

1 **Leveraging information in spatial transcriptomics to predict super-resolution**

2 **gene expression from histology images in tumors**

3 Minxing Pang¹, Kenong Su², Mingyao Li²

4

5 1. Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah

6 University of Science and Technology (KAUST), Thuwal, Saudi Arabia.

7 2. Department of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine,

8 University of Pennsylvania, USA.

9

10 **Correspondence:**

11 suken@pennmedicine.upenn.edu

12 mingyao@pennmedicine.upenn.edu

13

14 **Keywords:** spatial transcriptomics; gene expression prediction; histology; vision transformer.

15 ABSTRACT

16 Recent developments in spatial transcriptomics (ST) technologies have enabled the profiling of
17 transcriptome-wide gene expression while retaining the location information of measured genes
18 within tissues. Moreover, the corresponding high-resolution hematoxylin and eosin-stained
19 histology images are readily available for the ST tissue sections. Since histology images are easy
20 to obtain, it is desirable to leverage information learned from ST to predict gene expression for
21 tissue sections where only histology images are available. Here we present HisToGene, a deep
22 learning model for gene expression prediction from histology images. To account for the spatial
23 dependency of measured spots, HisToGene adopts Vision Transformer, a state-of-the-art method
24 for image recognition. The well-trained HisToGene model can also predict super-resolution gene
25 expression. Through evaluations on 32 HER2+ breast cancer samples with 9,612 spots and 785
26 genes, we show that HisToGene accurately predicts gene expression and outperforms ST-Net
27 both in gene expression prediction and clustering tissue regions using the predicted expression.
28 We further show that the predicted super-resolution gene expression also leads to higher
29 clustering accuracy than observed gene expression. Gene expression predicted from HisToGene
30 enables researchers to generate virtual transcriptomics data at scale and can help elucidate the
31 molecular signatures of tissues.

32 INTRODUCTION

33 Knowledge of the spatial organization of cells and the spatial variation of gene expression in
34 tissues is important in understanding the complex transcriptional architecture of multicellular
35 organisms. For example, in a heterogeneous tissue such as tumor, cancer cells can differ vastly
36 from each other in their gene expression profiles and cellular properties due to residing in distinct
37 tumor microenvironments. Recent advances in spatial transcriptomics (ST) technologies have
38 made it possible to profile gene expression while retaining the spatial location information of the
39 measured genes within tissues (1-6). Applications of the ST technologies in diverse tissues and
40 diseases have transformed our views of transcriptome complexity (7-9). A popular ST technology
41 is based on spatial barcoding followed by next-generation sequencing in which transcriptome-
42 wide gene expression is measured in gene capture locations, referred to as spatially barcoded
43 spots. Such ST technologies include Spatial Transcriptomics (10), 10x Genomics Visium, SLIDE-
44 seq (11), SLIDE-seq2 (12), and many others (13,14). Data from such spatial barcoding-based
45 technology typically include a high-resolution hematoxylin and eosin (H&E)-stained histology
46 image of the tissue section from which the gene expression data are obtained.

47

48 Although ST offers rich information, such data are still expensive to generate, which prevents the
49 applications of ST in large-scale studies. On the other hand, H&E-stained histology images are
50 easier and cheaper to obtain than ST and are routinely generated in clinics. It is desirable to
51 leverage information learned from ST to predict gene expression from histology images. Such
52 predictions can generate virtual ST data, which will enable the study of spatial variations of gene
53 expression at scale. Indeed, several studies have shown that tumor related genes are highly
54 correlated with histological features, suggesting that gene expression can be predicted from
55 histology images. HE2RNA (15), a model based on the integration of multiple data modes, is
56 trained to systematically predict gene expression profiles from whole-slide images without the

57 reliance on expert annotation. ST-Net (16) predicts spatially variable gene expression from
58 histology images using a supervised convolutional neural network (CNN) trained from breast
59 cancer ST data.

60

61 While these methods have shown promising performance, they are not short of limitations.
62 HE2RNA was developed for bulk RNA sequencing and lacks the ability to learn from ST data.
63 While ST-Net is specifically designed for ST, it does not model the spatial location information of
64 each measured spot in their CNN model. Since gene expression often displays local patterns,
65 which are often manifested in the histology images, it is necessary to explicitly model the spatial
66 location information when predicting gene expression. Although CNN-based models have
67 dominated the field of computer vision for years, different parts of an image must be processed
68 in the same way regardless of their position. This intrinsic bias in CNN makes it less ideal for ST
69 data. However, this bias has been recently alleviated by Vision Transformer (17), which internally
70 utilizes self-attention mechanism for divided image patches and has shown strong performance
71 on many tasks, including medical image classification, segmentation (18), and registration (19).

72

73 To utilize these advances in Vision Transformer, we developed HisToGene, an attention-based
74 model that aims to predict gene expression from H&E-stained histology images based on the
75 relationship between histological features and gene expression features learned from a training
76 ST dataset. To account for the spatial dependency of measured spots in ST, HisToGene employs
77 a modified Vision Transformer model, which can naturally model the positional relationship
78 between spots through appropriate positional embedding. Compared to ST-Net (16), our
79 attention-based model considers the spot dependency together with histological features when
80 predicting gene expression. After model training, HisToGene can further predict super-resolution
81 gene expression by averaging predicted gene expression from densely sampled histology image
82 patches. To the best of our knowledge, it is the first time that gene expression can be predicted

83 at such high resolution based on histology images. Gene expression predicted from HisToGene
84 enables researchers to generate virtual transcriptomics data at scale and can help elucidate the
85 molecular signatures of tissues.

86

87 MATERIALS AND METHODS

88 Overview of HisToGene

89 HisToGene takes a ST dataset, possibly with multiple tissue sections obtained from different
90 subjects, as input for model training. For each tissue section in the ST data, it starts by extracting
91 patches from the histology image according to the spatial coordinates and size of each spot in
92 the ST data. The patch embedding and position embedding are then aggregated by learnable
93 linear layers through a modified Vision Transformer model. Next, HisToGene utilizes multi-head
94 attention layers to generate latent embeddings (**Figure. 1a**). With the well-trained model,
95 HisToGene can predict gene expression for each sampled patch from the histology image in a
96 test dataset that only has histology images. Furthermore, using a dense image patch sampling
97 strategy, HisToGene can predict super-resolution gene expression with 4 times of the original
98 patch/spot level resolution by default (**Figure. 1b**), but the resolution can be increased to an even
99 higher level when using more densely sampled patches.

100

101 Data preprocessing

102 HisToGene involves a training step and a prediction step. The training step takes a ST dataset
103 as input, which includes histology images, the gene expression data, and the spatial coordinates
104 for the spatial barcodes. In the training stage, it uses the histology images and the spatial
105 coordinates of the spatially barcoded spots as input, and the corresponding spatial gene
106 expression data as labels. The spatial gene expression data are stored in an $N \times D$ matrix of

107 unique molecular identifier (UMI) counts with N spots and D genes. The coordinates of the spots
108 are stored in an $N \times 2$ matrix indicating the (x, y) location of each spot.

109

110 For the histology image, we extract patches according to the size and location of each spot. We
111 assemble all patches in a tissue section and flatten them into an $N \times (3 \times W \times H)$ matrix as the
112 input for the Vision Transformer, where 3 is the number of channels, and W and H represent the
113 width and height of the patch. In our experiment on the HER2+ breast cancer dataset (20), $W =$
114 $H = 112$ pixels, which correspond to the diameter of each spot in the ST data.

115

116 For the spatial gene expression data, we first identify common genes across all tissue sections in
117 the training ST data. We then select the top 1,000 highly variable genes in each tissue section
118 and eliminate genes that are expressed in less than 1,000 spots across all tissue sections. The
119 gene expression values in each spot are normalized such that the UMI count for each gene is
120 divided by the total UMI counts across all genes in that spot, multiplied by 1,000,000, and then
121 transformed to a natural log scale.

122

123 **The modified Vision Transformer model for gene expression prediction**

124 Model architecture

125 Vision Transformer has been widely used in computer vision for image recognition and
126 outperformed other state-of-the-art methods in the ImageNet Large-Scale Visual Recognition
127 Challenge. The standard Vision Transformer model splits an image into a fixed number of
128 patches. However, in ST data, the number of spots that cover the captured tissue area is not
129 fixed. This property is similar to problems in natural language processing in which the lengths of
130 sentences are also variant. To accommodate variable numbers of spots in ST, we redesign the
131 encoding part of the Vision Transformer model with details described below.

132

133 Encoding of histology image and position features

134 After preprocessing, we use a learnable linear layer \mathbf{W}_h to transform the histology image features
135 from an $N \times (3 \times W \times H)$ matrix \mathbf{F}_h to an $N \times 1024$ matrix \mathbf{E}_h , i.e., $\mathbf{E}_h = \mathbf{F}_h \cdot \mathbf{W}_h$. Another part of
136 the input is the $N \times 2$ matrix, which represents the (x, y) coordinates of each spot in the ST data.
137 The x -coordinate information is represented by an $N \times 1$ vector, which is transformed into a one-
138 hot encoding matrix \mathbf{P}_x with size $N \times m$, where m is the maximum number of x -coordinates
139 among all tissue sections. For the HER2+ breast cancer dataset, $m = 30$. Next, we use a
140 learnable linear layer $\mathbf{W}_x \in \mathbb{R}^{30 \times 1024}$ to transform \mathbf{P}_x into an $N \times 1024$ matrix \mathbf{E}_x , i.e., $\mathbf{E}_x = \mathbf{P}_x \cdot$
141 \mathbf{W}_x . Similar transformations are performed for the y -coordinate vector to obtain an $N \times 1024$
142 encoding matrix \mathbf{E}_y . Finally, we obtain the embedding matrix by summing up the image feature
143 encoding matrix, the x -coordinate encoding matrix, and the y -coordinate encoding matrix,

144
$$\mathbf{E} = \mathbf{E}_h + \mathbf{E}_x + \mathbf{E}_y.$$

145 This embedding matrix is used as the input for the multi-head attention layers as described below.

146

147 Multi-Head Attention layers

148 The Multi-Head Attention module can automatically learn the attention for a "sequence". In
149 language data, the "sequence" is sequence of words in a sentence. In ST data, the "sequence"
150 is a sequence of spots/patches in a tissue section. The multi-head attention is a linear combination
151 of multiple attention heads,

152
$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = [\text{head}_1, \dots, \text{head}_n] \mathbf{W}_0,$$

153 where \mathbf{W}_0 is a learnable 1024×1024 parameter matrix that is used to aggregate the attention
154 heads, n is the number of heads, and $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ represent Query, Key, and Value. In our model, the
155 input matrix is the $N \times 1024$ embedding matrix \mathbf{E} obtained in the previous step. The attention
156 mechanism is defined as

157 $\text{head}_i = \text{Attention}(\mathbf{QW}_i^Q, \mathbf{KW}_i^K, \mathbf{VW}_i^V)$

158
$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V}$$

159 where \mathbf{W}_i^Q , \mathbf{W}_i^K , and \mathbf{W}_i^V are all learnable 1024×1024 parameter matrices. The shape of the input

160 for the attention is $N \times 1024$. In the attention operation, we have two parts, $\text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)$ and \mathbf{V} .

161 The former part is called Attention Map, whose shape is $N \times N$. The latter part is the Value of the
162 self-attention mechanism, where $\mathbf{Q} = \mathbf{K} = \mathbf{V}$. Each column of the Attention Map represents the
163 attention weight contributed from other spots. The Attention Map provides useful information on
164 how the model works. The result of the attention operation is an $N \times 1024$ matrix. We use the
165 output of each multi-head Attention layer as the input for the next layer and repeat this calculation
166 sequentially.

167

168 **Details of the model implementation**

169 We implement the HisToGene model using PyTorch with the following hyper-parameters: learning
170 rate is 10^{-5} , the number of training epochs is 100, drop-out ratio is 0.1, the number of Multi-Head
171 Attention layers is 8, and the number of attention heads is 16.

172

173 **Predicting gene expression at super-resolution**

174 The above trained Vision Transformer model can predict gene expression from histology images
175 with spot level resolution as the training ST data only contain gene expression measured within
176 spatially barcoded spots. However, since the histology image does not have tissue gaps, it is
177 possible to densely sample histology image patches and use predicted gene expression from
178 overlapping patches to estimate gene expression at a resolution that is higher than the original
179 spot. This is analogous to natural language processing, where the Transformer is trained using
180 short sentences but can make predictions for long sentences. In our case, the “sentence” is the

181 sequence of “spots”. Therefore, using the trained model, we can predict the expression for more
182 “spots” within a tissue section.

183

184 The key for our super-resolution gene expression prediction lies in the dense sampling scheme
185 of the histology image patches. First, we uniformly sample patches from the tissue area according
186 to a grid with size of each cell determined by the spot size in the training ST data. For example,
187 in the HER2+ breast cancer dataset, each patch is 112×112 pixels. We sample patches from the
188 histology image such that the entire tissue area is covered by non-overlapping patches of size
189 112×112 pixels. By sampling the patches this way, each sub-patch is covered by 4 patches. Using
190 the trained model, we can predict the gene expression for each sampled patch. For each sub-
191 patch, its gene expression is predicted by the average of the predicted expressions for the 4
192 patches that cover it. Since the size of each sub-patch is only $\frac{1}{4}$ of the original patch, this sampling
193 scheme can increase the gene expression resolution by 4 times. We note that our sampling
194 scheme can be easily modified to increase gene expression prediction resolution to a higher level.

195

196 **RESULTS**

197 **Overview of evaluations**

198 To evaluate the performance of HisToGene, we analyzed the HER2+ breast cancer dataset (20),
199 which includes 36 tissue sections collected from 8 HER2+ breast cancer patients. We retained
200 32 sections from 7 patients that have at least 180 spots per section in the analysis. To evaluate
201 the gene expression prediction accuracy, we conducted leave-one-out (32-fold) cross validation.
202 Specifically, for each section, we used the other 31 sections to train the model and make spatial
203 gene expression predictions for that section. To select genes for prediction, we first considered
204 the top 1,000 highly variable genes for each section and then filtered those that were expressed
205 in less than 1,000 spots across all tissue sections. This filtering left with 9,612 spots and 785

206 genes for model training. We compared HisToGene with ST-Net for gene expression prediction.
207 Since the source codes of ST-Net released by the authors are not maintained, we were not able
208 to analyze the HER2+ breast cancer data using their codes. For comparison, we used our own
209 implementation of ST-Net.

210

211 **HisToGene enables super-resolution gene expression prediction and consistently**
212 **outperforms ST-Net**

213 Since there are no tissue gaps in a histology image, it is possible to densely sample patches from
214 the image, predict gene expression for each sampled patch, and then use the average of the
215 predicted expression from overlapping patches to predict the gene expression for the overlapping
216 tissue area. This allows us to increase the gene expression prediction resolution as the
217 overlapping area among patches is much smaller than the size of the original patch. By averaging
218 predicted gene expression across spatially close patches also reduces prediction uncertainty.
219 Based upon this intuition, we implemented a super-resolution gene expression prediction
220 algorithm in which the modified Vision Transformer in HisToGene can take image patches with
221 variable lengths as input. With the patterned dense sampling of image patches shown in **Figure**
222 **1b**, we can increase the gene expression prediction resolution by 4 times. Using a similar
223 patterned image patch sampling scheme, the gene expression prediction resolution can be
224 increased by 9 times, 25 times, or higher.

225

226 For illustration, we sampled in the image patches such that the gene expression resolution
227 prediction is increased by 4 times. An ideal super-resolution gene expression prediction method
228 should increase the gene expression resolution while retaining the original expression pattern at
229 the patch level, i.e, spot level, as this will ensure no artificial patterns are introduced during the
230 super-resolution gene expression prediction. To evaluate whether HisToGene has this property,
231 we obtained the patch/spot level gene expression from the super-resolution expression predicted

232 by HisToGene. Specifically, we summed up the expression values for 4 adjacent sub-patches to
233 “recover” the patch/spot-level gene expression. Results obtained from this super-resolution
234 expression “recovered” approach were denoted by HisToGene*. We conducted the leave-one-
235 out cross validation for the 32 tissue sections in the HER2+ breast cancer dataset. For each tissue
236 section, we calculated the correlations between the observed gene expression and the predicted
237 gene expression. **Figure 2a** shows that among the 32 tissue sections, HisToGene* predicted
238 patch/spot level gene expression has significantly higher correlations with the observed spot-level
239 gene expression than HisToGene for 19 (59%) sections, whereas HisToGene has significantly
240 higher correlations than HisToGene* for 6 (19%) sections. These results indicate that with the
241 densely sampled image patches as input in the trained prediction model, we can not only increase
242 gene expression prediction resolution, but also the patch/spot-level gene expression prediction
243 accuracy. Such increased accuracy is due to the flexibility of the attention mechanism in handling
244 longer sequences of image patches, which makes the prediction benefit from information in
245 additional batches included in the longer sequences. The increased accuracy is also due to the
246 use of average predicted expression across nearby patches as the random error of the mean is
247 less than that of an individual prediction.

248
249 We also performed gene expression prediction using ST-Net but found its predictions generally
250 yielded low correlations with the observed expression. In fact, for most of the tissue sections, the
251 mean correlations are around zero, and the correlations are not much better even for patient B in
252 which both HisToGene and HisToGene* yielded much higher correlations. We suspect the
253 relatively poor performance of ST-Net is due to its failure in considering the spatial dependency
254 of spots when building the prediction model. As such, patches obtained from different patients are
255 treated in the same way. As reported in the original study (20), there are strong subject-to-subject
256 differences among patients, thus ignoring such differences would lead to less accurate prediction.
257 By contrast, the modified Vision Transformer in HisToGene considers a tissue section as the

258 modeling unit, thus the histology and gene expression relationships are learned within each tissue
259 section, which helps alleviate the subject-to-subject differences among patients. These results
260 demonstrate the importance of considering the spatial dependency of spots when training the
261 prediction model.

262
263 To show that both HisToGene and HisToGene* can predict biologically meaningful information,
264 we conducted gene set enrichment analysis using fgsea (21). Inspired by iPath (22), which sorts
265 genes by positive values, for each approach, we ranked the genes by the average -log10 p-values
266 across all 32 tissue sections, where the p-value for each tissue section was obtained by testing
267 whether the correlation between the observed and the predicted expression values was
268 significantly different from zero. We used the top 100 genes to calculate the enrichment score for
269 each pathway from the C2 canonical pathways in MSigDB (23). Then, the significance for each
270 pathway was assessed by permutations (n=10,000) of the gene list. The enrichment analysis
271 results demonstrate that the highly correlated genes in HisToGene and HisToGene* are enriched
272 in breast-cancer-related pathways (**Figure 2b**). For example, HisToGene*'s top enriched
273 pathways include human thyroid stimulating hormone pathway and REACTOME integration of
274 energy metabolism pathway. Previous studies have reported that thyroid hormones are
275 associated with the risk of breast cancer (24), and energy metabolism is a hallmark of cancer
276 cells and links with the breast cancer brain metastases (25). By contrast, the top enriched
277 pathways for ST-Net show less relevance with breast cancer.

278
279 **Visualization of the predicted gene expression**
280 To gain a better understanding of the predicted gene expression, we next selected the top
281 predicted genes obtained from each method for visualization. For each gene in a tissue section,
282 we calculated the correlation between the observed and the predicted expression values and
283 tested whether the correlation is significantly different from zero. We then ranked the genes by

284 the average -log10 p-values across all 32 tissue sections. **Figure 3a (Supplementary Table 1)**
285 shows the top 4 genes (*GNAS*, *MYL12B*, *FASN*, and *CLDN4*) predicted by HisToGene, where
286 the expression for the best predicted tissue section by HisToGene was visualized. *GNAS* (mean
287 $R = 0.32$) encodes the stimulatory G-protein alpha subunit and regulates the production of the
288 second messenger cyclic AMP. Elevated expression of *GNAS* has been found in several tumors
289 including breast cancer and promotes breast cancer cell proliferation (26). *MYL12B* (mean $R =$
290 0.27) encodes a myosin regulatory subunit that plays an important role in the regulation of non-
291 muscle cell contractile activity via its phosphorylation. A recent study showed that the activity of
292 myosin II in cancer cells drives tumor progression, where the activation of myosin II in non-muscle
293 cells is regulated by phosphorylation of a regulatory light chain such as *MYL12B* (27). *FASN*
294 (mean $R = 0.27$) encodes a key enzyme that is involved in the biogenesis of membrane lipids in
295 proliferating cells and is closely associated with the occurrence and development of tumors (28).
296 Inhibition of *FASN* induces apoptosis in breast cancer cells, making it a potential therapeutic
297 target for breast cancer (29). *CLDN4* (mean $R = 0.26$) encodes a tight junction protein that is
298 required for cell adhesion. It is frequently expressed in primary breast cancers, especially in their
299 metastases, thus is a promising membrane bound molecular imaging and drug target for breast
300 cancer (30-32). As a comparison, we also included the predicted gene expression obtained from
301 HisToGene*, the super-resolution gene expression (denoted by HisToGene_SR), and ST-Net.
302 Although the mean -log10 p-values for genes obtained from HisToGene* are not as significant as
303 HisToGene, the general predicted expression patterns are similar to HisToGene, and for the
304 selected tissue sections, the correlations are similar in magnitude to HisToGene. By contrast, the
305 ST-Net predicted expression shows little correlation with the observed expression.

306

307 **Figure 3b (Supplementary Table 2)** shows the top 4 genes (*GNAS*, *FN1*, *MYL12B*, and *FASN*)
308 predicted by HisToGene*, 3 of them (*GNAS*, *MYL12B*, and *FASN*) were also predicted by
309 HisToGene as the top genes. *FN1* is a gene that shows higher correlation in HisToGene* (mean

310 R = 0.24) than in HisToGene (mean R = 0.22). *FN1* encodes fibronectin, a glycoprotein that is
311 present in a soluble dimeric form in plasma, and in a dimeric or multimeric form at the cell surface
312 and in extracellular matrix. Fibronectin is involved in cell adhesion and migration processes, and
313 high expression of *FN1* is associated with breast cancer invasion and metastasis (33).
314 Interestingly, although *FASN* is among the top 4 best predicted genes by both HisToGene* (mean
315 R = 0.24) and HisToGene (mean R = 0.27), the best predicted tissue sections are different. For
316 the best tissue section predicted by HisToGene* (R = 0.47), the HisToGene correlation is 0.42,
317 only slightly worse than HisToGene*, whereas the correlation obtained from ST-Net prediction is
318 close to 0. In general, we found that HisToGene* has higher correlations than HisToGene,
319 whereas the correlations for ST-Net are often close to 0. For *GNAS*, *FN1*, *MYL12B*, and *FASN*,
320 we further examined the super-resolution gene expression prediction, which revealed fine grained
321 spatial expression patterns that are missed in the original patch/spot level gene expression
322 prediction.

323
324 As a comparison, we also visualized the top 4 genes (*IGHM*, *PPP1R1B*, *IGLC2*, and *PNMT*)
325 predicted by ST-Net (**Supplementary Figure 1 and Supplementary Table 3**). The average
326 correlations for these 4 genes are much lower than the top 4 genes predicted by HisToGene and
327 HisToGene*.

328
329 **HisToGene predicted gene expression can recover pathologists annotated spatial**
330 **domains**

331 Next, we examined if the predicted gene expression can be used to recover the pathologists
332 annotated spatial domains. The HER2+ breast cancer data included 6 tissue sections with
333 pathologists' annotation, which allowed us to further evaluate if the predicted gene expression
334 patterns are biologically meaningful. If the predictions are useful in revealing the underlying
335 biology, we would expect the clusters obtained using the predicted gene expression to agree well

336 with the pathologists annotated spatial domains. We performed K-Means clustering using the
337 predicted gene expression obtained from HisToGene, HisToGene*, and ST-Net. The clustering
338 results were evaluated using Adjusted Rand Index (ARI) by treating pathologists annotated spatial
339 domains as the ground truth. As a comparison, we also performed clustering analysis using the
340 observed gene expression for each tissue section.

341

342 **Figure 4** shows the clustering results and the corresponding ARIs for each method and the results
343 obtained using the observed gene expression. Among the 6 tissue sections with pathologists'
344 annotation, HisToGene* yielded the highest ARIs for 4 sections (B1, C1, D1, and F1), and for
345 sections D1 and F1, the HisToGene*'s ARIs are much higher than those obtained from the
346 observed gene expression and ST-Net. For E1, ST-Net had the highest ARI. For G2, the observed
347 gene expression had the highest ARI. Clustering analysis using observed gene expression is a
348 commonly conducted task in spatial transcriptomics (34-36). Interestingly, HisToGene* had even
349 higher ARIs than the observed gene expression for 4 out of the 6 tissue sections. Since
350 HisToGene* is based on the aggregated super-resolution gene expression, we next performed
351 clustering analysis using the super-resolution gene expression, denoted by HisToGene_SR.
352 Although we cannot directly calculate the ARIs for HisToGene_SR, visual examination indicates
353 that the clustering results agreed well with the pathologists annotated spatial domains, with the
354 tumor region clearly separated from the background.

355

356 **Understanding the HisToGene prediction with attention map**

357 It is intriguing that the predicted super-resolution gene expression led to higher clustering ARIs
358 than the observed gene expression. We next sought to investigate how the super-resolution gene
359 expression prediction works. Attention is a key feature in HisToGene's modified Vision
360 Transformer model. To understand how attention contributes to the HisToGene predicted gene

361 expression, we examined the attention maps. HisToGene's modified Vision Transformer model
362 has 8 layers and each layer has 16 heads, leading to 128 attention maps. For visualization, we
363 randomly chose the attention map from the first, fourth, and eighth layer, which represent the
364 shallow, median, and deep layers. **Figure 5** shows three representative attention maps when
365 HisToGene predicts the expression for a given target spot (the yellow spot in each plot) under the
366 original spot level resolution and the super-resolution settings. The results indicate that 1) in the
367 shallow layer, the modified Vision Transformer model mainly pays attention to the target spot; 2)
368 in the median layer, the model starts to pay attention to some distant spots, and the pattern is
369 especially clear in the super-resolution setting; 3) in the deep layer, the model pays more attention
370 to distant spots that are tumor related. During the model training process, HisToGene can
371 adaptively change the scale of weights when the input sample size changes; for example, the
372 average weight of the super-resolution attention is about 1/10 of the original-resolution attention.
373 It is also evident that in the super-resolution setting, the model utilizes more refined information
374 provided by the neighboring patches. The difference in the attention weights for input with different
375 sample sizes explains why the gene expression prediction for the same image patch can be
376 different when the number of patches is different.

377

378 **DISCUSSION**

379 In this paper, we presented HisToGene, a deep learning method that predicts super-resolution
380 gene expression from histology images in tumors. Trained in a ST dataset, HisToGene models
381 the spatial dependency in gene expression and histological features among spots through a
382 modified Vision Transformer model. HisToGene has been evaluated in 32 heterogeneous HER2+
383 breast cancer tissue sections with 9,612 spots and 785 genes obtained from different patients.
384 Our results consistently show that HisToGene outperformed ST-Net in the spot level gene
385 expression prediction. Additionally, HisToGene can predict gene expression at super-resolution,

386 a feature that ST-Net does not have. To the best of our knowledge, HisToGene is the first method
387 for super-resolution gene expression prediction using histology images. The subsequent
388 clustering analysis using predicted gene expression also shows that HisToGene consistently
389 yielded higher ARIs than ST-Net, and for many of the tissue sections that we analyzed, the ARIs
390 were even higher than those obtained from the observed gene expression. This is likely due to
391 the use of attention, which has the ability to naturally learn from neighborhood. Since the predicted
392 gene expression is based on the histology images, which do not have tissue gaps, it is possible
393 that the consideration of all captured tissue areas in the prediction helped recover expression
394 patterns that are not captured in the observed gene expression.

395
396 Compared to ST-Net, HisToGene benefits from the consideration of spots' dependency and the
397 advanced network architecture, which makes HisToGene robust to heterogeneity among patients.
398 Being robust to batch effects, especially the systematic differences between the training and
399 testing data is an advantage of HisToGene because due to experimental and technical
400 constraints, batch effects are often unavoidable in real studies. HisToGene is robust to
401 heterogeneity among patients due to the following reasons. First, the multi-head attention matrix
402 in HisToGene utilizes the histological features from all spots, implying that when predicting the
403 gene expression for one spot, image features from neighboring spots also contribute.
404 Furthermore, the attention matrix is updated during the training stage, which ensures appropriate
405 adjustment of the neighboring spots' contributions. Second, HisToGene predicts the gene
406 expression for all spots within a tissue section together. These mechanisms enable HisToGene
407 to model the relationship between histology images and the spatial gene expression data for an
408 entire tissue section, hence minimizing batch effects in histology and gene expression features
409 when learning their relationships. By contrast, CNN-based models such as ST-Net consider each
410 spot independently, making these models more sensitive to batch effects.

411

412 HisToGene is computationally fast. To show the computational advantages of HisToGene, we
413 compared its running time for training 31 tissue sections of HER2+ dataset with ST-Net. Our
414 experiment was conducted on CentOS 7 with 24 cores Intel Xeon 8260 CPU and a single NVIDIA
415 V100 (32GB) GPU. On average, it took HisToGene 11 minutes but 27 minutes for ST-Net.

416

417 We mainly focused our analyses on the HER2+ breast cancer dataset in this paper, because this
418 dataset involves a relatively large number of tissue sections and patients. It provides an
419 opportunity to evaluate the performance of HisToGene and ST-Net in the presence of patient
420 heterogeneity. To show the generalizability of HisToGene to other data, we also analyzed a
421 cutaneous squamous cell carcinoma (cSCC) dataset (37), which includes 12 tissue sections
422 obtained from 4 patients, with each patient having 3 sections. Unlike the HER2+ breast cancer
423 dataset, where all tissue sections were generated using the same ST platform, the 12 tissue
424 sections in the cSCC data were generated using a mixture of the relatively low-resolution Spatial
425 Transcriptomics and the more recent 10x Visium platforms. Using the same filtering criteria as the
426 HER2+ breast cancer dataset, 6,630 spots and 134 genes remained for model training and
427 prediction in the cSCC dataset. We also conducted the leave-one-out cross-validation experiment
428 in this dataset, and the results are shown in **Supplementary Note 1**. Due to the relatively small
429 number of tissue sections and patients and the platform heterogeneity among the 12 tissue
430 sections, neither HisToGene nor ST-Net can reliably predict the gene expression with high
431 accuracy. However, HisToGene still yielded higher prediction accuracy than ST-Net. While the
432 requirement of a relatively large training set is a potential limitation of deep learning-based models
433 such as HisToGene, we anticipate that as more and more training ST data become available in
434 the near future, the performance and robustness of HisToGene can be further improved.

435

436 **ACKNOWLEDGEMENTS**

437 The authors would like to thank Alma Andersson and Joakim Lundberg for providing the histology
438 images for the HER2+ breast cancer dataset.

439

440 **AUTHOR CONTRIBUTIONS**

441 This study was conceived of and led by M.L.. M.P. designed the model and algorithm. M.P.
442 implemented the HisToGene software and led the data analysis with input from M.L. and K.S..
443 M.P., K.S., and M.L. wrote the paper.

444

445 **FUNDING**

446 This work was supported by the following grant: R01GM125301 (to M.L.).

447

448 **COMPETING FINANCIAL INTERESTS**

449 The authors declare no competing interests.

450 **FIGURE LEGENDS**

451 **Figure 1. Workflow of HisToGene.** **a**, The modified Vision Transformer in HisToGene starts from
452 sequences of extracted patches from histology images in the training ST dataset. Added by the
453 position embedding obtained from the spatial coordinates of the spots, the patch embedding goes
454 through the Multi-Head Attention layers and is transformed by a linear layer. The final output of
455 the modified Vision Transformer is the spatial gene prediction. This modified Vision Transformer
456 will be trained based on the observed gene expression in the training ST dataset. **b**, Using the
457 trained model, HisToGene can predict super-resolution gene expression. HisToGene first predicts
458 gene expression for each sampled patch from the histology image in a test dataset that only has
459 histology images. Using a dense image patch sampling strategy, HisToGene then predicts the
460 super-resolution gene expression by averaging the predicted gene expression among
461 overlapping patches. With the patch sampling strategy shown in **b**, the resolution is increased 4
462 times, but the resolution can be increased to an even higher level when using more densely
463 sampled patches.

464

465 **Figure 2. Evaluation of gene expression prediction for the HER2+ breast cancer dataset.** **a**,
466 Boxplot of the Pearson correlations between the predicted and observed gene expression for the
467 785 genes predicted by HisToGene, HisToGene*, and ST-Net. HisToGene* is based on the
468 recovered patch/spot level gene expression obtained from the super-resolution gene expression
469 prediction, denoted by HisToGene_SR. **b**, Enrichment analysis for the top 100 predicted genes
470 by HisToGene, HisToGene*, and ST-Net.

471

472 **Figure 3. Visualization of the top predicted genes in the HER2+ breast cancer dataset.** **a**,
473 Relative expression of the top 4 genes predicted by HisToGene. The genes were selected based
474 on the average -log10 p-values across all 32 tissue sections, where the p-value for each tissue

475 section was obtained by testing whether the correlation between the predicted and observed gene
476 expression was significantly different from zero. For each of the 4 genes, the tissue section that
477 had the smallest p-value by HisToGene was selected for visualization. HisToGene* was based
478 on the recovered patch/spot level gene expression obtained from the super-resolution gene
479 expression prediction, denoted by HisToGene_SR. **b**, Relative expression of the top 4 genes
480 predicted by HisToGene*. The genes were selected based on the average -log10 p-values across
481 all 32 tissue sections, where the p-value for each tissue section was obtained by testing whether
482 the correlation between the predicted and observed gene expression was significantly different
483 from zero. For each of the 4 genes, the tissue section that had the smallest p-value by
484 HisToGene* was selected for visualization. HisToGene* was based on the recovered patch/spot
485 level gene expression obtained from the super-resolution gene expression prediction, denoted by
486 HisToGene_SR.

487

488 **Figure 4. Clustering analysis using predicted gene expression in the HER2+ breast cancer**
489 **dataset.** 6 of the 32 tissue sections had pathologists' annotation which allowed us to evaluate
490 whether the predicted gene expression can reveal the pathologists annotated spatial domains.
491 The first column shows the histology image with the pathologists' annotation, where the red lines
492 represent invasive cancer, green lines represent breast glands, yellow lines represent immune
493 infiltrate, and the blue lines represent connective tissue. The remaining columns show the
494 clustering results generated from the observed, ST-Net predicted gene expression, HisToGene
495 predicted gene expression, HisToGene* predicted gene expression, and HisToGene_SR using
496 K-Means clustering algorithm (k=4). Clustering accuracy was evaluated by the Adjusted Rand
497 Index (ARI) between the pathology annotations and the clusters obtained from the predicted gene
498 expression.

499

500 **Figure 5. Attention maps in HisToGene's modified Vision Transformer.** Displayed are the
501 attention weights in the modified Vision Transformer when making gene expression predictions
502 for the target spot (the yellow spot in each figure) in the HER2+ breast cancer dataset. The first
503 row shows the attention maps when predicting the gene expression at the original patch/spot
504 level. The second row shows the attention maps when predicting the gene expression at the
505 super-resolution level.

506

507 **DATA AVAILABILITY**

508 We analyzed two publicly available ST datasets. These data were acquired from the following
509 websites or accession numbers: (1) human HER2-positive breast tumor ST data
510 (<https://github.com/almaan/her2st>); (2) human cutaneous squamous cell carcinoma 10x Visium
511 data (GSE144240).

512

513 **SOFTWARE AVAILABILITY**

514 An open-source implementation of HisToGene can be downloaded from
515 <https://github.com/maxpmx/HisToGene>

516 **REFERENCES**

517 1. Liao, J., Lu, X., Shao, X., Zhu, L. and Fan, X. (2021) Uncovering an Organ's Molecular
518 Architecture at Single-Cell Resolution by Spatially Resolved Transcriptomics. *Trends
519 Biotechnol*, **39**, 43-58.

520 2. Waylen, L.N., Nim, H.T., Martelotto, L.G. and Ramialison, M. (2020) From whole-mount
521 to single-cell spatial assessment of gene expression in 3D. *Commun Biol*, **3**, 602.

522 3. Burgess, D.J. (2019) Spatial transcriptomics coming of age. *Nat Rev Genet*, **20**, 317.

523 4. Asp, M., Bergenstrahle, J. and Lundeberg, J. (2020) Spatially Resolved
524 Transcriptomes-Next Generation Tools for Tissue Exploration. *Bioessays*, **42**,
525 e1900221.

526 5. Crosetto, N., Bienko, M. and van Oudenaarden, A. (2015) Spatially resolved
527 transcriptomics and beyond. *Nat Rev Genet*, **16**, 57-66.

528 6. Moor, A.E. and Itzkovitz, S. (2017) Spatial transcriptomics: paving the way for tissue-
529 level systems biology. *Curr Opin Biotechnol*, **46**, 126-133.

530 7. Asp, M., Giacomello, S., Larsson, L., Wu, C., Furth, D., Qian, X., Wardell, E., Custodio, J.,
531 Reimegard, J., Salmen, F. *et al.* (2019) A Spatiotemporal Organ-Wide Gene Expression
532 and Cell Atlas of the Developing Human Heart. *Cell*, **179**, 1647-1660 e1619.

533 8. Maniatis, S., Aijo, T., Vickovic, S., Braine, C., Kang, K., Mollbrink, A., Fagegaltier, D.,
534 Andrusivova, Z., Saarenpaa, S., Saiz-Castro, G. *et al.* (2019) Spatiotemporal dynamics
535 of molecular pathology in amyotrophic lateral sclerosis. *Science*, **364**, 89-93.

536 9. Chen, W.T., Lu, A., Craessaerts, K., Pavie, B., Sala Frigerio, C., Corthout, N., Qian, X.,
537 Lalakova, J., Kuhnemund, M., Voytyuk, I. *et al.* (2020) Spatial Transcriptomics and In
538 Situ Sequencing to Study Alzheimer's Disease. *Cell*, **182**, 976-991 e919.

539 10. Stahl, P.L., Salmen, F., Vickovic, S., Lundmark, A., Navarro, J.F., Magnusson, J.,
540 Giacomello, S., Asp, M., Westholm, J.O., Huss, M. *et al.* (2016) Visualization and analysis
541 of gene expression in tissue sections by spatial transcriptomics. *Science*, **353**, 78-82.

542 11. Rodrigues, S.G., Stickels, R.R., Goeva, A., Martin, C.A., Murray, E., Vanderburg, C.R.,
543 Welch, J., Chen, L.M., Chen, F. and Macosko, E.Z. (2019) Slide-seq: A scalable
544 technology for measuring genome-wide expression at high spatial resolution. *Science*,
545 **363**, 1463-1467.

546 12. Stickels, R.R., Murray, E., Kumar, P., Li, J., Marshall, J.L., Di Bella, D.J., Arlotta, P.,
547 Macosko, E.Z. and Chen, F. (2020) Highly sensitive spatial transcriptomics at near-
548 cellular resolution with Slide-seqV2. *Nature biotechnology*.

549 13. Liu, Y., Yang, M., Deng, Y., Su, G., Enninful, A., Guo, C.C., Tebaldi, T., Zhang, D., Kim, D.,
550 Bai, Z. *et al.* (2020) High-Spatial-Resolution Multi-Omics Sequencing via Deterministic
551 Barcoding in Tissue. *Cell*, **183**, 1665-1681 e1618.

552 14. Vickovic, S., Eraslan, G., Salmen, F., Klughammer, J., Stenbeck, L., Schapiro, D., Aijo, T.,
553 Bonneau, R., Bergenstrahle, L., Navarro, J.F. *et al.* (2019) High-definition spatial
554 transcriptomics for in situ tissue profiling. *Nat Methods*, **16**, 987-990.

555 15. Schmauch, B., Romagnoni, A., Pronier, E., Saillard, C., Maille, P., Calderaro, J., Kamoun,
556 A., Sefta, M., Toldo, S., Zaslavskiy, M. *et al.* (2020) A deep learning model to predict
557 RNA-Seq expression of tumours from whole slide images. *Nat Commun*, **11**, 3877.

558 16. He, B., Bergenstrahle, L., Stenbeck, L., Abid, A., Andersson, A., Borg, A., Maaskola, J.,
559 Lundeberg, J. and Zou, J. (2020) Integrating spatial gene expression and breast
560 tumour morphology via deep learning. *Nat Biomed Eng*, **4**, 827-834.

561 17. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T.,
562 Dehghani, M., Minderer, M., Heigold, G., Gelly, S. *et al.* (2020) An image is worth 16x16
563 words: transformer for image recognition at scale. *ICLR 2021 Proceedings*.

564 18. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L. and Zhou, Y. (2021)
565 TransUNet: transformers make strong encoders for medical image segmentation.
566 *arxiv*; <https://arxiv.org/abs/2102.04306>.

567 19. Chen, J., He, Y., Frey, E.C., Li, Y. and Du, Y. (2021) ViT-V-Net: vision transformer for
568 unsupervised volumeric medical image registration. *arXiv:2104.06468*.

569 20. Andersson, A., Larsson, L., Stenbeck, L., Salmen, F., Ehinger, A., Wu, S.Z., Al-Eryani, G.,
570 Roden, D., Swarbrick, A., Borg, A. *et al.* (2021) Spatial deconvolution of HER2-positive
571 breast cancer delineates tumor-associated cell type interactions. *Nat Commun*, **12**,
572 6012.

573 21. Sergushichev, A.A. (2016) An algorithm for fast preranked gene set enrichment
574 analysis using cumulative statistic calculation. *bioRxiv*; *doi*:
575 <https://doi.org/10.1101/060012>.

576 22. Su, K., Yu, Q., Shen, R., Sun, S.Y., Moreno, C.S., Li, X. and Qin, Z.S. (2021) Pan-cancer
577 analysis of pathway-based gene expression pattern at the individual level reveals
578 biomarkers of clinical prognosis. *Cell Rep Methods*, **1**.

579 23. Liberzon, A., Birger, C., Thorvaldsdottir, H., Ghandi, M., Mesirov, J.P. and Tamayo, P.
580 (2015) The Molecular Signatures Database (MSigDB) hallmark gene set collection.
581 *Cell Syst*, **1**, 417-425.

582 24. Ortega-Olvera, C., Ulloa-Aguirre, A., Angeles-Llerenas, A., Mainero-Ratchelous, F.E.,
583 Gonzalez-Acevedo, C.E., Hernandez-Blanco, M.L., Ziv, E., Aviles-Santa, L., Perez-
584 Rodriguez, E. and Torres-Mejia, G. (2018) Thyroid hormones and breast cancer
585 association according to menopausal status and body mass index. *Breast Cancer Res*,
586 **20**, 94.

587 25. Chen, E.I., Hewel, J., Krueger, J.S., Tiraby, C., Weber, M.R., Kralli, A., Becker, K., Yates,
588 J.R., 3rd and Felding-Habermann, B. (2007) Adaptation of energy metabolism in
589 breast cancer brain metastases. *Cancer Res*, **67**, 1472-1486.

590 26. Jin, X., Zhu, L., Cui, Z., Tang, J., Xie, M. and Ren, G. (2019) Elevated expression of GNAS
591 promotes breast cancer cell proliferation and migration via the PI3K/AKT/Snail1/E-
592 cadherin axis. *Clin Transl Oncol*, **21**, 1207-1219.

593 27. Georgouli, M., Herraiz, C., Crosas-Molist, E., Fanshawe, B., Maiques, O., Perdrix, A.,
594 Pandya, P., Rodriguez-Hernandez, I., Ilieva, K.M., Cantelli, G. *et al.* (2019) Regional
595 Activation of Myosin II in Cancer Cells Drives Tumor Progression via a Secretory
596 Cross-Talk with the Immune Microenvironment. *Cell*, **176**, 757-774 e723.

597 28. Flavin, R., Peluso, S., Nguyen, P.L. and Loda, M. (2010) Fatty acid synthase as a
598 potential therapeutic target in cancer. *Future Oncology*, **6**, 551-562.

599 29. Menendez, J. and Lupu, R. (2017) Fatty acid synthase (FASN) as a therapeutic target
600 in breast cancer. *Expert Opinion in Therapeutic Targets*, **21**, 1001-1016.

601 30. Jiwa, L.S., van Diest, P.J., Hoefnagel, L.D., Wesseling, J., Wesseling, P., Dutch Distant
602 Breast Cancer Metastases, C. and Moelans, C.B. (2014) Upregulation of Claudin-4,
603 CAIX and GLUT-1 in distant breast cancer metastases. *BMC Cancer*, **14**, 864.

604 31. Lanigan, F., McKiernan, E., Brennan, D.J., Hegarty, S., Millikan, R.C., McBryan, J.,
605 Jirstrom, K., Landberg, G., Martin, F., Duffy, M.J. *et al.* (2009) Increased claudin-4
606 expression is associated with poor prognosis and high tumour grade in breast cancer.
607 *Int J Cancer*, **124**, 2088-2097.

608 32. Szasz, A.M., Tokes, A.M., Micsinai, M., Krenacs, T., Jakab, C., Lukacs, L., Nemeth, Z.,
609 Baranyai, Z., Dede, K., Madaras, L. *et al.* (2011) Prognostic significance of claudin
610 expression changes in breast cancer with regional lymph node metastasis. *Clin Exp
611 Metastasis*, **28**, 55-63.

612 33. Wang, Y., Xu, H., Zhu, B., Qiu, Z. and Lin, Z. (2018) Systematic identification of the key
613 candidate genes in breast cancer stroma. *Cell Mol Biol Lett*, **23**, 44.

614 34. Sun, S., Zhu, J. and Zhou, X. (2020) Statistical analysis of spatial expression patterns
615 for spatially resolved transcriptomic studies. *Nat Methods*, **17**, 193-200.

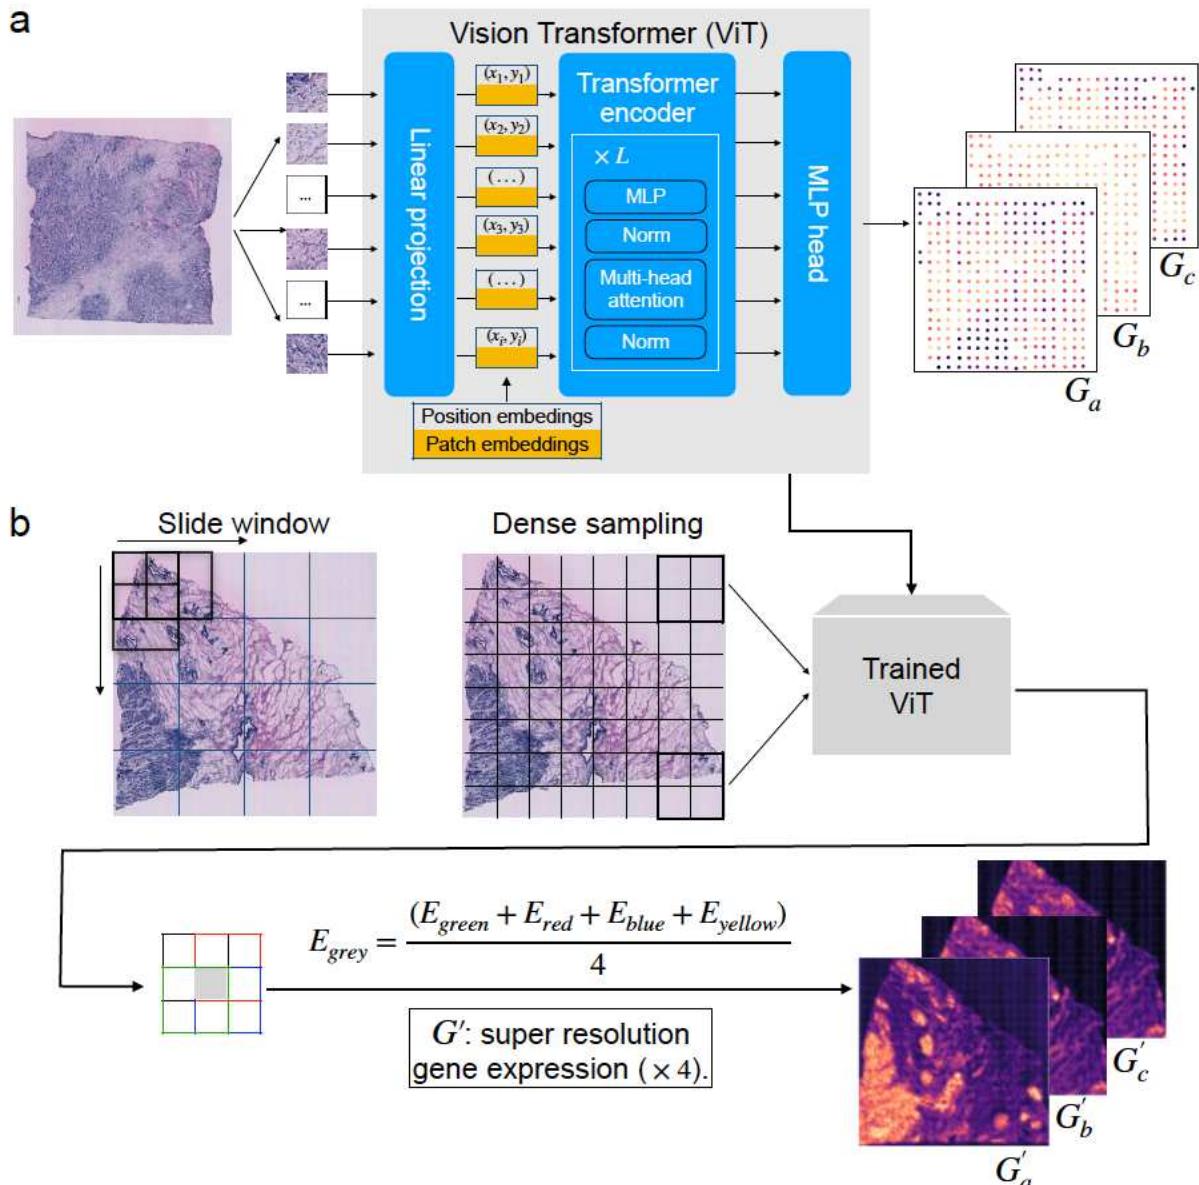
616 35. Hu, J., Li, X., Coleman, K., Schroeder, A., Ma, N., Irwin, D.J., Lee, E.B., Shinohara, R.T. and
617 Li, M. (2021) SpaGCN: Integrating gene expression, spatial location and histology to
618 identify spatial domains and spatially variable genes by graph convolutional network.
619 *Nat Methods*, **18**, 1342-1351.

620 36. Hu, J., Schroeder, A., Coleman, K., Chen, C., Auerbach, B.J. and Li, M. (2021) Statistical
621 and machine learning methods for spatially resolved transcriptomics with histology.
622 *Comput Struct Biotechnol J*, **19**, 3829-3841.

623 37. Ji, A.L., Rubin, A.J., Thrane, K., Jiang, S., Reynolds, D.L., Meyers, R.M., Guo, M.G., George,
624 B.M., Mollbrink, A., Bergenstrahle, J. *et al.* (2020) Multimodal Analysis of Composition
625 and Spatial Architecture in Human Squamous Cell Carcinoma. *Cell*, **182**, 497-514
626 e422.

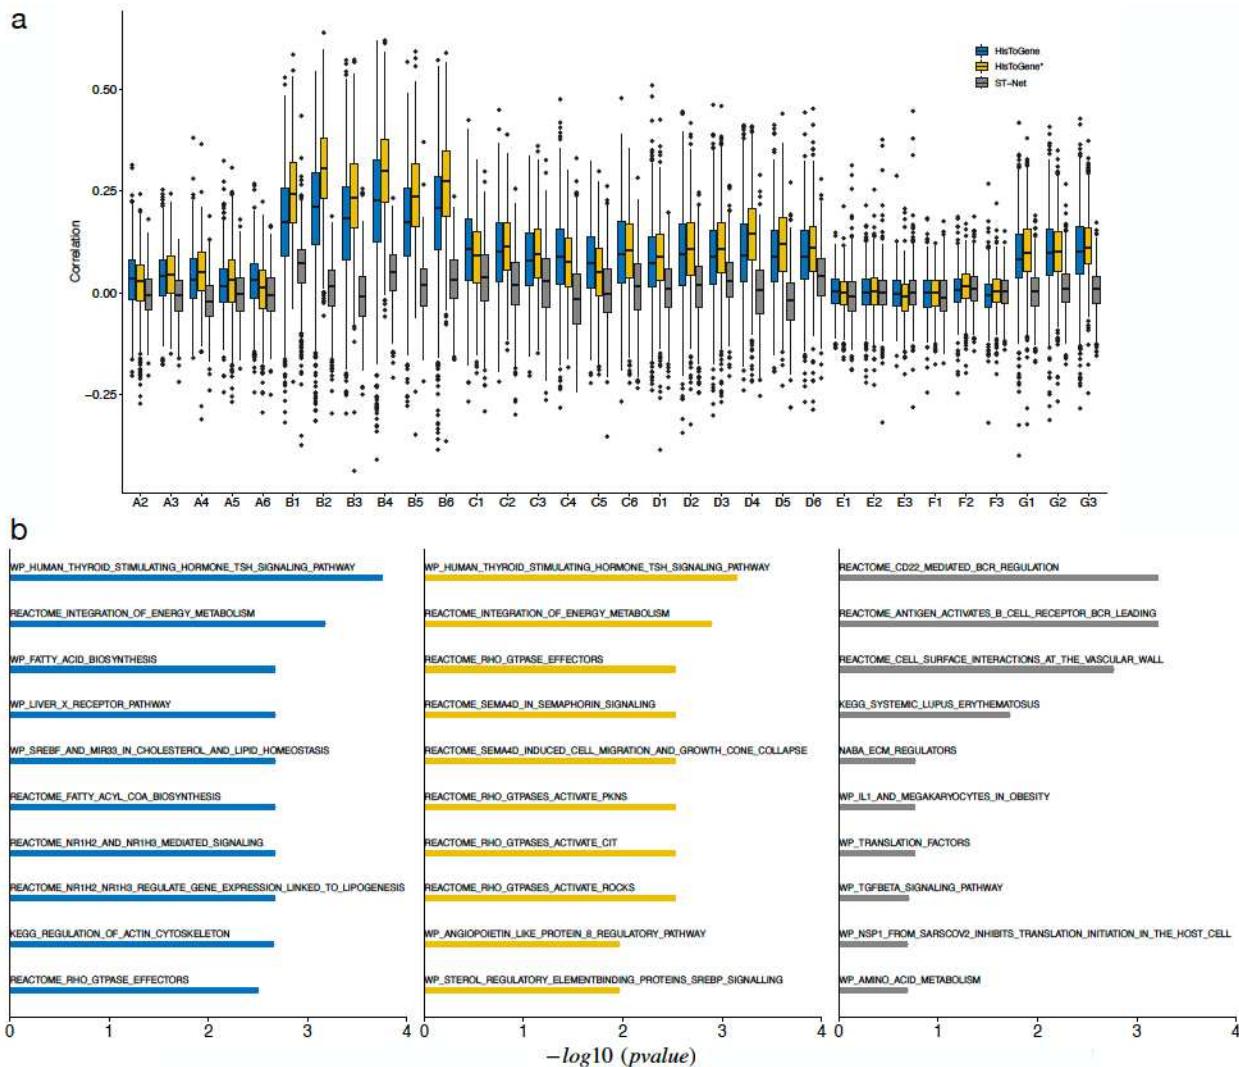
627

628 **Figure 1**



629

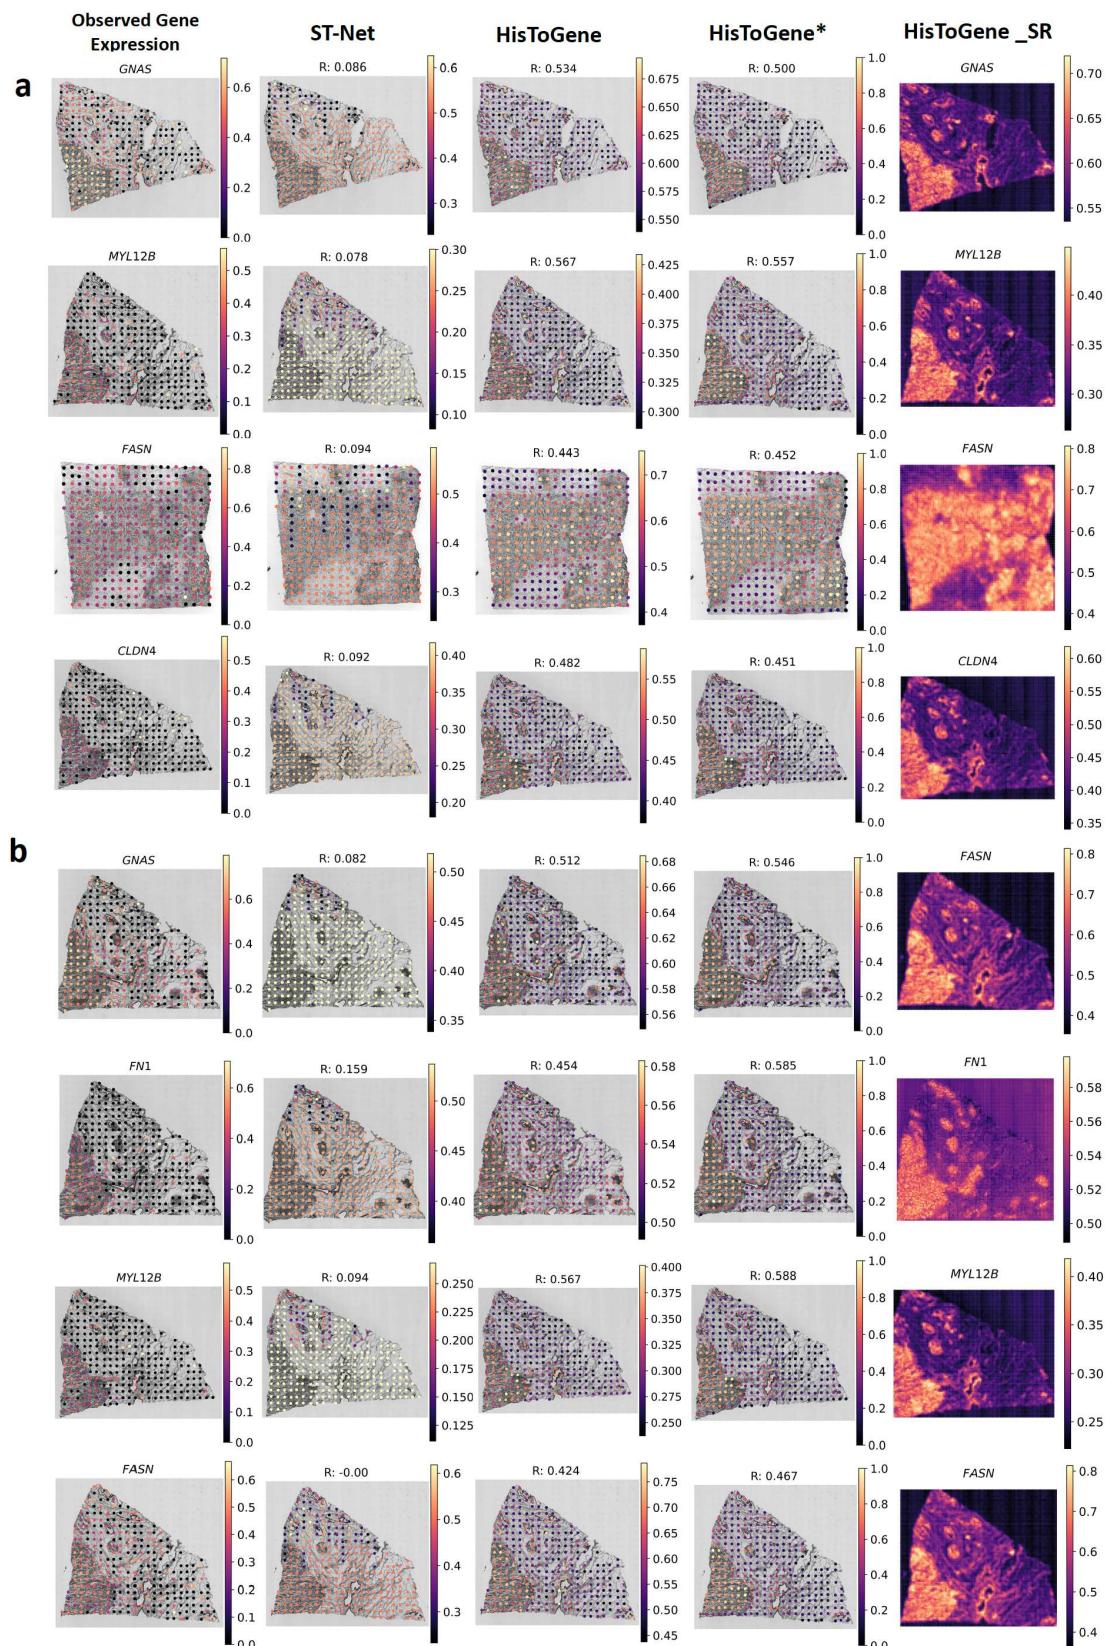
630 **Figure 2**



631

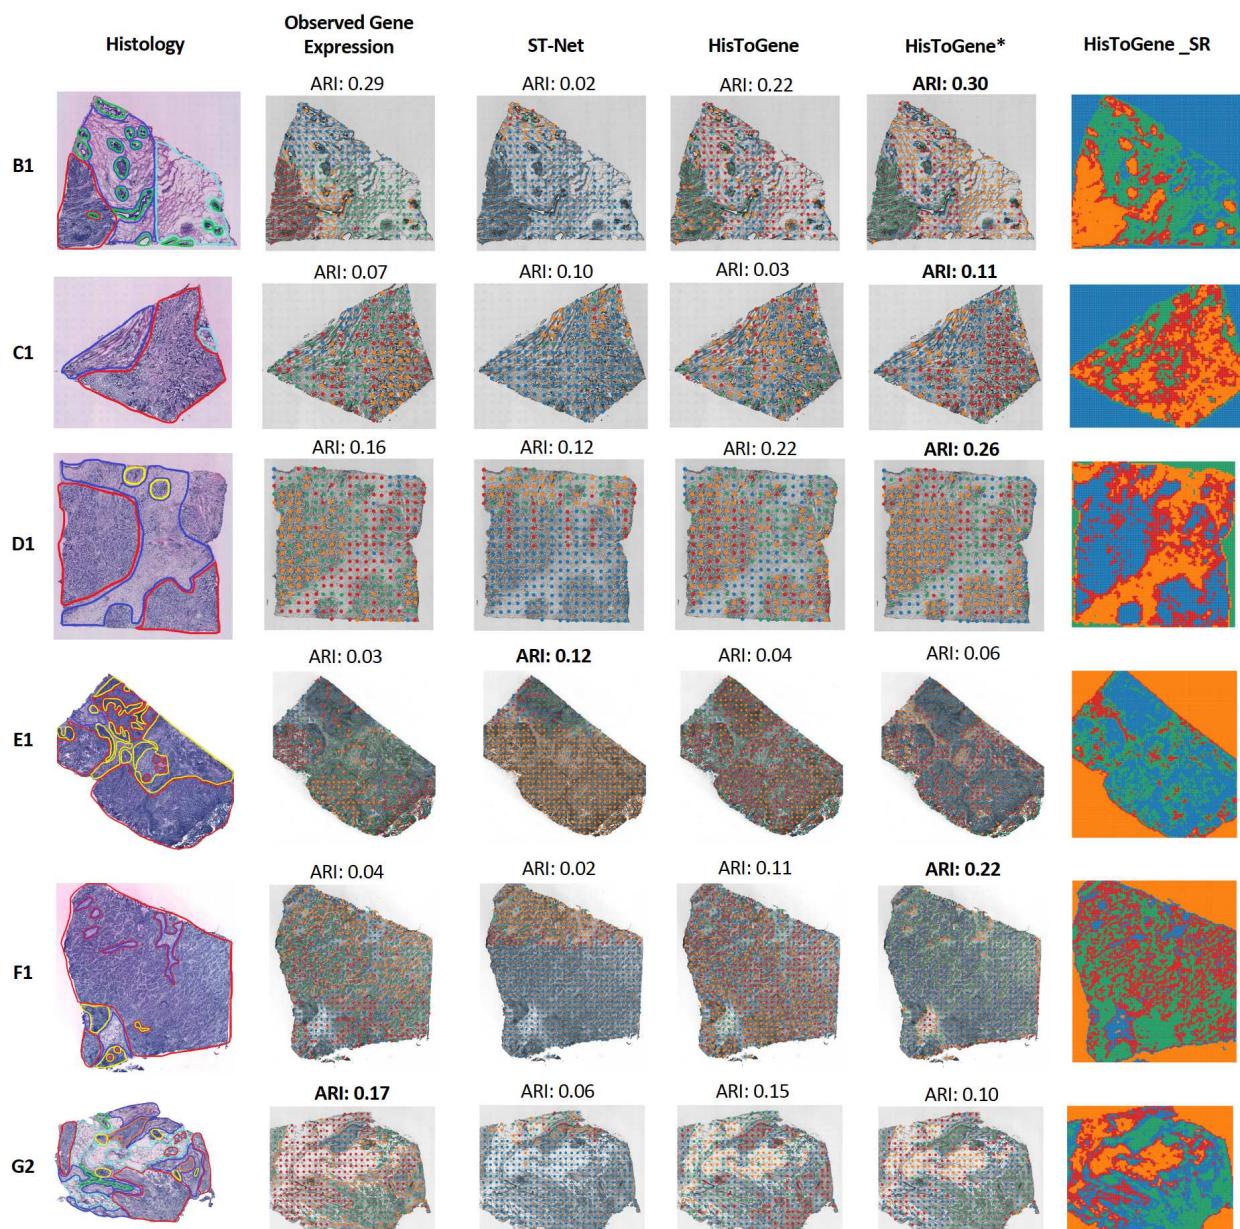
632

633 **Figure 3**



634

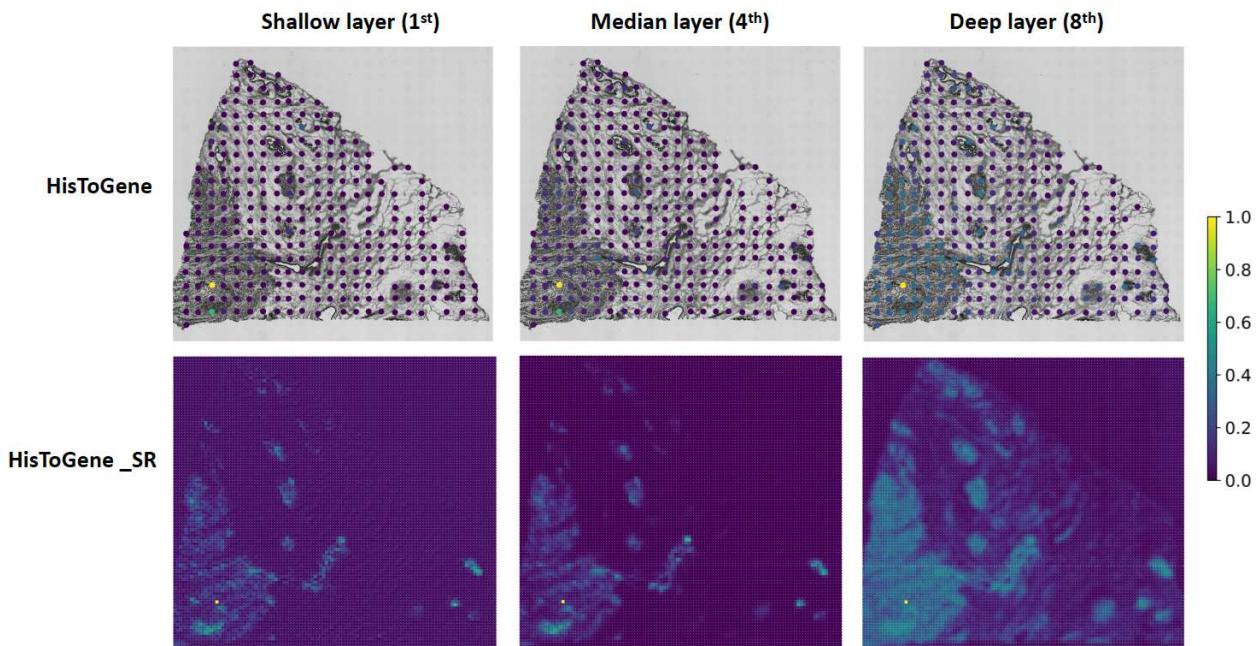
635 **Figure 4**



636

637

638 **Figure 5**



639