1655. PMLR,
 553 2018.

554 Bryan He, Ludvig Bergenstråhle, Linnea Stenbeck, Abubakar Abid, Alma Andersson, Åke Borg,
 555 Jonas Maaskola, Joakim Lundeberg, and James Zou. Integrating spatial gene expression and breast
 556 tumour morphology via deep learning. *Nature biomedical engineering*, 4(8):827–834, 2020.

557

558 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 559 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 560 pp. 770–778, 2016.

561

562 Joy Hsu, Jeffrey Gu, Gong Wu, Wah Chiu, and Serena Yeung. Capturing implicit hierarchical
 563 structure in 3d biomedical images with self-supervised hyperbolic representations. *Advances in
 564 neural information processing systems*, 34:5112–5123, 2021.

565

566 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 567 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

568

569 Guillaume Jaume, Paul Doucet, Andrew H. Song, Ming Y. Lu, Cristina Almagro-Perez, Sophia J.
 570 Wagner, Anurag J. Vaidya, Richard J. Chen, Drew F. K. Williamson, Ahrong Kim, and Faisal
 Mahmood. Hest-1k: A dataset for spatial transcriptomics and histology image analysis. In
 571 *Advances in Neural Information Processing Systems*, December 2024a.

572

573 Guillaume Jaume, Anurag Vaidya, Richard J Chen, Drew FK Williamson, Paul Pu Liang, and Faisal
 574 Mahmood. Modeling dense multimodal interactions between biological pathways and histology
 575 for survival prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
 576 Pattern Recognition*, pp. 11579–11590, 2024b.

577

578 Valentin Khrulkov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky.
 579 Hyperbolic image embeddings. In *Proceedings of the IEEE/CVF conference on computer vision
 580 and pattern recognition*, pp. 6418–6428, 2020.

581

582 Blue B Lake, Rajasree Menon, Seth Winfree, Qiwen Hu, Ricardo Melo Ferreira, Kian Kalhor, Daria
 583 Barwinska, Edgar A Otto, Michael Ferkowicz, Dinh Diep, et al. An atlas of healthy and injured
 584 cell states and niches in the human kidney. *Nature*, 619(7970):585–594, 2023.

585

586 Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436–444, 2015.

587

588 Hao Li, Ying Chen, Yifei Chen, Rongshan Yu, Wenxian Yang, Liansheng Wang, Bowen Ding, and
 589 Yuchen Han. Generalizable whole slide image classification with fine-grained visual-semantic inter-
 590 action. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 591 pp. 11398–11407, 2024.

592

593 Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven
 594 Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum
 595 distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021.

596

597 Yuxiang Lin, Ling Luo, Ying Chen, Xushi Zhang, Zihui Wang, Wenxian Yang, Mengsha Tong, and
 598 Rongshan Yu. St-align: A multimodal foundation model for image-gene alignment in spatial
 599 transcriptomics. *arXiv preprint arXiv:2411.16793*, 2024.

594 Ming Y Lu, Bowen Chen, Drew FK Williamson, Richard J Chen, Ivy Liang, Tong Ding, Guillaume
 595 Jaume, Igor Odintsov, Long Phi Le, Georg Gerber, et al. A visual-language foundation model for
 596 computational pathology. *Nature Medicine*, 30(3):863–874, 2024.

597

598 Pascal Mettes, Mina Ghadimi Atigh, Martin Keller-Ressel, Jeffrey Gu, and Serena Yeung. Hyperbolic
 599 deep learning in computer vision: A survey. *International Journal of Computer Vision*, 132(9):
 600 3484–3508, 2024.

601

602 Taku Monjo, Masaru Koido, Satoi Nagasawa, Yutaka Suzuki, and Yoichiro Kamatani. Efficient
 603 prediction of a spatial transcriptomics profile better characterizes breast cancer tissue sections
 604 without costly experimentation. *Scientific reports*, 12(1):4133, 2022.

605

606 Maximillian Nickel and Douwe Kiela. Learning continuous hierarchies in the lorentz model of
 607 hyperbolic geometry. In *International conference on machine learning*, pp. 3779–3788. PMLR,
 608 2018.

609

610 Ajit J Nirmal, Zoltan Maliga, Tuulia Vallius, Brian Quattrochi, Alyce A Chen, Connor A Jacobson,
 611 Roxanne J Pelletier, Clarence Yapp, Raquel Arias-Camison, Yu-An Chen, et al. The spatial
 612 landscape of progression and immunoediting in primary melanoma at single-cell resolution.
Cancer discovery, 12(6):1518–1541, 2022.

613

614 Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive
 615 coding. *arXiv preprint arXiv:1807.03748*, 2018.

616

617 Avik Pal, Max van Spengler, Guido Maria D’Amely di Melendugno, Alessandro Flaborea, Fabio
 618 Galasso, and Pascal Mettes. Compositional entailment learning for hyperbolic vision-language
 619 models. *arXiv preprint arXiv:2410.06912*, 2024.

620

621 Martyna Pawlak, Żaneta Kałuzińska-Kołat, Zbigniew W Pasieka, Damian Kołat, and Elżbieta
 622 Płuciennik. The critical role of col1a1 revealed by integrated bioinformatics analysis of
 623 differentially-expressed genes in colorectal cancer and inflammatory bowel disease. *Computers
 in Biology and Medicine*, 190:110116, 2025.

624

625 Marija Pizurica, Yuanning Zheng, Francisco Carrillo-Perez, Humaira Noor, Wei Yao, Christian
 626 Wohlfart, Antoaneta Vladimirova, Kathleen Marchal, and Olivier Gevaert. Digital profiling of
 627 gene expression from histology images with linearized attention. *Nature Communications*, 15(1):
 628 9886, 2024.

629

630 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 631 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 632 models from natural language supervision. In *International conference on machine learning*, pp.
 8748–8763. PmLR, 2021.

633

634 Alexander Schäbitz, C Hillig, M Mubarak, Manja Jargosch, Ali Farnoud, Emanuele Scala, Nils
 635 Kurzen, Anna Caroline Pilz, Nayanika Bhalla, Jenny Thomas, et al. Spatial transcriptomics
 636 landscape of lesions from non-communicable inflammatory skin diseases. *Nature Communications*,
 13(1):7729, 2022.

637

638 Patrik L Ståhl, Fredrik Salmén, Sanja Vickovic, Anna Lundmark, José Fernández Navarro, Jens Mag-
 639 nusson, Stefania Giacomello, Michaela Asp, Jakub O Westholm, Mikael Huss, et al. Visualization
 640 and analysis of gene expression in tissue sections by spatial transcriptomics. *Science*, 353(6294):
 78–82, 2016.

641

642 Yang Sun, Chunlin Zhao, Yanwei Ye, Zhen Wang, Yuanhang He, Yulin Li, and Haoxun Mao. High
 643 expression of fibronectin 1 indicates poor prognosis in gastric cancer. *Oncology Letters*, 19(1):
 644 93–102, 2020.

645

646 Zhonghui Tang, Yinqi Bai, Qi Fang, Yuchen Yuan, Qianwen Zeng, Shuling Chen, Tianyi Xu, Jianyu
 647 Chen, Li Tan, Chunqing Wang, et al. Spatial transcriptomics reveals tryptophan metabolism
 648 restricting maturation of intratumoral tertiary lymphoid structures. *Cancer Cell*, 2025.

648 Alberto Valdeolivas, Bettina Amberg, Nicolas Giroud, Marion Richardson, Eric JC Gálvez, Solveig
 649 Badillo, Alice Julien-Laferrière, Demeter Túró, Lena Voith von Voithenberg, Isabelle Wells,
 650 et al. Profiling the heterogeneity of colorectal cancer consensus molecular subtypes using spatial
 651 transcriptomics. *NPJ precision oncology*, 8(1):10, 2024.

652 Chuhan Wang, Adam S Chan, Xiaohang Fu, Shila Ghazanfar, Jinman Kim, Ellis Patrick, and Jean YH
 653 Yang. Benchmarking the translational potential of spatial gene expression prediction from histology.
 654 *Nature Communications*, 16(1):1544, 2025a.

655 Hongyi Wang, Xiuju Du, Jing Liu, Shuyi Ouyang, Yen-Wei Chen, and Lanfen Lin. M2ost: Many-to-
 656 one regression for predicting spatial transcriptomics from digital pathology images. In *Proceedings*
 657 of the AAAI Conference on Artificial Intelligence, volume 39, pp. 7709–7717, 2025b.

658 Cameron G Williams, Hyun Jae Lee, Takahiro Asatsuma, Roser Vento-Tormo, and Ashraful Haque.
 659 An introduction to spatial transcriptomics for biomedical research. *Genome medicine*, 14(1):68,
 660 2022.

661 Sijia Wu, Jiajin Zhang, Yanfei Wang, Xinyu Qin, Zhaocan Zhang, Zhennan Lu, Pora Kim, Xiaobo
 662 Zhou, and Liyu Huang. metsdb: a knowledgebase of cancer metastasis at bulk, single-cell and
 663 spatial levels. *Nucleic Acids Research*, 53(D1):D1427–D1434, 2025.

664 Ronald Xie, Kuan Pang, Sai Chung, Catia Perciani, Sonya MacParland, Bo Wang, and Gary Bader.
 665 Spatially resolved gene expression prediction from histology images via bi-modal contrastive
 666 learning. *Advances in Neural Information Processing Systems*, 36:70626–70637, 2023.

667 Jiazhou Ye, Yan Lin, Zhiling Liao, Xing Gao, Cheng Lu, Lu Lu, Julu Huang, Xi Huang, Shilin
 668 Huang, Hongping Yu, et al. Single cell-spatial transcriptomics and bulk multi-omics analysis of
 669 heterogeneity and ecosystems in hepatocellular carcinoma. *NPJ Precision Oncology*, 8(1):262,
 670 2024.

671 Linlin Zhang, Dongsheng Chen, Dongli Song, Xiaoxia Liu, Yanan Zhang, Xun Xu, and Xiangdong
 672 Wang. Clinical and translational values of spatial transcriptomics. *Signal Transduction and*
 673 *Targeted Therapy*, 7(1):111, 2022.

674 Yikun Zhang, Geyan Ye, Chaohao Yuan, Bo Han, Long-Kai Huang, Jianhua Yao, Wei Liu, and
 675 Yu Rong. Atomas: Hierarchical adaptive alignment on molecule-text for unified molecule under-
 676 standing and generation. In *The Thirteenth International Conference on Learning Representations*,
 677 2025.

678 Zheyang Zhang, Yongxia Wang, Jinghang Zhang, Jiateng Zhong, and Rui Yang. Col1a1 promotes
 679 metastasis in colorectal cancer by regulating the wnt/pcp pathway. *Molecular Medicine Reports*,
 680 17(4):5037–5042, 2018.

681 Sichen Zhu, Yuchen Zhu, Molei Tao, and Peng Qiu. Diffusion generative modeling for spatially
 682 resolved gene expression inference from histology images. *arXiv preprint arXiv:2501.15598*,
 683 2025.

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 Appendix

703 CONTENTS

707 A Exponential Map Derivations	15
--	-----------

709 B Implementation Details	15
-------------------------------------	-----------

711 B.1 Implementation Details for Gene Selection	15
712 B.2 Implementation Details for Experiments	15
714 B.3 Implementation Details for Ablation Studies	18
715 B.4 Implementation Details for Evaluation Metrology	18

717 C Additional Experiment Results	19
--	-----------

719 C.1 Statistical Significance Analysis	19
720 C.2 Efficiency Benchmark	19
722 C.3 Cross-Laboratory Generalization	19
723 C.4 Robustness on HVG Gene Set	20
725 C.5 Evaluation with Patient-Level Data Splitting	20
726 C.6 Scalability with Number of Target Genes	21
728 C.7 Sensitivity to Neighborhood Size	21
729 C.8 Quantitative Gene-wise Analysis	22
730 C.9 Biological Validation on Key Marker Genes	22

732 D Additional Visualization	22
---------------------------------------	-----------

734 E Additional Ablation Study	25
--	-----------

736 F Limitation	27
-------------------------	-----------

739 G The Use of Large Language Models (LLMs)	27
--	-----------

741 H Broader Impact	27
-----------------------------	-----------

743 H.1 Impact on Real-world Applications	27
744 H.2 Impact on Future Research	27

756 **A EXPONENTIAL MAP DERIVATIONS**

758 In this section, we present the derivation of the exponential map (Equation 5) in our approach Desai
 759 et al. (2023). Let the vector $\mathbf{v} = [0, \mathbf{v}_{euc}] \in \mathbb{R}^{n+1}$ denotes the extension of the Euclidean embedding
 760 $\mathbf{v}_{euc} \in \mathbb{R}^n$. This vector belongs in the tangent space at the origin $\mathbf{O} = [\sqrt{1/c}, 0, \dots, 0] \in \mathbb{R}^{n+1}$
 761 of the hyperboloid as the Lorentzian inner product of these two vectors is zero, where $-c < 0$ is
 762 the curvature of the hyperboloid. Based on Equation 4, we can simplify the exponential map by
 763 considering only the space components:

$$764 \mathbf{x}_{space} = \cosh(\sqrt{c}\|\mathbf{v}\|_{\mathbb{L}})0 + \frac{\sinh(\sqrt{c}\|\mathbf{v}\|_{\mathbb{L}})}{\sqrt{c}\|\mathbf{v}\|_{\mathbb{L}}}\mathbf{v}_{euc}, \quad (12)$$

767 where the first term of this equation is zero. The Lorentzian norm of \mathbf{v} is equal to the Euclidean norm
 768 of space components:

$$769 \|\mathbf{v}\|_{\mathbb{L}} = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle_{\mathbb{L}}} = \sqrt{0 + \langle \mathbf{v}, \mathbf{v} \rangle_{\mathbb{E}}} = \|\mathbf{v}_{euc}\|, \quad (13)$$

771 where $\langle \cdot, \cdot \rangle_{\mathbb{E}}$ denotes standard Euclidean inner product. Therefore, this exponential map can be
 772 formulated as:

$$773 \mathbf{x}_{space} = \exp_{\mathbf{O}}^c(\mathbf{v}_{euc}) = \frac{\sinh(\sqrt{c}\|\mathbf{v}_{euc}\|_{\mathbb{E}})}{\sqrt{c}\|\mathbf{v}_{euc}\|_{\mathbb{E}}}\mathbf{v}_{euc}. \quad (14)$$

775 **B IMPLEMENTATION DETAILS**

776 **B.1 IMPLEMENTATION DETAILS FOR GENE SELECTION**

777 The gene expression profiles of Spatial Transcriptomics (ST) typically contain approximately 20,000
 778 to 30,000 genes, most of which exhibit low variability. Directly feeding all these genes into a
 779 deep learning model would lead to the severe curse of dimensionality. To address this issue, we
 780 implemented two gene selection strategies based on previous work (Zhu et al., 2025): 1) The top
 781 200 genes with both high mean expression and high variability (MHHVG); 2) The top 300 genes
 782 selected from all highly variable genes (HVG) ranked by mean expression level. Specifically, for each
 783 WSI, we extracted the top 2,000 highly variable genes based on its corresponding gene expression
 784 profile. These highly variable gene sets are then pooled across all WSIs to form a union set. From
 785 this union set, we rank genes by mean expression and variance independently to identify the top 300
 786 genes in each category. The top 300 genes with the highest mean expression are designated as highly
 787 expressed genes (HEG), and the top 300 genes with the highest variance are designated as highly
 788 variable genes (HVG). We then take the intersection of these two sets to define the Highly Mean and
 789 Highly Variable Genes (HMHVG). This approach ensures a robust and consistent gene selection for
 790 downstream analyses. Figure 5 displays the Highly Mean and Highly Variable Genes (HMHVG) set,
 791 while Figure 6 shows the selected Highly Variable Genes (HVG) set.

794 **B.2 IMPLEMENTATION DETAILS FOR EXPERIMENTS**

795 We compare HiST against four state-of-the-art methods: TRIPLEX (Chung et al., 2024), StNet (He
 796 et al., 2020), BLEEP (Xie et al., 2023), and Stem (Zhu et al., 2025). To ensure a fair and rigorous
 797 comparison, we re-implemented all baseline models and trained them from scratch using the exact
 798 same data splits described in our evaluation protocol of Section 4.1. For each method, we meticulously
 799 followed the hyperparameter tuning strategies outlined in their original publications, adapting
 800 their official public code where available. This standardized setup guarantees that all performance
 801 differences can be attributed to model architecture and learning strategy rather than variations in data
 802 or implementation.

803 Our model is trained with AdamW with an initial learning rate of 0.0001. The batch size is 128 and
 804 the hidden embedding channel is 1024. The hyperparameters of α and β are 0.2 and 0.4 respectively.
 805 As for LoRA, we only adapt the attention weights of the foundational model with rank 4, applying it
 806 to the last 11 attention layers. Besides, we implement mixed precision training using PyTorch’s AMP
 807 for accelerated computation, with all experiments seeded at 42 for reproducibility. The epochs of
 808 training are up to 200 with early stopping, and we assume models have converged when the validation
 809 loss fails to improve for 10 consecutive epochs. All experiments are trained on RTX4090 GPUs.

810
811
812
813
814
815
816
817
818

819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863	Dataset 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863	Genes to be predicted 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863
	Colorectum	'A2M', 'ACTA2', 'AEBP1', 'AGR2', 'AHNAK', 'ANXA11', 'APOE', 'ASS1', 'ATP1B1', 'ATP5ME', 'ATP6V0C', 'B2M', 'BCAP31', 'BGN', 'BST2', 'BTF3', 'C15orf48', 'C19orf33', 'C1OA', 'C1QB', 'C1R', 'C1S', 'C3', 'CALD1', 'CALM1', 'CCN2', 'CD24', 'CD44', 'CD55', 'CD59', 'CD74', 'CD81', 'CD99', 'CDH17', 'CEACAM5', 'CEACAM6', 'CEACAM7', 'CKB', 'CLCA1', 'CLDN3', 'CLDN4', 'CLDN7', 'COL12A1', 'COL18A1', 'COL1A1', 'COL1A2', 'COL3A1', 'COL4A1', 'COL4A2', 'COL5A1', 'COL5A2', 'COL6A1', 'COL6A2', 'COL6A3', 'COMM66', 'COX5A', 'COX5B', 'COX6B1', 'COX7B', 'CRIP1', 'CSTB', 'CTSB', 'CTSD', 'CXCL1', 'CXCL10', 'DBI', 'DCN', 'DEFA5', 'DMBT1', 'DUOX2', 'DUSP1', 'EEF1A1', 'EGR1', 'EIF3', 'ENO1', 'EPCAM', 'FABP1', 'FAM3D', 'FBLN1', 'FCGBP', 'FLNA', 'FN1', 'FOS', 'FSTL1', 'FTH1', 'FTL', 'FYXD3', 'GPRC5A', 'GPX2', 'GREM1', 'GSN', 'HLA-A', 'HLA-B', 'HLA-C', 'HLA-DPA1', 'HLA-DPB1', 'HLA-DQB1', 'HLA-DRA', 'HLA-DRB1', 'HMGN2', 'HNRNPH1', 'HNRNPU', 'HSPB1', 'HSPD1', 'ID1', 'IDO1', 'IER2', 'IER3', 'IF130', 'IF16', 'IGFBP4', 'IGFBP5', 'IGFBP7', 'IGHA1', 'IGHG1', 'IGHG3', 'IGHG4', 'IGHM', 'IGKC', 'IGLC1', 'IGLC2', 'IGLC3', 'IL32', 'IRF1', 'ISG15', 'ITLN1', 'ITM2B', 'JCHAIN', 'JUNB', 'JUND', 'JUP', 'KDEL2R', 'KLF5', 'KRT18', 'KRT19', 'KRT8', 'LCN2', 'LDHB', 'LGALS1', 'LGALS3', 'LGALS4', 'LUM', 'LYZ', 'MALAT1', 'MARCKSL1', 'MDK', 'MGP', 'MMP1', 'MMP11', 'MMP2', 'MUC1', 'MUC12', 'MUC13', 'MUC2', 'MUC5B', 'MYL12B', 'MYL9', 'NBL1', 'NDUFA13', 'NORAD', 'NRA4A', 'OLFM4', 'PABPC1', 'PERP', 'PGK1', 'PHGR1', 'PI3', 'PIGR', 'PLA2G2A', 'POSTN', 'PP1B', 'PRSS8', 'PSMB8', 'PTMA', 'RBM3', 'REG1A', 'REG3A', 'REG4', 'RHOA', 'RHOC', 'ROM01', 'RRBP1', 'S100A10', 'S100A6', 'S100A8', 'S100A9', 'S100P', 'SAT1', 'SELENOP', 'SERPINA1', 'SERPING1', 'SFN', 'SLC12A2', 'SMIM22', 'SNRPP', 'SOD2', 'SPARC', 'SPINK1', 'SPINK4', 'SPINT2', 'SPTB1N', 'ST14', 'STAT1', 'SULF1', 'TAGLN', 'TFF1', 'TFI3', 'TGFBI', 'THBS1', 'THY1', 'TIMP1', 'TIMP2', 'TM16B', 'TMEM54', 'TMEM59', 'TMSB4X', 'TPM2', 'TPT1', 'TSPAN1', 'TSPAN3', 'TSPAN8', 'TSPN', 'TYMP', 'UBD', 'UQCR11', 'UQCRH', 'VIM', 'YWHAB', 'ZFP36', 'ZG16'
	Skin	'ACTB', 'ACTG1', 'AHNAK', 'ANXA1', 'ANXA2', 'APRT', 'AQP3', 'ARPC2', 'ASPRV1', 'ATP1B3', 'ATP5F1A', 'ATP5F1B', 'ATP5F1E', 'ATP5MC2', 'ATP5MC3', 'ATP5ME', 'ATP5MF', 'ATP5MG', 'B2M', 'BTF3', 'C19orf33', 'CALM1', 'CALML3', 'CALML5', 'CASP14', 'CCL27', 'CD24', 'CD44', 'CD74', 'CD9', 'CDSN', 'CFL1', 'CHCHD2', 'CLTB', 'CNBP', 'CNFN', 'COL17A1', 'COL1A1', 'COL1A2', 'COL3A1', 'COL6A1', 'COL6A2', 'COX41', 'COX5B', 'COX6A1', 'COX6B1', 'COX6C', 'COX7A2', 'COX7B', 'COX7C', 'COX8A', 'CRABP2', 'CSNK1A1', 'CST3', 'CST6', 'CSTA', 'CSTB', 'CTNNBIP1', 'CTS1', 'CXCL14', 'DBI', 'DCD', 'DCN', 'DEFB4A', 'DEGS1', 'DMKN', 'DSC1', 'DSC3', 'DSC4', 'DSC5', 'DSP', 'DSTN', 'DYNLL1', 'EEF1A1', 'EEF1B2', 'EEF1D', 'EEF2', 'EIF1', 'EIF3E', 'EIF3F', 'EIF4G2', 'EIF5A', 'ELOB', 'ENO1', 'FABP5', 'FADS2', 'FAM25A', 'FAU', 'FLG', 'FLG2', 'FTH1', 'FTL', 'GAPDH', 'GJA1', 'GJB2', 'GLTP', 'GPNNMB', 'GSTP1', 'GUK1', 'H3-3A', 'H3-3B', 'HINT1', 'HLA-B', 'HLA-C', 'HLA-DPA1', 'HLA-DPB1', 'HLA-DQB1', 'HLA-DRA', 'HLA-E', 'HMG81', 'HNRNPA1', 'HNRNPA2B1', 'HNRNPK', 'HOPX', 'HSP90AA1', 'HSP90AB1', 'HSPA8', 'HSPB1', 'IF127', 'IF16', 'IFITM3', 'IGKC', 'ITM2B', 'IVL', 'KLF5', 'KLK1', 'KRT1', 'KRT14', 'KRT16', 'KRT2', 'KRT5', 'KRT6A', 'KRT6B', 'KRT7C', 'KRTDAP', 'LAD1', 'LCE1B', 'LCE3D', 'LDHA', 'LGALS1', 'LGALS3', 'LGALS7B', 'LMNA', 'LY6D', 'LYPD3', 'LYZ', 'MIF', 'MUC1', 'MYL3', 'NACA', 'NCCR14', 'NDUFA4', 'NDUFS5', 'NQP53', 'NP1M1', 'OAZ1', 'P4HB', 'PABPC1', 'PERP', 'PFDN5', 'PFN1', 'PGM1', 'PI3', 'PKM', 'PKP1', 'PLP2', 'POLR2L', 'PPDPF', 'PPIA', 'PPL', 'PRDX1', 'PSAP', 'PSMA7', 'PTMA', 'RAB11A', 'RAC1', 'RACK1', 'RAN', 'RBM3', 'ROM01', 'RTN4', 'S100A10', 'S100A11', 'S100A14', 'S100A16', 'S100A2', 'S100A4', 'S100A9', 'SBSN', 'SCGB2A2', 'SDC1', 'SELENOW', 'SERBP1', 'SERF1', 'SERPINB3', 'SERPINB4', 'SERPINB5', 'SFN', 'SH3BGRL3', 'SLC25A3', 'SLC25A5', 'SLC25A6', 'SLC2A1', 'SLP1', 'SLURP1', 'SNRDP2', 'SPARC', 'SPINK5', 'SPINT2', 'SPRR1B', 'SPRR2A', 'SPRR2B', 'SPRR2D', 'SPRR2E', 'SPRR2G', 'SUB1', 'TACSTD2', 'TAGLN2', 'TMA7', 'TMBM6', 'TMEM45A', 'TMSB10', 'TMSB4X', 'TOMM7', 'TPH1', 'TP1', 'TRIM29', 'TSPN', 'TUBA1B', 'TUBA1C', 'TUBB4B', 'TXN', 'TYMP', 'UBA52', 'UBB', 'UBC', 'UBL5', 'UQCR10', 'UQCR2', 'UQCRB', 'UQCRD', 'VIM', 'YBX1', 'YBX3', 'YWHAB', 'YWHAZ'
	Kidney	'A2M', 'ACADVL', 'ACTA2', 'ACTB', 'ACTG1', 'ADGRG1', 'ADIRF', 'AEBP1', 'ALDOB', 'ANPEP', 'ANXA2', 'APLP2', 'APOE', 'APP', 'AQP1', 'AQP2', 'ANSA', 'ASS1', 'ATP1A1', 'ATP1B1', 'ATP5F1D', 'ATP5MC3', 'ATP5ME', 'ATP5MF', 'ATP6V0C', 'B2M', 'BCAM', 'BGN', 'C1OA', 'C1R', 'C2', 'CALB1', 'CALD1', 'CALM1', 'CALM2', 'CANX', 'CD151', 'CD24', 'CD81', 'CDK11', 'CDK16', 'CHCHD2', 'CIRBP', 'CKB', 'CLCNK2', 'CLU', 'COL1A2', 'COL3A1', 'COL4A1', 'COL4A2', 'COL5A1', 'COL6A1', 'COX5B', 'COX6A1', 'COX6C', 'COX7A2', 'COX7B', 'COX7C', 'COX8A', 'CRIM1', 'CRYAB', 'CST3', 'CTS1', 'CXCL14', 'CYSTM1', 'DCN', 'DXD17', 'DDX6', 'DEFB1', 'DSTN', 'DUSP1', 'DYNLL1', 'EEF1A1', 'EEF1D', 'EEF2', 'EFHD1', 'EIF4A1', 'ENG', 'EPAS1', 'EZR', 'FABP1', 'FAU', 'FLNA', 'FTH1', 'FTL', 'FYXD2', 'FYXD4', 'GABARAP', 'GATM', 'GHITM', 'GPX3', 'GSTM3', 'GSTP1', 'GTF2I', 'HINT1', 'HLA-A', 'HLA-B', 'HLA-C', 'HLA-DPA1', 'HLA-DRA', 'HLA-DRB1', 'HLA-E', 'HNRNPA1', 'HNRNPA2B1', 'HSD11B2', 'HSP90AB1', 'HSPAB', 'HSPB1', 'HTR1A', 'IDH2', 'IFITM2', 'IFITM3', 'IGFBP2', 'IGFBP4', 'IGFBP7', 'IGHA1', 'IGHG1', 'IGHG4', 'IGKC', 'IGLC1', 'IGLC2', 'IGLC3', 'ITGA3', 'ITGB1', 'ITM2B', 'IVNS1ABP', 'KCNJ1', 'KCNJ15', 'KNG1', 'LAMP1', 'LAMTOR5', 'LAPTM4A', 'LDHA', 'LGALS1', 'LRP2', 'LUM', 'MAL', 'MALAT1', 'MGP', 'MGST1', 'MGST3', 'MIF', 'MIOX', 'MMP7', 'MUC1', 'MYL12A', 'MYL6', 'MYL9', 'MZT2B', 'NAT8', 'NDRG1', 'NDUFA1', 'NDUFA13', 'NDUFA2', 'NDUFA2', 'NDUFA4', 'NDUFA6', 'NDUFB2', 'NDUFB7', 'NDUFB8', 'NDUFB9', 'NDUFC1', 'NDUFS5', 'NDUFS6', 'NDUFS7', 'NEAT1', 'NME2', 'NPC2', 'OAZ1', 'OGDH', 'OST4', 'P4HB', 'PBP1', 'PCK1', 'PDZK1IP1', 'PEBP1', 'PEPD', 'PFN1', 'PGK1', 'PHTP1', 'PIGR', 'PODXL', 'POLR2L', 'PPP1R1A', 'PTGDS', 'PTH1R', 'RER', 'RHCG', 'RHOA', 'RNASE1', 'ROM01', 'RTNA', 'S100A10', 'S100A2', 'S100A6', 'SAT1', 'SDC1', 'SELENOP', 'SERPINA1', 'SERPINA5', 'SFRP1', 'SLC12A1', 'SLC12A3', 'SLC25A3', 'SLC25A5', 'SLC25A6', 'SLC2A1', 'SLC6A12', 'SMIM24', 'SNHG26', 'SOD1', 'SOD2', 'SPARC', 'SPINK1', 'SPP1', 'SRP14', 'SSR4', 'SUCG1', 'TAGLN', 'TAGLN2', 'THY1', 'TIMP1', 'TIMP2', 'TIMP3', 'TINAGL1', 'TMA7', 'TMSB10', 'TMSB4X', 'TPH1', 'TPM1', 'TPT1', 'TSPN', 'TUBB', 'TXN', 'UBA52', 'UGT2B7', 'UMOD', 'UQCRB', 'UQCRF1', 'VIM', 'WFDC2'

Figure 5: HMHG gene selection in each dataset

Dataset	Genes to be predicted
Colorectum	'A2M', 'ACTA2', 'AEBP1', 'AGR2', 'AHNAK', 'ANXA11', 'ANXA5', 'APL2P', 'APOE', 'ARPC2', 'ASS1', 'ATP1B1', 'ATP5F1A', 'ATP5F1A', 'ATP5F1B', 'ATP5ME', 'ATP6V0B', 'B2M', 'BCAP31', 'BGN', 'BST2', 'BTF3', 'C15orf48', 'C19orf33', 'C1QA', 'C1QB', 'C1R', 'C1S', 'C3', 'CALD1', 'CALM1', 'CALM8', 'CAPZB', 'CCN2', 'CCND1', 'CD24', 'CD44', 'CD55', 'CD59', 'CD74', 'CD81', 'CD99', 'CDH17', 'CEACAM5', 'CEACAM7', 'CHCHD10', 'CIRBP', 'CKB', 'CLCA1', 'CLDN3', 'CLDN4', 'CLDN7', 'CLTA', 'COL12A1', 'COL18A1', 'COL1A1', 'COL1A2', 'COL3A1', 'COL4A1', 'COL4A2', 'COL5A1', 'COL5A2', 'COL6A1', 'COL6A2', 'COL6A3', 'COMM6', 'COX5A', 'COX5B', 'COX6B1', 'COX7B', 'CRIP1', 'CSTB', 'CTSB', 'CTSC', 'CTSD', 'CTSS', 'CXCL1', 'CXCL10', 'CXCL14', 'CYCS', 'FABP1', 'FAM3D', 'FBNL1', 'FCGPB', 'FLNA', 'FN1', 'FOS', 'FSTL1', 'FT1', 'FXYD3', 'GAS5', 'GLUL', 'ENO1', 'EPCAM', 'FABP1', 'FBNL1', 'FCGPB', 'FLNA', 'FN1', 'FOS', 'FSTL1', 'FT1', 'FXYD3', 'GAS5', 'GLUL', 'GNB1', 'GPI', 'GPRC5A', 'GPX2', 'GREM1', 'GSN', 'HLA-A', 'HLA-B', 'HLA-C', 'HLA-DPA1', 'HLA-DPB1', 'HLA-DQB1', 'HLA-DRA', 'HLA-DRA1', 'HMGN1', 'HMGN2', 'HNRPAB', 'HNRPNH1', 'HNRPNU', 'HSBP1', 'HSPPD1', 'HSPPG2', 'ID1', 'IDO1', 'IDR2', 'IER2', 'IEF1', 'IEF13', 'IF16', 'IGFBP4', 'IGFBP5', 'IGFBP7', 'IGHA1', 'IGHG1', 'IGHG3', 'IGHG4', 'IGKC', 'IGLC1', 'IGLC2', 'IGLC3', 'IL32', 'IRF1', 'ISG15', 'ITLN1', 'ITMB2', 'ITM2C', 'JCHAIN', 'JTB', 'JUNB', 'JUND', 'JUNP', 'KDEL2R', 'KLF5', 'KTR18', 'KTR7', 'KRTCP2', 'LCN2', 'LDHB', 'LGALS1', 'LGALS3', 'LGALS4', 'LUM', 'LY6E', 'LYZ', 'MALAT1', 'MARCKSL1', 'MDK', 'MGP', 'MGST1', 'MLEC', 'MMP1', 'MMP11', 'MMP2', 'MORF4L2', 'MUC1', 'MUC12', 'MUC13', 'MUC2', 'MUC5B', 'MYL12B', 'MYL9', 'NAMPT', 'NBL1', 'NDUFA13', 'NDUFB1', 'NDUFS6', 'NOP53', 'NORAD', 'NR4A1', 'OLFM4', 'PABPC1', 'PDI4', 'PERP', 'PFKL', 'PGK1', 'PHGR1', 'PI3', 'PIGR', 'PLA2G2A', 'PLOLR21', 'POSTN', 'PP1B', 'PRDX2', 'PRDX6', 'PRR13', 'PRSS8', 'PSMA1', 'PSMA4', 'PSMB8', 'PTPB1', 'PTGES3', 'PTMA', 'PTMS', 'QSOX1', 'RAB1A', 'RBMS', 'REG1A', 'REG3A', 'REG4', 'RHOA', 'RHOB', 'RHOC', 'RNASE1', 'RNASEK', 'ROM01', 'RRBP1', 'S100A14', 'S100A4', 'S100A6', 'S100A8', 'S100A9', 'S100P', 'SAT1', 'SCD4', 'SEC61A1', 'SELENOP', 'SELENOW', 'SERBP1', 'SERPINA1', 'SERPING1', 'SFN', 'SKP1', 'SLC12A2', 'SMIM22', 'SNRPB', 'SOD2', 'SPARC', 'SPINK1', 'SPINK4', 'SPINT2', 'SPTBN1', 'SRP14', 'SRRM2', 'SRSF3', 'SSR3', 'ST14', 'STAT1', 'SULF1', 'TAGLN', 'TFF1', 'TFF3', 'TGFBI', 'THBS1', 'THY1', 'TIMP1', 'TIMP2', 'TM9SF2', 'TMED10', 'TME176B', 'TME54M', 'TME59M', 'TMSB4X', 'TPM2', 'TP11', 'TSPAN1', 'TSPAN3', 'TSPAN8', 'TSPO', 'TST', 'TXNDC5', 'TYMP', 'UBD', 'UQCR10', 'UQCR11', 'UQCRH', 'VIM', 'XPB1', 'YWHAB', 'ZFA51', 'ZFP361L', 'ZG16'
Skin	'ACTB', 'ACTG1', 'ACTN4', 'AHNAK', 'AHNAK2', 'ANXA1', 'ANXA2', 'APOE', 'APRT', 'AQP3', 'ARPC2', 'ASPRV1', 'ATP1B3', 'ATP5F1A', 'ATP5F1B', 'ATP5MC2', 'ATP5MC3', 'ATP5ME', 'ATP5MF', 'ATP5MG', 'ATP5PD', 'B2M', 'BTF3', 'BTG1', 'C19orf33', 'C4orf3', 'CALM1', 'CALM2', 'CALM3', 'CALM5', 'CASP14', 'CLC7', 'CD24', 'CD44', 'CD63', 'CD74', 'CD81', 'CD9', 'CDSN', 'CFL1', 'CHCHD2', 'CLTB', 'CNBP', 'CNFN', 'COL17A1', 'COL1A1', 'COL1A2', 'COL3A1', 'COL6A1', 'COL6A2', 'COX41', 'COX5B', 'COX6A1', 'COX6B1', 'COX6C', 'COX7A2', 'COX7B', 'COX7C', 'COX8A', 'CRABP2', 'CSDE1', 'CSNK1A1', 'CST3', 'CST6', 'CSTA', 'CSTB', 'CTNNBP1', 'CTSB', 'CTSD', 'CXCL14', 'D1B', 'DCD', 'DCN', 'DDX5', 'DEFB4A', 'DEG51', 'DMKN', 'DSC1', 'DSC3', 'DSCG1', 'DSCG3', 'DSP', 'DSTN', 'DYNLL1', 'EEF1A1', 'EEF1B2', 'EEF1D', 'EEF2', 'EIF1', 'EIF3E', 'EIF3F', 'EIF3K', 'EIF4G2', 'EIF5A', 'ELOB', 'EMP2', 'ENO1', 'EZR', 'FABP5', 'FADS2', 'FAM25A', 'FAU', 'FLG', 'FLG2', 'FT1', 'FTL', 'GAPDH', 'GJA1', 'GJB2', 'GLTP', 'GP9MB', 'GPX4', 'GRN', 'GSN', 'GSTP1', 'GUK1', 'H3-3A', 'H3-3B', 'HINT1', 'HLA-A', 'HLA-B', 'HLA-C', 'HLA-DPA1', 'HLA-DPB1', 'HLA-DQB1', 'HLA-DRA', 'HLA-DRB1', 'HLA-E', 'HMGBL1', 'HMGN2', 'HNRPNA1', 'HNRPNA2B1', 'HNRPNU', 'HOPX', 'HSP90AA1', 'HSP90AB1', 'HSP90B1', 'HSP8A', 'HSBP1', 'IF12', 'IF16', 'IFITM3', 'IGFBP7', 'IGKC', 'ITM2B', 'IVL', 'KLF4', 'KLF5', 'KLT7', 'KRT1', 'KRT10', 'KRT14', 'KRT16', 'KRT17', 'KRT2', 'KRT5', 'KRT6A', 'KRT6B', 'KRT6C', 'KRTDAP', 'LAD1', 'LAPTM4A', 'LCE1B', 'LCE3D', 'LDHA', 'LGALS1', 'LGALS3', 'LGALS7', 'LMNA', 'LY6D', 'LYPD3', 'LYZ', 'MIF', 'MUC1', 'MYL12A', 'MYL6', 'MZYB2', 'NACA', 'NAP1L1', 'NCCR1', 'NDUFA1', 'NDUFA4', 'NDUFB1', 'NDUFB4', 'NDUFC1', 'NDUFS5', 'NOP53', 'NPM1', 'OAZ1', 'OST4', 'P4HB', 'PABPC1', 'PCBP1', 'PCBP2', 'PERP', 'PFDN5', 'PNF1', 'PGAM1', 'PGK1', 'P13', 'PKM', 'PKP1', 'PLP2', 'PLOLR21', 'PPDPF', 'PPIA', 'PP1B', 'PPL', 'PPP1R14B', 'PRDX1', 'PSAP1', 'PSMA7', 'PTMA', 'RAB11A', 'RAC1', 'RACK1', 'RAN', 'RBM3', 'RHOA', 'ROM01', 'RTN4', 'S100A10', 'S100A11', 'S100A14', 'S100A16', 'S100A2', 'S100A4', 'S100A6', 'S100A7', 'S100A8', 'S100A9', 'SSBN', 'SCGB2A2', 'SCDC1', 'SELENOW', 'SEM1', 'SERBP1', 'SERF2', 'SERPINB3', 'SERPINB4', 'SERPINB5', 'SFN', 'SH3BGLR3', 'SKP1', 'SLC25A3', 'SLC25A5', 'SLC25A6', 'SLC25A1', 'SLC38A2', 'SLPI', 'SLURP1', 'SNRPD2', 'SPARC', 'SPINK5', 'SPIN2', 'SPRR1B', 'SPRR2A', 'SPRR2B', 'SPRR2D', 'SPRR2E', 'SPRR2G', 'SRP14', 'SSR4', 'SUB1', 'TACSTD2', 'TAGLN2', 'TMA7', 'TMBIM6', 'TME45A', 'TMSB10', 'TMSB4X', 'TOMM7', 'TP11', 'TP11', 'TRIM29', 'TSP0', 'TUBA1B', 'TUBA1C', 'TUBB', 'TUBB4B', 'TXN', 'TXNIP', 'TYMP', 'UBA52', 'UBB', 'UBC', 'UBEB2D3', 'UBL5', 'UOCR10', 'UOCR11', 'UQCRRQ', 'VIM', 'YBX1', 'YBX3', 'YWHAB', 'YWHAZ', 'ZFP361L', 'ZFP362L'
Kidney	'A2M', 'ACADVL', 'ACAT1', 'ACO2', 'ACTA2', 'ACTB', 'ACTG1', 'ADGRG1', 'AD11', 'ADIRF', 'AEBP1', 'ALDOB', 'ANAPC16', 'ANPEP', 'ANXA2', 'ANXA5', 'APL2P', 'APOE', 'APP', 'AQP1', 'AQP2', 'ARHGDIA', 'ASAH1', 'ASS1', 'ATP1A1', 'ATP1B1', 'ATP5F1A', 'ATP5F1D', 'ATP5MC3', 'ATP5ME', 'ATP5MF', 'ATP6AP2', 'ATP6V0B', 'B2M', 'BCAM', 'BGN', 'BSG', 'C1QA', 'C1R', 'C7', 'CA2', 'CALB1', 'CALD1', 'CALM1', 'CALM2', 'CAN', 'CAPN2', 'CD151', 'CD24', 'CD74', 'CD81', 'CD9', 'CDH16', 'CDKN1C', 'CFL1', 'CHCHD10', 'CHCHD2', 'CIRBP', 'CKB', 'CLCNKB', 'CLTC', 'CLU', 'COL1A2', 'COL3A1', 'COL4A1', 'COL4A2', 'COX5A', 'COX5B', 'COX6A1', 'COX6B1', 'COX6C', 'COX7A2', 'COX7B', 'COX7C', 'COX8A', 'CRIM1', 'CRIP2', 'CRYAB', 'CSDE1', 'CSRP1', 'CST3', 'CTSB', 'CTSH', 'CXCL12', 'CXCL14', 'CYSTM1', 'DNC', 'DDT', 'DDX17', 'DDX5', 'DEFB1', 'DSTN', 'DUSP1', 'DYNLL1', 'DYNLL2', 'EEF1A1', 'EEF1D', 'EEF2', 'EEFH1', 'EIF3K', 'EIF4A1', 'EIF4A2', 'EIF4B', 'ENG', 'EPAS1', 'EZR', 'FABP1', 'FAU', 'FCGRT', 'FLNA', 'FT1', 'FTL', 'FXYD2', 'FXYD4', 'GABARAP', 'GATM', 'GHITM', 'GPX3', 'GSN', 'GSTM3', 'GSTP1', 'GTF21', 'HINT1', 'HLA-A', 'HLA-B', 'HLA-C', 'HLA-DPA1', 'HLA-DRB1', 'HLA-E', 'HNRPNA1', 'HNRPNA2B1', 'HNRPNH1', 'HSID1B2', 'HSP90AB1', 'HSP8A', 'HSBP1', 'HSPPD1', 'HTRA1', 'IDH2', 'IFITM2', 'IFITM3', 'IGFBP2', 'IGFBP4', 'IGFBP5', 'IGFBP7', 'IGHA1', 'IGHG1', 'IGHG3', 'IGHG4', 'IGKC', 'IGLC1', 'IGLC2', 'IGLC3', 'ITGB1', 'ITGB2', 'IVNS1ABP', 'JUND', 'KCNJ1', 'KCNJ15', 'KNG1', 'LAMB2', 'LAMP1', 'LAMTOR5', 'LAPTM4A', 'LDHA', 'LGALS1', 'LITAF', 'LRP2', 'LYM6', 'MALAT1', 'METTL7A', 'MGP', 'MGST1', 'MGST3', 'MIF', 'MIOX', 'MMP7', 'MUC1', 'MYL12A', 'MYL6', 'MYL9', 'M2T2B', 'NAT8', 'NDRG1', 'NDUFA1', 'NDUFA13', 'NDUFA2', 'NDUFA4', 'NDUFB1', 'NDUFB2', 'NDUFB7', 'NDUFB9', 'NDUFC1', 'NDUFS5', 'NDUFS6', 'NDUFS7', 'NEAT1', 'NME2', 'NPM53', 'NORAD', 'NPC2', 'NUCB1', 'OAZ1', 'OGDH', 'OST4', 'P4HB', 'PCBP1', 'PCPK1', 'PDZK1IP1', 'PEBP1', 'PEPD', 'PFKL', 'Pfn1', 'PGAM1', 'PGK1', 'PHPT1', 'PPIB', 'PODXL', 'PLOLR21', 'PP1B', 'PPP1R1A', 'PTGDS', 'PTHR1', 'RABAC1', 'RAC1', 'REN', 'RHCGB', 'RHOA', 'RNASE1', 'ROM01', 'RTN4', 'S100A10', 'S100A2', 'S100A6', 'SAT1', 'SCNN1A', 'SCPP', 'SCD1', 'SELENOM', 'SELENOP', 'SERPINA1', 'SERPINB5A', 'SERPINB5', 'SH3BGLR3', 'SLC12A1', 'SLC12A3', 'SLC13A3', 'SLC25A5', 'SLC25A6', 'SLC3A1', 'SLC5A12', 'SMIM24', 'SNHG25', 'SOD1', 'SOD2', 'SPARC', 'SPINK1', 'SP1', 'SPR14', 'SSR4', 'ST13', 'SUCLG1', 'SUMO2', 'TAGLN1', 'TAGLN2', 'TAPBP', 'THY1', 'TIMP1', 'TIMP2', 'TIMP3', 'TINAGL1', 'TMA7', 'TMSB10', 'TMSB4X', 'TP11', 'TPM1', 'TPM3', 'TP11', 'TRIR', 'TSC22D1', 'TSPAN1', 'TUBA1A', 'TUBB', 'TXN', 'UBA52', 'UGT2B7', 'UMOD', 'UQCRR', 'UQCRR1', 'UQCRRFS1', 'VIM', 'WFDC2', 'ZFP362L'

Figure 6: HVG gene selection in each dataset

918 B.3 IMPLEMENTATION DETAILS FOR ABLATION STUDIES
919920 In this work, HiST consists of three main modules: Hierarchical Hyperbolic Alignment (HHA), Gene
921 Decoder and Multi-Level Representation Extractors. To validate the necessity of each component,
922 we design a series of ablation experiments for each module respectively.
923924 **Strategy of Alignment** We design five strategies of alignment to evaluate the necessity of the
925 HHA module. First, we remove only the gene-image regularization term of the HEA loss (w/o G-I
926 HEA) to investigate the impact of lacking the entailment loss between gene and image. Second, we
927 remove the Hierarchical Entailment Alignment component of HHA (w/o HEA) to investigate the
928 impact of lacking hierarchical constraints in hyperbolic space on model performance. Third, we
929 eliminate HHA module (w/o HEA + HCA), retaining only multi-scale image information to assess
930 its contribution. Fourth, we replace this module with MERU (Desai et al., 2023) (MERU), preserving
931 only the cross-model hierarchical structure while disregarding intra-modal multi-scale information
932 in ST data. Finally, we substitute this module with CLIP (Radford et al., 2021) (CLIP), aligning
933 only spot-level gene expression with spot-level images. These experiments collectively highlight
934 the critical roles of hierarchical constraints, multi-scale information integration, and cross-modal
935 alignment mechanisms in the module’s effectiveness.
936937 **Input of Decoder** In order to investigate the impact of input data on the overall performance of
938 the Gene Decoder, we conducted ablation experiments where we separately use the representations
939 learned from spot-level images and niche-level images as inputs to the Gene Decoder to predict genes.
940941 **Choice of LoRA** To investigate the significance of the LoRA (Low-Rank Adaptation) component
942 within the Multi-Level Representation Extractors, we design experiments to evaluate its role in
943 efficient finetuning the pre-trained model to current task for multi-scale inputs, particularly niche-
944 level images, which may differ in resolution from standard inputs. By adjusting the number of last
945 attention layers adapted by LoRA, we can control the extent of fine-tuning. Setting the number of
946 adapted attention layers to 0 (i.e., freezing all pre-trained weights) allowed us to establish a baseline
947 where no fine-tuning occurs.
948

949 B.4 IMPLEMENTATION DETAILS FOR EVALUATION METROLOGY

950 We evaluate model performance using the top-k mean Pearson Correlation Coefficient (PCC@k),
951 mean squared error (MSE), and mean absolute error (MAE). For the j -th gene at the i -th spot, the
952 PCC of the j -th gene (PCC_j) is formulated as:
953

954
$$PCC_j = \frac{\sum_{i=1}^n (\hat{y}_{i,j} - \bar{\hat{y}}_{\cdot,j})(y_{i,j} - \bar{y}_{\cdot,j})}{\sqrt{\sum_{i=1}^n (\hat{y}_{i,j} - \bar{\hat{y}}_{\cdot,j})^2} \sqrt{\sum_i^n (y_{i,j} - \bar{y}_{\cdot,j})^2}}, \quad (15)$$

955

956 where $\hat{y}_{i,j}$ and $y_{i,j}$ represent the predicted and actual gene expression of the j -th gene at the i -th spot,
957 respectively, and $\bar{\hat{y}}_{\cdot,j}$ and $\bar{y}_{\cdot,j}$ denote the mean predicted and actual gene expression of the j -th gene
958 across spots. m and n are the numbers of genes and spots, separately. For PCC@k, the average value
959 across top-k PCC_j is calculated as:
960

961
$$PCC@k = \frac{1}{k} \sum_{j \in Topk} PCC_j. \quad (16)$$

962

963 Subsequently, MSE and MAE can be defined as:
964

965
$$MAE = \frac{1}{n \times m} \sum_{i=1}^n \sum_{j=1}^m |y_{i,j} - \hat{y}_{i,j}|, \quad (17)$$

966

967
$$MSE = \frac{1}{n \times m} \sum_{i=1}^n \sum_{j=1}^m (y_{i,j} - \hat{y}_{i,j})^2. \quad (18)$$

968

972 C ADDITIONAL EXPERIMENT RESULTS
973974 C.1 STATISTICAL SIGNIFICANCE ANALYSIS
975

976 To assess the statistical significance of the observed performance improvements, we conducted
977 paired t-tests across 5 independent runs. As shown in Table 4, HiST shows statistically significant
978 improvements ($p < 0.05$) on the vast majority of metrics when compared to baselines. For the
979 comparison against TRIPLEX, although a few metrics did not reach the significance threshold, HiST
980 still demonstrated a consistent performance trend rather than an incidental improvement. We attribute
981 these few instances of non-significance to the high biological variance inherent in the datasets, yet
982 the overall results strongly support the robustness of our method.

983 Table 4: HiST performance improvement (P-value) compared to baseline models on three datasets
984 for HMHVGs.
985

986 Dataset	987 Model	PCC@10 ↑	PCC@50 ↑	PCC@200 ↑	MSE ↓	MAE ↓
988 Colorectum	989 TRIPLEX	2.83% (0.119)	2.94% (0.135)	3.24% (0.083)	19.84% (0.048)	9.26% (0.034)
	990 StNet	11.62% (0.012)	12.62% (0.010)	13.8% (0.019)	11.12% (<0.01)	6.31% (<0.01)
	991 BLEEP	13.15% (<0.01)	15.47% (<0.01)	24.7% (<0.01)	26.48% (<0.01)	12.57% (<0.01)
	992 Stem	7.62% (0.030)	12.10% (<0.01)	19.52% (<0.01)	16.19% (0.015)	7.13% (0.031)
993 Skin	994 TRIPLEX	0.91% (0.057)	1.59% (0.018)	2.52% (<0.01)	5.08% (0.066)	5.32% (<0.01)
	995 StNet	4.34% (<0.01)	4.24% (<0.01)	4.52% (<0.01)	6.22% (0.081)	4.73% (<0.01)
	996 BLEEP	6.48% (<0.01)	6.63% (<0.01)	7.76% (<0.01)	16.59% (<0.01)	6.29% (<0.01)
	997 Stem	7.26% (<0.01)	8.43% (<0.01)	10.35% (<0.01)	27.00% (<0.01)	10.07% (<0.01)
998 Kidney	999 TRIPLEX	6.63% (<0.001)	8.30% (<0.01)	10.95% (<0.01)	3.94% (0.051)	4.41% (0.043)
	1000 StNet	17.94% (<0.01)	20.77% (<0.01)	27.85% (<0.01)	7.71% (<0.01)	3.58% (<0.01)
	1001 BLEEP	19.20% (<0.01)	21.16% (<0.01)	25.68% (<0.01)	12.59% (<0.01)	5.54% (<0.01)
	1002 Stem	15.43% (<0.01)	26.85% (<0.01)	43.98% (<0.01)	21.90% (<0.01)	10.31% (<0.001)

1000 C.2 EFFICIENCY BENCHMARK
1001

1002 We evaluated the computational overhead of our proposed model by comparing the memory usage and
1003 runtime performance of HiST against other Euclidean models. While HiST’s hyperbolic computations
1004 incur higher overhead without specific hardware optimization, we find this acceptable given its
1005 superior performance advantages (as shown in Table 1 in the main text), its generalization power, and
1006 its future scalability potential for biological modeling. Table 5 details the training memory, training
1007 time per epoch, and inference time per epoch on the HCC dataset for 200 target genes.

1008 Table 5: Memory usage and runtime performance comparison of HiST and baseline models on the
1009 HCC dataset with 200 target genes.
1010

1011 Model	1012 Memory Usage (GB) ↓	1013 Training Time (s/epoch) ↓	1014 Inference Time (s/epoch) ↓
1013 TRIPLEX	13.34	30	2
1014 StNet	18.44	27	1
1015 BLEEP	5.89	10	1
1016 Stem	12.96	22	1553
1017 HiST (Ours)	18.39	74	3

1018 C.3 CROSS-LABORATORY GENERALIZATION
1019

1020 We conducted a cross-laboratory validation experiment to assess the model’s generalization capability
1021 across different laboratory settings and potential stain variations. We used a Whole Slide Image
1022 (WSI) (NCB1563) from a new, independent kidney dataset (Canela et al., 2023). As shown in Table
1023 6, our model, HiST (Ours), achieved significantly superior performance compared to the baseline
1024 models. This result highlights HiST’s robustness and its ability to capture transferable biological
1025 features that generalize beyond the training domain.

1026
1027
1028 Table 6: Performance comparison on a cross-laboratory whole slide image.
1029
1030
1031
1032
1033

Model	PCC@10 ↑	PCC@50 ↑	PCC@200 ↑	MSE ↓	MAE ↓
TRIPLEX	0.374±0.056	0.265±0.035	0.162±0.030	1.595±0.121	1.015±0.038
StNet	0.215±0.039	0.111±0.033	-0.003±0.031	1.758±0.183	1.045±0.046
BLEEP	0.281±0.042	0.199±0.053	0.096±0.056	1.646±0.215	1.019±0.057
Stem	0.357±0.015	0.254±0.008	0.145±0.007	1.822±0.143	1.033±0.042
HiST	0.510±0.016	0.361±0.022	0.197±0.024	1.574±0.128	0.983±0.039

1034
1035
1036 C.4 ROBUSTNESS ON HVG GENE SET
1037

1038 We evaluated HiST on the HVG gene set, as described in Section B.1, to further demonstrate the
1039 robustness of our model. As presented in Table 7, HiST outperforms other methods, achieving
1040 PCC@200 values of 0.505, 0.821, and 0.450 across all three datasets (Colorectum (Valdeolivas et al.,
1041 2024), Skin (Schäbitz et al., 2022), Kidney (Lake et al., 2023)), respectively. Similar to the results for
1042 HMHVG in Table 1, TRIPLEX also achieves the second-best overall performance, demonstrating
1043 that the integration of multi-scale image features consistently enhances model performance across
1044 different gene selection criteria.

1045
1046 Table 7: Performance comparison on the HVG gene set of three spatial transcriptomics datasets.
1047 Higher values on PCC@10, PCC@50, PCC@200 are better. Lower values on MAE and MSE are
1048 better.

Dataset	Model	PCC@10 ↑	PCC@50 ↑	PCC@200 ↑	MSE ↓	MAE ↓
Colorectum	TRIPLEX	0.685±0.154	0.613±0.181	0.484±0.238	1.830±0.826	1.042±0.257
	StNet	0.656±0.122	0.578±0.138	0.447±0.181	1.642±0.427	1.009±0.154
	BLEEP	0.662±0.119	0.568±0.122	0.422±0.170	1.891±0.673	1.064±0.211
	Stem	0.679±0.111	0.574±0.124	0.415±0.173	1.799±0.613	1.045±0.205
	HiST (Ours)	0.716±0.117	0.641±0.137	0.504±0.195	1.464±0.564	0.951±0.197
Skin	TRIPLEX	0.826±0.093	0.795±0.112	0.745±0.136	0.957±0.488	0.675±0.219
	StNet	0.808±0.100	0.782±0.113	0.739±0.135	0.970±0.402	0.677±0.192
	BLEEP	0.782±0.108	0.755±0.120	0.708±0.139	1.117±0.559	0.696±0.224
	Stem	0.781±0.097	0.749±0.116	0.698±0.140	1.196±0.626	0.707±0.241
	HiST (Ours)	0.838±0.089	0.812±0.104	0.766±0.127	0.882±0.387	0.631±0.181
Kidney	TRIPLEX	0.541±0.094	0.453±0.087	0.331±0.068	1.138±0.255	0.843±0.098
	StNet	0.537±0.111	0.454±0.100	0.332±0.079	1.084±0.153	0.822±0.052
	BLEEP	0.508±0.123	0.426±0.114	0.306±0.087	1.181±0.219	0.853±0.080
	Stem	0.501±0.112	0.404±0.097	0.273±0.067	1.288±0.199	0.888±0.071
	HiST (Ours)	0.618±0.099	0.530±0.090	0.399±0.071	1.079±0.198	0.821±0.082

1066
1067
1068 C.5 EVALUATION WITH PATIENT-LEVEL DATA SPLITTING
1069

1070 We implemented a strict patient-level data splitting protocol to address potential data leakage and
1071 ensure a more rigorous evaluation of generalization performance. This ensures that all Whole Slide
1072 Images (WSIs) from a single patient are confined to a single data partition (train, validation, or test),
1073 adopting a stricter and more clinically meaningful evaluation standard than prior works (Zhu et al.,
1074 2025; Xie et al., 2023).

1075 As shown in Table 8, this more challenging setup led to a performance drop across all methods, which
1076 confirms the presence of patient-specific features and highlights the importance of this strict separation.
1077 Crucially, even under these stringent conditions, HiST maintains its position as the top-performing
1078 model. It consistently achieves state-of-the-art results across most Pearson Correlation (PCC) metrics
1079 and datasets, demonstrating that its architectural advantages provide greater generalization capability
to new, unseen patients, making it more robust for real-world applications.

1080 Table 8: Performance comparison on three spatial transcriptomics datasets using strict patient-level
 1081 splits. **Bold** indicates the best performance, and underline indicate the second-best performance.
 1082

1083	Dataset	Model	PCC@10 \uparrow	PCC@50 \uparrow	PCC@200 \uparrow	MSE \downarrow	MAE \downarrow
1084	Colorectum	TRIPLEX	<u>0.685\pm0.115</u>	<u>0.613\pm0.128</u>	<u>0.418\pm0.116</u>	2.423 \pm 0.744	1.187 \pm 0.187
		StNet	0.654 \pm 0.111	0.585 \pm 0.114	0.386 \pm 0.091	<u>2.284\pm0.319</u>	<u>1.155\pm0.098</u>
		BLEEP	0.612 \pm 0.119	0.538 \pm 0.117	0.340 \pm 0.081	2.551 \pm 0.513	1.213 \pm 0.137
		Stem	0.653 \pm 0.136	0.560 \pm 0.126	0.357 \pm 0.086	2.351 \pm 0.504	1.166 \pm 0.144
		HiST	0.712\pm0.096	0.642\pm0.093	0.448\pm0.078	2.152\pm0.515	1.131\pm0.152
1085	Skin	TRIPLEX	0.816 \pm 0.089	0.783 \pm 0.108	<u>0.726\pm0.131</u>	<u>1.361\pm0.701</u>	0.812 \pm 0.196
		StNet	0.786 \pm 0.102	0.759 \pm 0.112	0.704 \pm 0.128	1.407 \pm 0.652	0.832 \pm 0.194
		BLEEP	0.789 \pm 0.102	0.761 \pm 0.116	0.705 \pm 0.132	1.479 \pm 0.627	0.826 \pm 0.204
		Stem	0.765 \pm 0.103	0.734 \pm 0.117	0.673 \pm 0.132	1.743 \pm 1.040	0.898 \pm 0.264
		HiST	0.828\pm0.079	0.799\pm0.093	0.746\pm0.111	1.310\pm0.727	0.792\pm0.210
1086	Kidney	TRIPLEX	<u>0.511\pm0.107</u>	0.418 \pm 0.109	0.288 \pm 0.102	1.207\pm0.249	0.868\pm0.100
		StNet	0.503 \pm 0.079	<u>0.425\pm0.079</u>	<u>0.303\pm0.077</u>	1.332 \pm 0.311	0.917 \pm 0.166
		BLEEP	0.495 \pm 0.083	0.407 \pm 0.085	0.279 \pm 0.086	1.368 \pm 0.365	0.915 \pm 0.137
		Stem	0.468 \pm 0.073	0.365 \pm 0.063	0.227 \pm 0.059	1.449 \pm 0.323	0.942 \pm 0.115
		HiST	0.587\pm0.065	0.495\pm0.069	0.357\pm0.074	<u>1.283\pm0.362</u>	0.889 \pm 0.140

1100 C.6 SCALABILITY WITH NUMBER OF TARGET GENES

1101
 1102 We investigated the model’s scalability by analyzing how its performance, memory usage, and
 1103 training time scale with an increasing number of target genes. As shown in Table 9, increasing
 1104 the number of target genes from 200 to 3000 has a minimal impact on memory and training time.
 1105 This is because only the first and last linear layers of the gene encoder and decoder are affected by
 1106 the number of genes. While the computational cost of these layers scales linearly, this overhead is
 1107 insignificant compared to the model’s overall complexity, demonstrating HiST’s excellent scalability
 1108 for larger-scale genomic analyses.

1109 Table 9: Analysis of performance, memory, and training time scale with the number of target genes
 1110 in the HiST Model on the HCC dataset.

1112	Metric	Target Genes		
		200	1000	3000
1113	PCC@10 \uparrow	0.721 \pm 0.105	0.723 \pm 0.114	0.713 \pm 0.107
1114	PCC@50 \uparrow	0.642 \pm 0.128	0.655 \pm 0.133	0.649 \pm 0.123
1115	PCC@200 \uparrow	0.477 \pm 0.184	0.557 \pm 0.165	0.564 \pm 0.153
1116	PCC@500 \uparrow	-	0.472 \pm 0.187	0.492 \pm 0.167
1117	PCC@1000 \uparrow	-	0.387 \pm 0.184	0.429 \pm 0.169
1118	PCC@2000 \uparrow	-	-	0.353 \pm 0.156
1119	PCC@3000 \uparrow	-	-	0.143 \pm 0.294
1120	MSE \downarrow	1.498 \pm 0.456	0.933 \pm 0.215	0.632 \pm 0.095
1121	MAE \downarrow	0.958 \pm 0.158	0.766 \pm 0.092	0.621 \pm 0.039
1122	Memory Usage (GB)	18.39	18.43	18.76
1123	Training Time (s/epoch)	74	73	73
1124	Inference Time (s/epoch)	3	3	2

1125 C.7 SENSITIVITY TO NEIGHBORHOOD SIZE

1126
 1127 The construction of a niche is dependent on the neighborhood size, a key hyperparameter in our model.
 1128 To evaluate the model’s sensitivity to this hyperparameter, we varied the neighborhood sizes (6, 18,
 1129 and 36) for niche construction on the HCC dataset. As shown in Table 10, the model’s performance
 1130 remains highly stable across the different settings. This demonstrates that HiST is robust and not
 1131 sensitive to the choice of neighborhood size.

1134 Table 10: Performance comparison on different neighborhood sizes on the HCC dataset.
1135

1136 Neighborhood sizes	1137 PCC@10 ↑	1138 PCC@50 ↑	1139 PCC@200 ↑	1140 MSE ↓	1141 MAE ↓
1142 6	1143 0.721±0.105	1144 0.642±0.128	1145 0.477±0.184	1146 1.498±0.456	1147 0.958±0.158
1148 18	1149 0.714±0.112	1150 0.633±0.132	1151 0.474±0.173	1152 1.569±0.582	1153 0.978±0.189
1154 36	1155 0.716±0.109	1156 0.636±0.126	1157 0.474±0.172	1158 1.616±0.613	1159 0.993±0.201

1142 C.8 QUANTITATIVE GENE-WISE ANALYSIS

1144 Our quantitative, per-gene analysis reveals that the enhanced performance of HiST on spatially-
1145 constrained genes (e.g., immune, structural, and tissue-specific markers) validates our hypothesis that
1146 morphological context provides meaningful constraints for these gene categories. Table 11 presents a
1147 detailed comparison of the gene-wise Pearson Correlation Coefficient (PCC) improvement of HiST
1148 over TRIPLEX in the tumor region of a representative sample, highlighting significant performance
1149 gains on genes with known biological importance.

1150 Table 11: Gene-wise PCC improvement of HiST vs. TRIPLEX in the tumor region of sample ZEN42.
1151 Genes with significant biological importance are highlighted in **bold**.

1153 Gene	1154 PCC (HiST)	1155 PCC (TRIPLEX)	1156 PCC Improvement
1157 HLA-DQB1	0.592	0.237	0.355
1158 S100A14	0.575	0.224	0.351
1159 GREM1	0.625	0.296	0.329
1160 FABP1	0.161	-0.166	0.327
1161 PIGR	0.810	0.511	0.299
1162 MUC5B	0.580	0.291	0.289
1163 REG3A	0.536	0.279	0.257
1164 REG1A	0.476	0.226	0.250
1165 LCN2	0.497	0.258	0.240
1166 FXYD3	0.461	0.230	0.232
1167 HLA-DPB1	0.539	0.315	0.223
1168 TFF3	0.769	0.552	0.217
1169 CD74	0.589	0.379	0.210
1170 SPINK1	0.531	0.346	0.186

1169 C.9 BIOLOGICAL VALIDATION ON KEY MARKER GENES

1171 To demonstrate the translational value and biological insights of our method, we conducted a
1172 region-specific gene prediction analysis targeting the tumor-stromal interface in colorectal cancer,
1173 a critical area for prognosis and therapeutic response (Valdeolivas et al., 2024; Feng et al., 2024).
1174 We selected three clinically validated marker genes representing distinct functional compartments:
1175 **KRT18** (epithelial tumor core), **ACTA2** (stromal activation), and **IGKC** (B-cell/plasma cell immune
1176 response).

1177 As shown in Table 12, we measured the region-specific prediction accuracy improvement over
1178 baseline methods. The results indicate that HiST more accurately captures the spatial expression
1179 patterns of these critical genes, which is essential for downstream applications such as therapeutic
1180 target identification, immune infiltration assessment, and tumor margin evaluation.

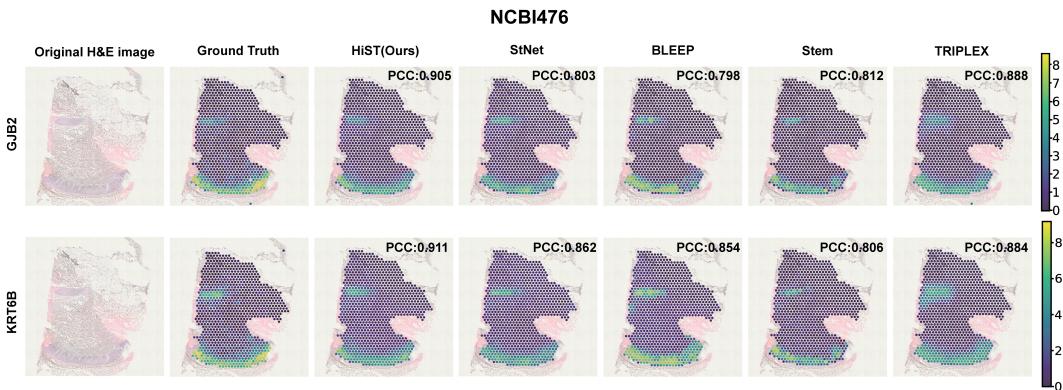
1181 D ADDITIONAL VISUALIZATION

1184 In this section, we present visualizations of additional marker genes across different tissues to
1185 demonstrate the robustness of our findings. We selected three samples (MICS27, NCBI476, and
1186 NCBI704) from datasets of lung, skin, and kidney tissues, respectively, and plotted the predicted gene
1187 expressions on their H&E images for the following marker genes: ALDH1A1, ELF3, GJB2, KRT6B,
1188 UMOD, and PODXL. These genes are highly correlated with specific tissue cell types. Specifically,

1188 Table 12: Region-specific gene prediction accuracy improvement over baseline models. P-values
 1189 were determined by a paired Student’s t-test across tissue regions.
 1190

1191	Gene	Model	MAE Improv. (%)	MSE Improv. (%)	Correlation Improv. (%)
1192	KRT18	BLEEP	8.80 (0.1084)	23.13 (0.0698)	18.92
1193		iSTAR	52.42 (<0.001)	131.37 (<0.001)	12.91
1194		Stem	17.94 (<0.001)	42.71 (<0.001)	7.56
1195		StNet	9.86 (0.0875)	31.95 (0.0171)	14.04
1196		TRIPLEX	84.57 (<0.001)	238.58 (<0.001)	10.40
1197	ACTA2	BLEEP	36.66 (0.0064)	70.55 (0.0582)	5.34
1198		iSTAR	43.89 (0.0016)	108.59 (0.0105)	24.04
1199		StNet	10.67 (0.4406)	16.43 (0.6867)	2.45
1200		TRIPLEX	28.76 (0.0414)	62.14 (0.1025)	1.08
1201		StNet	36.89 (0.0413)	72.50 (0.1153)	7.72
1202	IGKC	BLEEP	53.18 (0.0037)	231.16 (0.0169)	41.84
1203		iSTAR	4.26 (0.8164)	52.25 (0.4332)	1.96
1204		StNet	53.06 (0.0030)	250.32 (0.0072)	31.50
1205		TRIPLEX	36.73 (0.0299)	206.47 (0.0417)	32.97
1206		StNet	143.21 (<0.001)	691.59 (<0.001)	24.43

1207
 1208 ALDH1A1 and ELF3 are marker genes for lung tissue, GJB2 and KRT6B are marker genes for skin
 1209 tissue, and UMOD and PODXL are marker genes for kidney tissue. Visualizations of predicted gene
 1210 expressions by different models, along with their PCC comparisons, are presented in Figures 7 to 10.
 1211 The figures clearly demonstrate that HiST’s gene expression predictions exhibit strong consistency
 1212 with the ground truth, achieving higher PCC than existing models.
 1213



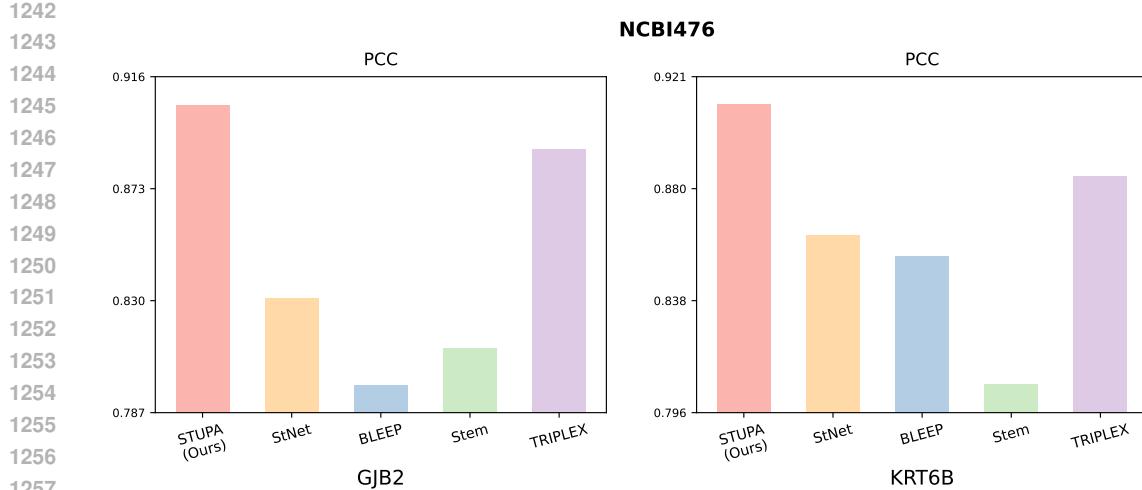


Figure 8: PCC comparison of GJB2 and KRT6B gene predictions (Higher PCC reflects greater accuracy)

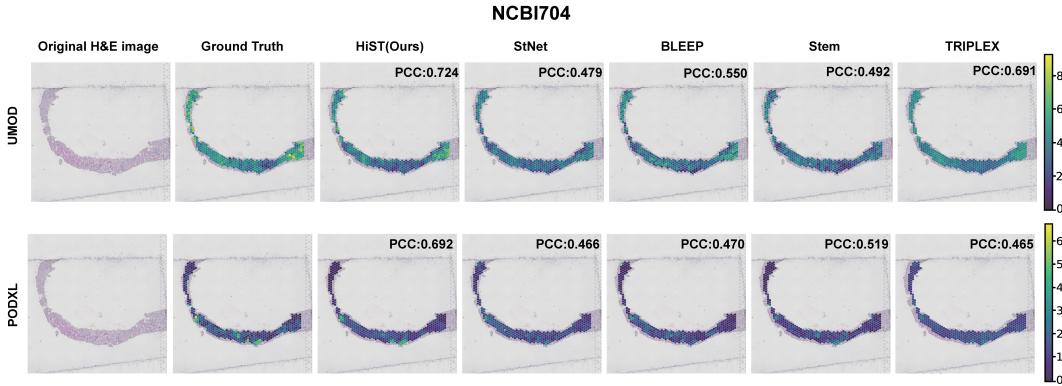


Figure 9: Visualization of UMOD and PODXL gene predictions in the NCBI704 sample from the Kidney dataset. We present the PCC values comparing the ground truth with the gene expression predictions generated by each model.

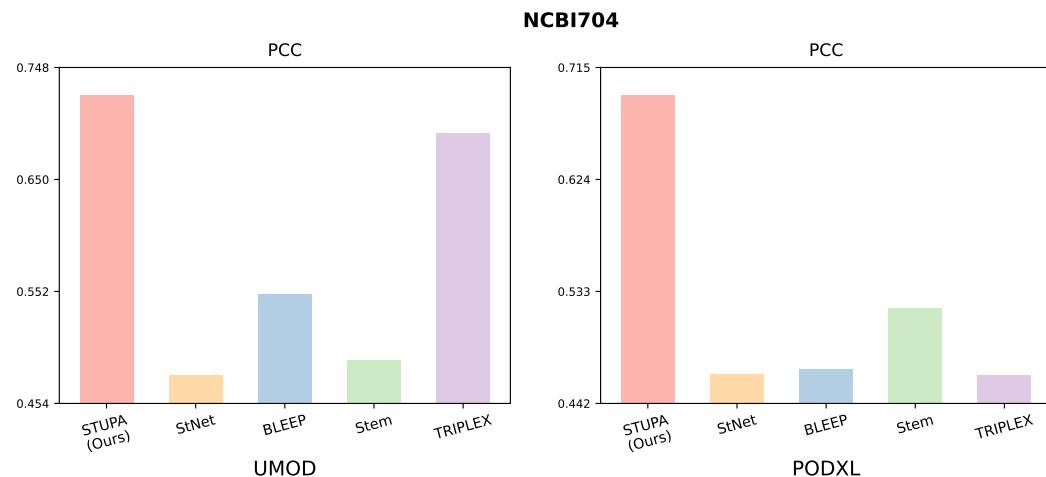


Figure 10: PCC comparison of UMOD and PODXL gene predictions (Higher PCC reflects greater accuracy)

1296 **E ADDITIONAL ABLATION STUDY**
1297

1298 We conduct further ablation studies on the HMHVG gene set for the Skin and Kidney datasets,
1299 with results for alignment strategy, LoRA settings, and decoder input detailed in Tables 13, 14,
1300 and 15, respectively. Through these ablation studies, we rigorously evaluate the contribution of
1301 HiST’s individual modules to system performance. These experiments demonstrate that HiST
1302 consistently achieves optimal performance in the majority of experimental conditions, highlighting
1303 the effectiveness and robustness of these modules.

1304
1305 **Table 13: Additional ablation study of the alignment strategy**
1306

Dataset	Alignment	PCC@10 ↑	PCC@50 ↑	PCC@200 ↑	MSE ↓	MAE ↓
Skin	w/o G-I HEA	0.833±0.084	0.805±0.100	0.753±0.126	0.981±0.431	0.672±0.184
	w/o HEA	0.832±0.090	0.803±0.108	0.746±0.136	0.974±0.418	0.681±0.181
	w/o HEA + HCA	0.809±0.099	0.783±0.113	0.731±0.137	1.006±0.437	0.695±0.187
	MERU	0.812±0.098	0.784±0.115	0.730±0.142	0.986±0.420	0.672±0.181
	CLIP	0.810±0.098	0.784±0.113	0.733±0.137	0.980±0.393	0.681±0.176
	Ours	0.839±0.086	0.812±0.102	0.758±0.129	0.932±0.418	0.657±0.182
Kidney	w/o G-I HEA	0.610±0.101	0.514±0.096	0.378±0.074	1.147±0.188	0.839±0.065
	w/o HEA	0.607±0.092	0.513±0.085	0.377±0.066	1.131±0.195	0.837±0.074
	w/o HEA + HCA	0.576±0.099	0.484±0.089	0.344±0.064	1.134±0.168	0.837±0.058
	MERU	0.586±0.099	0.494±0.090	0.355±0.062	1.148±0.205	0.842±0.071
	CLIP	0.558±0.098	0.462±0.087	0.321±0.058	1.220±0.293	0.867±0.097
	Ours	0.617±0.094	0.526±0.088	0.390±0.07	1.077±0.155	0.817±0.058

1320
1321 **Table 14: Ablation study of the LoRA settings**
1322

Dataset	LoRA Layers	PCC@10 ↑	PCC@50 ↑	PCC@200 ↑	MSE ↓	MAE ↓
Colorectum	0	0.701±0.126	0.622±0.151	0.458±0.200	1.662±0.638	1.007±0.212
	5	0.704±0.131	0.619±0.158	0.456±0.201	1.565±0.612	0.976±0.207
	7	0.713±0.117	0.627±0.146	0.468±0.187	1.562±0.606	0.978±0.209
	Ours (11)	0.721±0.105	0.642±0.128	0.477±0.184	1.498±0.456	0.958±0.158
Skin	0	0.831±0.088	0.804±0.104	0.751±0.130	0.977±0.437	0.677±0.185
	5	0.833±0.088	0.806±0.105	0.754±0.130	0.971±0.438	0.670±0.189
	7	0.838±0.083	0.810±0.100	0.756±0.126	0.944±0.413	0.654±0.184
	Ours (11)	0.839±0.086	0.812±0.102	0.758±0.129	0.932±0.418	0.657±0.182
Kidney	0	0.587±0.099	0.498±0.088	0.366±0.071	1.110±0.162	0.828±0.059
	5	0.603±0.096	0.510±0.088	0.376±0.071	1.100±0.207	0.826±0.083
	7	0.606±0.105	0.511±0.100	0.378±0.084	1.106±0.193	0.827±0.077
	Ours (11)	0.617±0.094	0.526±0.088	0.390±0.070	1.077±0.155	0.817±0.058

1336
1337 **Table 15: Additional ablation study of the decoder input**
1338

Dataset	Decoder Input	PCC@10 ↑	PCC@50 ↑	PCC@200 ↑	MSE ↓	MAE ↓
Colorectum	only spot	0.709±0.098	0.622±0.115	0.450±0.169	1.564±0.334	0.973±0.119
	only niche	0.661±0.179	0.570±0.203	0.405±0.238	1.801±0.878	1.031±0.268
	Ours (spot+niche)	0.721±0.105	0.642±0.128	0.477±0.184	1.498±0.456	0.958±0.158
Skin	only spot	0.814±0.096	0.787±0.111	0.735±0.134	1.003±0.417	0.673±0.187
	only niche	0.831±0.092	0.803±0.107	0.748±0.132	0.979±0.441	0.670±0.188
	Ours (spot+niche)	0.839±0.086	0.812±0.102	0.758±0.129	0.932±0.418	0.657±0.182
Kidney	only spot	0.584±0.093	0.492±0.088	0.353±0.067	1.171±0.184	0.849±0.065
	only niche	0.588±0.097	0.491±0.090	0.356±0.073	1.110±0.139	0.828±0.051
	Ours (spot+niche)	0.617±0.094	0.526±0.088	0.390±0.070	1.077±0.155	0.817±0.058

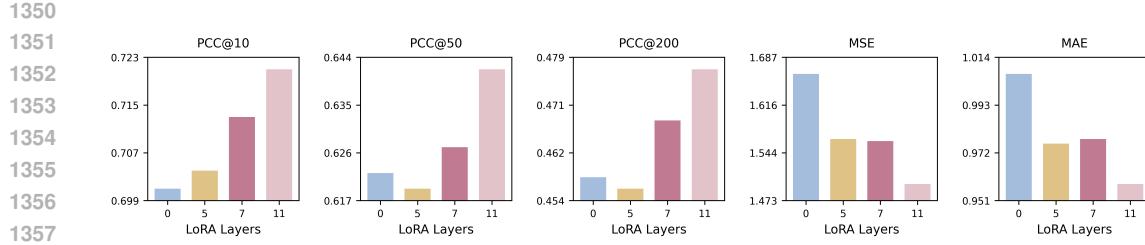


Figure 11: Ablation study on the choice of the last layers of LoRA for Colorectum

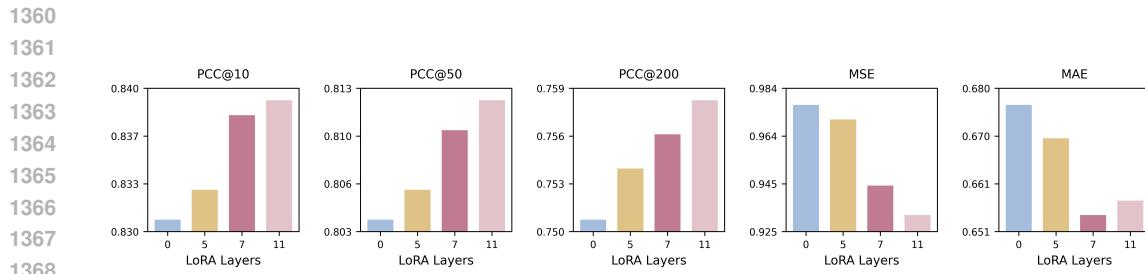


Figure 12: Ablation study on the choice of the last layers of LoRA for Skin

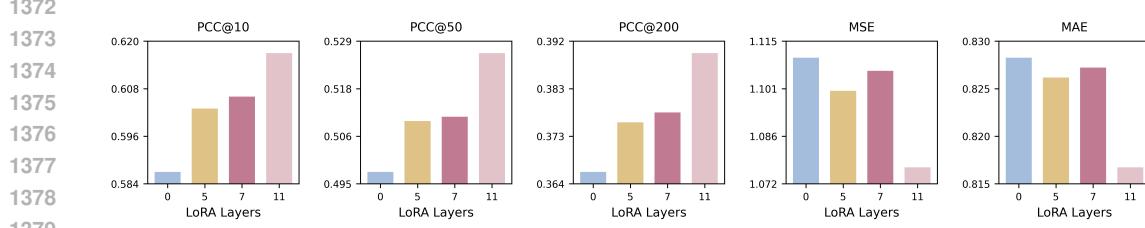


Figure 13: Ablation study on the choice of the last layers of LoRA for Kidney

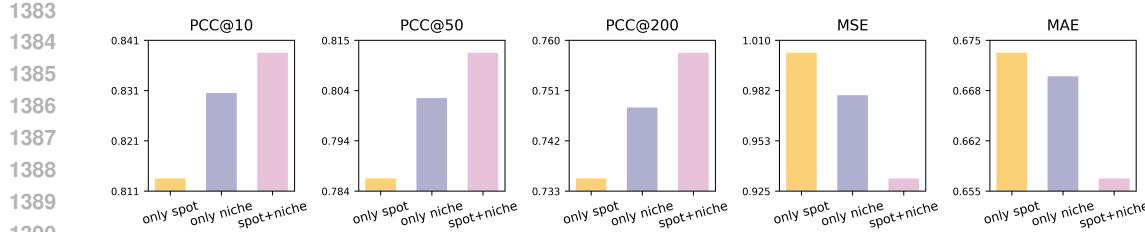


Figure 14: Ablation study of the input data of decoder for Skin

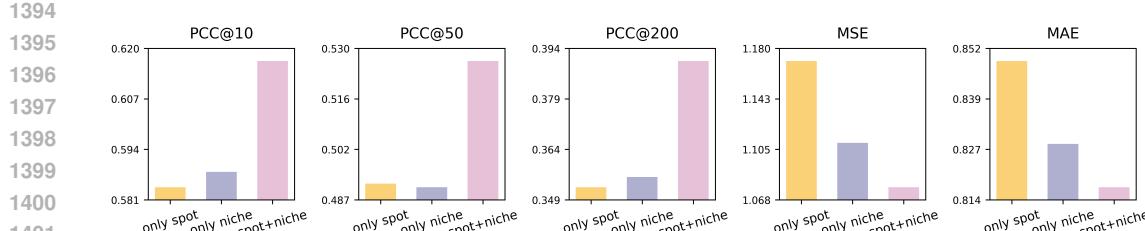


Figure 15: Ablation study of the input data of decoder for Kidney

1404 **F LIMITATION**
 1405

1406 While HiST advances spatial transcriptomics prediction through the hierarchical hyperbolic structure,
 1407 several limitations merit discussion. First, the gene decoder relies on a simple multilayer perceptron
 1408 (MLP) to predict expression from aligned image representations, potentially limiting its capacity to
 1409 model complex gene-gene interactions. Second, the integration of multi-level features (spot-level
 1410 and niche-level images) is achieved through concatenation, which may overlook fine-grained feature
 1411 fusion. Third, the model is trained and fine-tuned on a limited dataset, and its generalization capability
 1412 could be further enhanced by training on larger, more diverse cohorts. In the future, we plan to
 1413 address these limitations by training on a large-scale dataset and transforming HiST into a foundation
 1414 model for broader applicability.

1415 **G THE USE OF LARGE LANGUAGE MODELS (LLMs)**
 1416

1417 In the writing of this paper, we used large language models only for aiding and polishing the text.
 1418 We did not use them for research ideation or to make a substantive contribution to the content of the
 1419 paper.

1420 **H BROADER IMPACT**
 1421

1422 **H.1 IMPACT ON REAL-WORLD APPLICATIONS**
 1423

1424 For decades, WSIs have served as a foundation in biomedical research and clinical diagnostics.
 1425 However, the expensive and labor-intensive nature of Spatial Transcriptomics (ST) limits its broad
 1426 use, driving the demand for deep learning methods that directly predict spatially resolved gene
 1427 expression from these images. To tackle the challenges of incorporating broader pathological
 1428 and genetic contexts while effectively capturing target-modality information, we present HiST, an
 1429 innovative framework for predicting spatial transcriptomics. HiST leverages multi-level hyperbolic
 1430 image-gene representations, which model hierarchical structures and improve the integration of
 1431 cross-modal features. Evaluated across three diverse tissue datasets, HiST consistently surpasses
 1432 current state-of-the-art models, showcasing its performance in predicting spatial gene expression.
 1433 Therefore, the proposed method, which reduces research costs and enhances efficiency, shows promise
 1434 in advancing the development and application of spatial transcriptomics. The primary focus of our
 1435 work is to advance scientific methodology, and we foresee no direct negative societal consequences.
 1436

1437 **H.2 IMPACT ON FUTURE RESEARCH**
 1438

1439 In this paper, we observe that current spatial transcriptomics prediction methods, which operate
 1440 within traditional Euclidean space, neglect the inherent multi-level structure of gene expression
 1441 and overlook biological heterogeneity. Given that hyperbolic space is more suitable for modeling
 1442 hierarchical structures and can capture richer information, it enhances the cross-modal representation
 1443 between pathological images and gene expression. Therefore, we propose HiST, which, in contrast to
 1444 traditional Euclidean approaches, incorporates Multi-Level Representation Extractors and Hierarchi-
 1445 cal Hyperbolic Alignment to fully exploit the intrinsic hierarchical structure of ST data. We hope this
 1446 framework will inspire future studies in multimodal learning that leverage geometric structures, such
 1447 as hyperbolic spaces, in the biological field, like single-cell RNA sequencing, proteomics, and other
 1448 data modalities. Blending the inherent geometry of biological systems with multimodal learning can
 1449 guide the development of future sophisticated representation models.

1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457