

1 CellPatch: a Highly Efficient Foundation Model for Single-Cell

2 Transcriptomics with Heuristic Patching

3 Hanwen Zhu^{1,†}, Yushun Yuan^{2,†}, Jiyuan Yang^{3,†}, Kangwen Cai⁴, Nana Wei⁴,
4 Senxin Zhang², Lu Wang³, Wen-Jie Jiang⁵, YuanChen Sun⁴, An Liu⁴, Futing
5 Lai¹, Yu-Juan Wang⁴, Zeyu Ma^{2,6,*}, Xiaoqi Zheng^{3,*}, Hua-Jun Wu^{1,4,7,*}

6 ¹Department of Biomedical Informatics, School of Basic Medical Sciences, Peking
7 University Health Science Center, Beijing 100191, China

8 ²Department of Mathematics, Shanghai Normal University, Shanghai 200234, China

9 ³Center for Single-Cell Omics, School of Public Health, Shanghai Jiao Tong
10 University School of Medicine, Shanghai 200025, China.

11 ⁴Key laboratory of Carcinogenesis and Translational Research (Ministry of
12 Education/Beijing), Peking University Cancer Hospital & Institute, Beijing, 100142,
13 China

14 ⁵Department of Cardiology and Institute of Vascular Medicine, Peking University Third
15 Hospital, State Key Laboratory of Vascular Homeostasis and Remodeling, NHC Key
16 Laboratory of Cardiovascular Molecular Biology and Regulatory Peptides, Beijing
17 Key Laboratory of Cardiovascular Receptors Research, Peking University, Beijing
18 100191, China.

19 ⁶Shanghai Key Laboratory of Computer Software Testing and Evaluating, Shanghai
20 Development Center of Computer Software Technology, Shanghai, 201112, China

21 ⁷Center for Precision Medicine Multi-Omics Research, Institute of Advanced Clinical
22 Medicine, Peking University, Beijing 100191, China

23 * Corresponding authors: mzy@sscenter.sh.cn (ZM); xqzheng@shsmu.edu.cn (XZ);
24 hjwu@pku.edu.cn (HJW).

25 [†] Co-first authors: Hanwen Zhu, Yushun Yuan, and Jiyuan Yang

26 **Abstract**

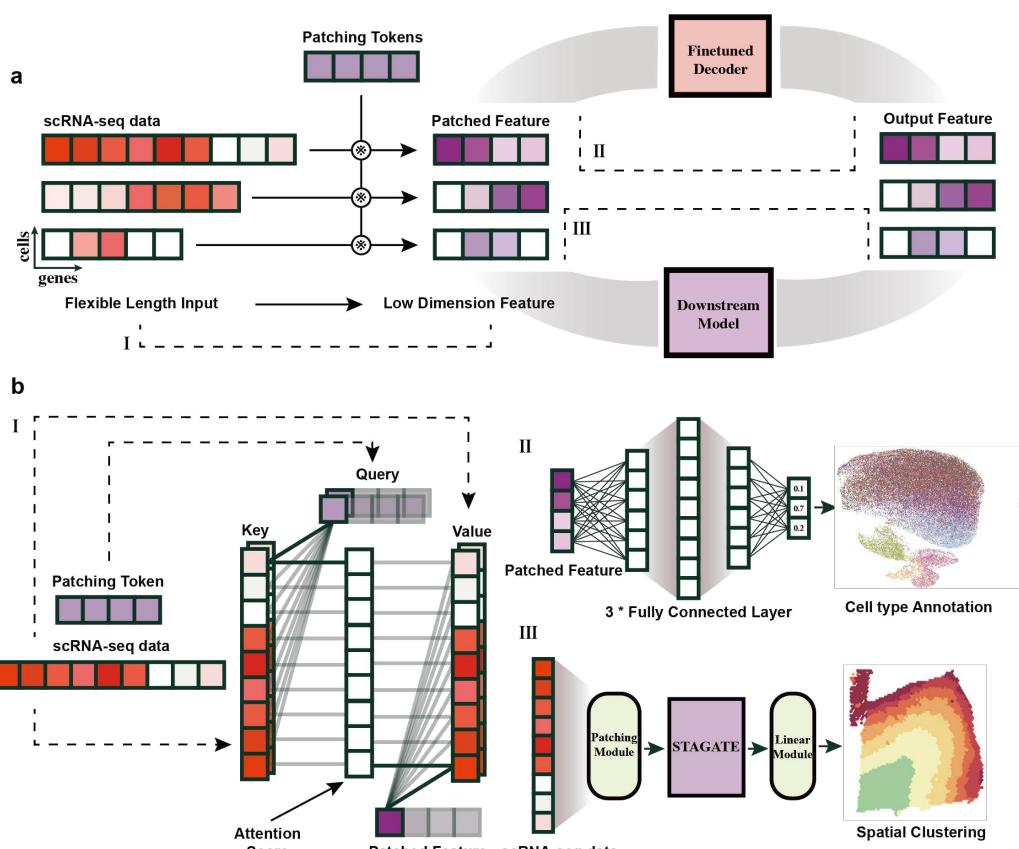
27 The rapid advancement of foundation models has significantly enhanced the
28 analysis of single-cell omics data, enabling researchers to gain deeper insights
29 into the complex interactions between cells and genes across diverse tissues.
30 However, existing foundation models often exhibit excessive complexity,
31 hindering their practical utility for downstream tasks. Here, we present
32 CellPatch, a lightweight foundation model that leverages the strengths of the

33 cross-attention mechanism and patch tokenization to reduce model complexity
34 while extracting efficient biological representations. Comprehensive
35 evaluations conducted on single-cell RNA-sequencing datasets across
36 multiple organs and tissue states demonstrate that CellPatch achieves
37 state-of-the-art performance in downstream analytical tasks while maintaining
38 ultra-low computational costs during both pretraining and finetuning phases.
39 Moreover, the flexibility and scalability of CellPatch allow it to serve as a
40 general framework that can be incorporated with other well established
41 single-cell analysis software, thereby enhancing their performance through
42 transfer learning on diverse downstream tasks.

43 **Main**

44 Single-cell RNA sequencing (scRNA-seq) technology has revolutionized
45 molecular biology in recent years through its capacity to generate
46 high-resolution transcriptomic profiles at unprecedented cellular scales¹⁻³. This
47 technological advancement enables precise measurements of gene
48 expression patterns in individual cells, yielding crucial insights into cellular
49 heterogeneity, developmental trajectories, and disease mechanisms^{4,5}.
50 Meanwhile, The establishment of the Human Cell Atlas⁶ has further facilitated
51 the systematic study of gene regulatory networks at single-cell resolution.
52 Deep learning approaches have emerged as powerful tools for integrating and
53 modeling large-scale scRNA-seq datasets, revealing complex cellular states
54 and dynamics^{7,8}. Recent Transformer-based foundation models of cells, such
55 as scBERT⁹, scGPT⁵, GeneFormer¹⁰, scFoundation¹¹, and LangCell¹²,
56 particular show promise in effective feature extraction and versatile
57 downstream applications, providing new insights into cellular functions. These
58 models demonstrate that Transformer-based architectures¹³ significantly
59 outperform traditional deep learning algorithms, including scVAE¹⁴ and scVI¹⁵,
60 in learning representations of single cells. However, despite their success in
61 processing single-cell data and addressing downstream analytical tasks¹⁶,
62 current foundation models face significant challenges. The high dimensionality
63 of transcriptomics data introduces considerable computational complexity
64 within attention-based architectures, while batch-specific variations in gene
65 detection necessitate complex preprocessing workflows. These challenges
66 collectively constrain the models' transferability, scalability, and flexibility
67 across a range of biological applications.

68 To address these challenges, we present CellPatch, a novel foundation
69 model that employs an effective gene patching strategy to reduce model
70 complexity. Unlike text sequences, which can be chunked based on syntactic
71 order, or images, where pixels can be patched according to spatial position,
72 there is currently no universal patching strategy specifically designed for gene
73 expression data¹⁷. CellPatch employs an innovative cross-attention
74 mechanism where integrated patch tokens automatically patch genes as prior
75 information and extract patch-level features. CellPatch is designed to directly
76 execute downstream tasks using an encoder coupled with a task-specific
77 decoder. Through comprehensive evaluation using extensive real datasets,
78 CellPatch demonstrates superior performance compared to state-of-the-art
79 (SOTA) algorithms. Moreover, the modular design of CellPatch allows for
80 integration with existing SOTA models to enhance their capabilities. When
81 integrated with STAGATE¹⁸, a prominent spatial transcriptomics analysis
82 algorithm, CellPatch significantly improves spatial domain identification,
83 leading to more precise hierarchical structural determinations.

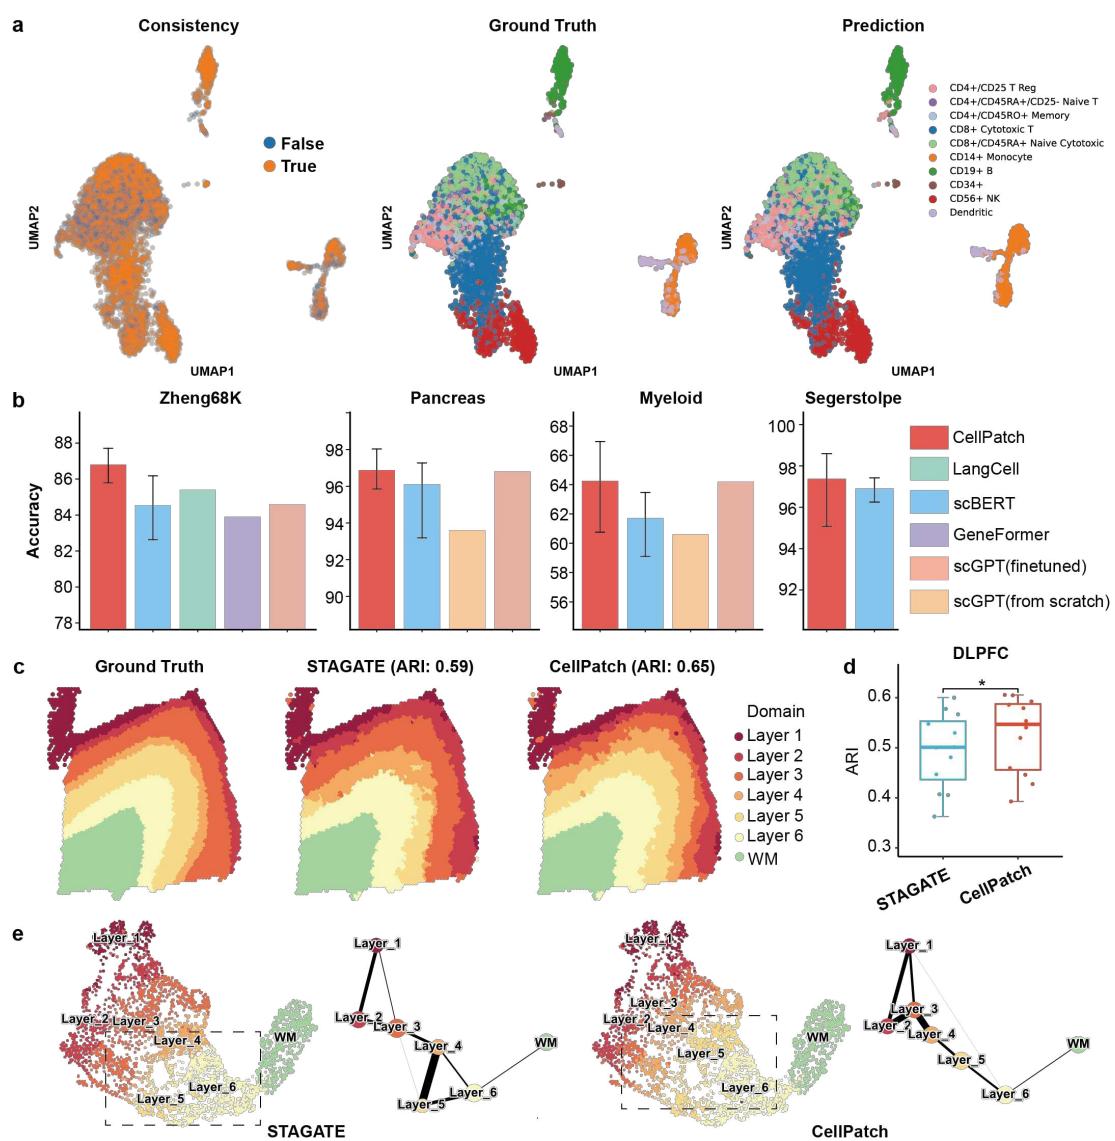


84
85 **Fig. 1 Model Framework.** **a**, Overall framework of CellPatch. CellPatch reduces the
86 dimensionality of single-cell RNA sequencing data by utilizing patch tokens obtained

87 through pretraining, enabling application to various downstream tasks via fine-tuning or
88 embedding. **b**, I Patch Module. Illustrating how a single cell derives patch features through
89 the cross-attention mechanism. II Cell type annotation. Cell features are classified through
90 three dense layers. III Model extension. The patching module and an additional fully
91 connected layer are integrated into the existing STAGATE framework.

92

93 The core innovation of CellPatch lies in its distinctive pretraining method
94 and feature extraction mechanism (Fig. 1a). CellPatch extracts patch features
95 from scRNA-seq data through unsupervised masked gene modeling
96 (Methods), utilizing a comprehensive dataset of 10M cells (described in Data
97 Availability). To optimize computational efficiency, we retained only non-zero
98 elements, with each assigned a gene token that encodes positional
99 information. These elements are subsequently integrated into input features
100 through a sophisticated combination of Patch Tokens and cross-attention
101 mechanisms (Fig. 1b). This approach enables CellPatch to extract gene
102 expression features across diverse cellular contexts and generate
103 standardized patch features of fixed dimensionality. The subsequent
104 application of the self-attention mechanism to these patch features facilitates
105 deep feature extraction, yielding higher-quality representations. A key
106 innovation in the decoder architecture is the incorporation of gene embeddings
107 as semantic prompts, which are not only utilized within the encoder but are
108 also reintroduced in the decoder for enhanced contextual integration. By
109 utilizing the cross-attention mechanism, the model effectively extracts
110 information from patch features and accurately reconstructs gene expression
111 profiles through the decoder. This dual-phase application of gene embeddings
112 enhances the semantic representation capabilities of the gene embeddings.



113

114 **Fig. 2 Performance of CellPatch on scRNA-seq and Spatial Transcriptomic Dataset.**

115 **a**, The cell type prediction performance of CellPatch on Zheng68K dataset. (Left)
116 Consistency between the ground truth and the predicted results of CellPatch; (Middle)
117 UMAP visualization of the ground truth; (Right) UMAP visualization of the predicted
118 results. **b**, Cell type annotation accuracies of various methods on real scRNA-seq
119 datasets, including Zheng68K, Pancreas, Myeloid, and Segerstolpe. Error bars indicate
120 95% confidence intervals across 10 replicates. **c**, Clustering results of original STAGATE
121 and CellPatch+STAGATE on sample 151675 from the DLPFC dataset. The ground truth
122 for tissue regions is based on the manual annotation of six cortex layers or white matter
123 (WM) provided by the original study. **d**, Boxplots illustrating clustering accuracies across
124 all 12 sections for original STAGATE and CellPatch+STAGATE. Asterisk (*) denotes
125 p-value less than 0.05 based on a paired one-sided t-test. **e**, UMAP visualizations and

126 PAGA graphs of sample 151675 generated from original STAGATE and
127 CellPatch+STAGATE embeddings, respectively.

128

129 To showcase the embedding capability of CellPatch, we evaluated its cell
130 type annotation performance against competitive algorithms across diverse
131 scRNA-seq datasets that include a variety of organs and tissue states. Taking
132 Zheng68K dataset as an example, we visualized the prediction results of
133 CellPatch (Fig. 2a) and annotated cells that were not consistently predicted.
134 We further calculated the confusion matrix to assess classification consistency
135 (Supplementary Figure 1). Model performance was systematically assessed
136 using accuracy and F1 score metrics. CellPatch demonstrated superior
137 performance across all benchmark datasets compared to existing methods,
138 achieving higher accuracy (Fig. 2b) and F1 scores (Supplementary Figure 2).
139 Additionally, we calculated the runtime required by different models to perform
140 the annotation task by incrementally increasing number of genes per cell as
141 input (Supplementary Figure 3). CellPatch excelled in both speed and
142 tolerance for higher number of input genes, outperforming other models and
143 demonstrating its efficiency, stability, and broad applicability across various
144 tissue states and experimental platforms.

145 To demonstrate the transfer learning capacity of CellPatch, we used
146 pretrained CellPatch model to assist STAGATE in capturing transcriptional
147 heterogeneity with spatial context. We assessed the spatial domain detection
148 performance of CellPatch combined with STAGATE (CellPatch+STAGATE)
149 and the original STAGATE using the human dorsolateral prefrontal cortex
150 (DLPFC) data by 10X Visium. In the following section, CellPatch+STAGATE is
151 referred to as CellPatch, and the original STAGATE is referred to as
152 STAGATE. In a representative analysis of slice 151675, which comprises
153 3,592 spots across 33,538 genes annotated into six neuronal layers and white
154 matter, CellPatch effectively delineates distinct cortex layers and accurately
155 recovers layer boundaries. In contrast, STAGATE struggles to accurately
156 recover cortical layers 4 and 5, resulting in inferior clustering performance
157 (STAGATE: ARI 0.59, CellPatch: ARI 0.65) (Fig. 2c). Moreover, CellPatch
158 significantly enhances clustering performance across the twelve samples
159 compared to STAGATE (Fig. 2d). Manual annotations and clustering results
160 for other DLPFC slices are shown in Supplementary Figure 4. Furthermore, we

161 performed UMAP and Partition-based graph abstraction (PAGA¹⁹) graph
162 visualization of low-dimensional embeddings generated by CellPatch and
163 STAGATE. CellPatch depicts the spatial trajectory from layer 1 to layer 6 and
164 white matter in the UMAP plot, as well as in the PAGA graph, reflecting the
165 established "inside-out" pattern of corticogenesis. However, since layers 4 and
166 5 cannot be effectively separated, the PAGA graph of STAGATE embedding
167 reveals a circular spatial trajectory pattern from layer 4 to layer 6 (Fig. 2e).
168 Gene set enrichment analysis (GSEA^{20,21}) revealed layer-specific enrichment
169 of key pathways, including myelin sheath, cytoplasmic translation, and
170 synapse-related pathways, highlighting critical transcriptional programs in
171 cortical development (Supplementary Figure 5). These findings were further
172 validated using the STARmap dataset with single-cell resolution
173 (Supplementary Figure 6). In summary, these results demonstrated
174 CellPatch's capacity to enhance the performance of downstream analysis
175 methods through transfer learning.

176 To investigate the interpretability of CellPatch, we analyzed the
177 correlations between attention scores from the first-layer patching process and
178 gene embeddings. Using Zheng68k dataset, we extracted and normalized the
179 attention weights during the transformation of gene expression data into patch
180 features. The normalized attention score matrix revealed distinct cell
181 type-specific patterns (Supplementary Figure 7). For example, in CD8+
182 Cytotoxic T cells, we observed elevated attention scores for the *IL32* gene, a
183 gene known to trigger cytokine production, including *TNF α* , which aligns with
184 the biological role of CD8+ Cytotoxic T cells in producing host defense
185 cytokines. On the heatmap, we highlighted representative cell type-specific
186 markers, demonstrating that the model's attention scores were significantly
187 higher for these genes within their respective cell populations, indicating a
188 targeted recognition of distinct cellular phenotypes.

189 Furthermore, we investigated whether the model-learned gene
190 embeddings contained semantic information beyond merely providing
191 positional context for expression values in downstream tasks. We extracted
192 the top 30 marker genes for each cell type in Zheng68k dataset, along with
193 marker genes reported in the original publication, and generated UMAP
194 visualizations of their embedding weights (Supplementary Figure 8). The
195 analysis revealed remarkable patterns, with embedding weights logically

196 representing gene-gene relationships. For example, multiple genes involved in
197 cytotoxic pathways expressed in CD56+ NK cells clustered near *GZMK* in the
198 lower left region, which was identified as an NK cell type-specific gene in
199 original Zheng68k study. Similarly, platelet-associated genes *PF4* and *PPBP*
200 clustered in the upper region, which also contained multiple CD34+ cell
201 type-specific marker genes. These analyses demonstrate that CellPatch
202 achieves cell type-specific recognition while effectively learning gene-gene
203 interaction relationships, validating both its robust learning capability and
204 post-training interpretability.

205 In summary, the comprehensive experimental results presented herein
206 validate the effectiveness of CellPatch across diverse single-cell
207 transcriptomic and related data types, establishing a robust foundation for its
208 further utilization. The model demonstrates exceptional interpretability and
209 exploration potential in both cell type classification and gene relationship
210 analysis, highlighting the advantages of the CellPatch architecture. While our
211 current analyses are based on training results from 10M data points over 50
212 epochs, UMAP visualization of the pre-training results suggests potential for
213 further model optimization. By increasing training iterations, and broadening
214 the sample space through additional data collection, CellPatch could be a
215 more comprehensive single-cell foundation model. In conclusion, the
216 CellPatch architecture provides an efficient training paradigm for gene-centric
217 data analysis, addresses the gap in scRNA-seq data patching, and achieves
218 excellent performance across multiple aspects of single-cell analysis.

219 **Methods**

220 **CellPatch overview**

221 CellPatch introduces the concept of patch tokens, which extract features from
222 input gene expression data through a cross-attention mechanism. Through
223 pretraining, we obtained a set of patch tokens with high generalization
224 capability, effectively reducing the dimension of gene expression data to a
225 unified meta-gene level. CellPatch improves upon masked gene modeling
226 (MGM), a commonly used training strategy, by extending the utility of gene
227 embeddings: beyond their conventional role as positional encodings in the
228 encoder, they are also employed as prompts in the decoder for expression
229 value reconstruction, thereby enriching their semantic representation.

230 **Gene and expression embedding**

231 CellPatch employs a distinctive embedding approach tailored to the unique
232 structure of single-cell data. Since each expression vector contains mixed
233 information of both gene IDs and their corresponding expression values,
234 mapping these components into feature vectors constitutes a critical step. In
235 this study, we assigned a learnable feature vector to the i -th gene:

236
$$\mathbf{g}_i = f_g(\text{gene_id}_i) \in \mathbb{R}^{\text{dim}}, i = 1, 2, \dots, n,$$

237 where \mathbf{g}_i represents the gene embedding vector, f_g denotes the
238 gene-to-embedding transformation layer, dim denotes the embedding
239 dimension, and n denotes the number of genes. This maps gene IDs into a
240 predefined feature space. For gene expression values, we implemented a
241 robust processing approach. We initially performed standard normalization on
242 the data. Following the methodology employed in scBERT, the expression
243 values are then clipped at a maximum of 5 and rounded to discrete values
244 before being fed into the model. Subsequently, these processed values are
245 transformed through an embedding layer into dim -dimensional expression
246 features

247
$$\mathbf{c}_i = f_c(\text{exp}_i) \in \mathbb{R}^{\text{dim}},$$

248 where \mathbf{c}_i represents the count embedding vector, f_c denotes the
249 count-to-embedding transformation layer, and exp_i represents the count
250 value of gene i .

251 Finally, the feature representation for a single cell at input is denoted as

252
$$\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{n_{\text{gene}}}\},$$

253 where $\mathbf{x}_i = \mathbf{g}_i + \mathbf{c}_i$. This design not only effectively captures information about
254 both gene ID and expression value, but also provides rich contextual
255 information for the model, thereby enhancing performance in downstream
256 tasks.

257 **Pretrain Encoder**

258 In the encoder design of CellPatch, we proposed a novel architecture that
259 differs from traditional Transformer structures. The model predefines a set of

260 Patch Tokens and generates patch features by extracting information from cell
261 features through cross-attention mechanism. Specifically, we first defined a set
262 of **PT** (Patch Tokens) as

263
$$\mathbf{PT} = \{\mathbf{pt}_1, \mathbf{pt}_2, \dots, \mathbf{pt}_{n_{patch}}\} \in \mathbb{R}^{dim},$$

264 where n_{patch} represents the number of patch tokens, a hyperparameter that
265 controls the dimension reduction ratio of gene features to patch features. The
266 cross-attention blocks (Supplementary Note 1) then transform the input cell
267 features \mathbf{X} into patch features **PF** (patch feature), which is formulated as

268
$$\mathbf{PF} = \{\mathbf{pf}_1, \dots, \mathbf{pf}_{n_{patch}}\} = CA(\mathbf{PT}, \mathbf{X}),$$

269 where CA denotes Cross-Attention blocks. Subsequently, these patch
270 features are further processed through self-attention blocks (Supplementary
271 Note 2) to obtain the final patch feature representation **PF'** , which serves as
272 input for various downstream tasks, formulated as

273
$$\mathbf{PF}' = SA(\mathbf{PF}),$$

274 where SA denotes Self-Attention blocks.

275 **Pretrain Decoder**

276 For the pretraining decoder design of CellPatch, we developed a
277 prompt-based cell reconstruction architecture. We innovatively utilized gene
278 embeddings as decoder prompts to reconstruct their corresponding gene
279 expressions. This process extracts information from **PF'** through
280 cross-attention mechanism and reconstructs gene expression via self-attention
281 modules. Specifically, the reconstruction process can be formulated as

282
$$\mathbf{h}_i = CA(\mathbf{g}_i, \mathbf{PF}'),$$

283
$$rec_i = SA(\mathbf{h}_i),$$

284 where \mathbf{h}_i represents the intermediate features extracted from **PF'** using gene
285 embedding \mathbf{g}_i as the query in cross-attention operations, and rec_i denotes
286 the reconstructed expression value for the i -th gene. The cross-attention block
287 and self-attention blocks consist of multi-head attention layers followed by
288 layer normalization and feed-forward networks.

289 During pretraining, we employed mean squared error (MSE) as the loss
290 function to measure the discrepancy between the reconstructed expression
291 values rec_i and the original expression values c_i , which can be formulated as

292
$$L_{Rec} = \text{MSE Loss}(\mathbf{rec}, \mathbf{c}) = \frac{1}{n} \sum_{i=1}^n (rec_i - c_i)^2,$$

293 where n denotes the number of gene prompts input to the decoder for each
294 cell. This reconstruction loss guides the model to learn meaningful
295 representations of gene expression patterns through the pretraining process.

296 Our experimental results demonstrate that, as the number of training
297 iterations increases, the intermediate representations \mathbf{PF}' generated by the
298 model progressively develop the ability to discriminate between different cell
299 types, validating the effectiveness of our training approach (Supplementary
300 Figure 9).

301 **Downstream task**

302 **Classification Task Decoder**

303 For cell type annotation tasks, we implemented a classification decoder
304 architecture that comprising stacked dense layers to predict cell type
305 assignments. Specifically, we first flattened the learned intermediate
306 representations \mathbf{PF}' , which are then processed through a Dense Block
307 consisting of three fully connected layers with dimensionality reduction.

308 The Dense Block projects the flattened features into a probability
309 distribution over the predefined cell type space for each input cell. This
310 transformation can be formulated as

311
$$\mathbf{p}_i = \{p_{i1}, p_{i2}, \dots, p_{i, n_{class}}\} = \text{Dense Block}(\mathbf{PF}'_i).$$

312 To quantify the performance, we adopt the cross-entropy criterion as our
313 loss function, which measures the dissimilarity between the predicted
314 probability distribution and ground truth annotations. The loss function is
315 defined as

316
$$L_i = \text{CrossEntropy}(\mathbf{p}_i, \mathbf{y}_i) = - \sum_{j=1}^{n_{class}} y_{ij} \log(p_{ij}),$$

317 where $\mathbf{p}_i \in R^{n_{class}}$ denotes the predicted probability distribution for the i -th cell,
318 n_{class} represents the cardinality of the cell type set, and \mathbf{y}_i is the one-hot
319 encoded ground truth label.

320 **Downstream Application: Enhancing Spatial Transcriptomics Analysis**
321 **through CellPatch-STAGATE Integration**

322 To further demonstrate the extensibility of CellPatch and its potential for
323 complex downstream tasks, we proposed an innovative integration with
324 STAGATE, a prominent spatial transcriptomics analysis framework.

325 Specifically, we utilized CellPatch to process spatial transcriptomic data
326 and generate \mathbf{PF} , which is then concatenated with the original input features
327 to serve as the enhanced input \mathbf{x}_i for the modified STAGATE module. The
328 downstream analysis is performed through the intermediate representations \mathbf{z}_i ,
329 which can be formulated as

330
$$\mathbf{z}_i = f_{se} \left(\text{concat}(\mathbf{x}_i, \mathbf{PF}) \right),$$

331 where f_{se} denotes the encoder of STAGATE. Given the graph autoencoder
332 architecture of STAGATE, we added an additional Dense layer at the end of
333 decoder to maintain consistency with the original feature space

334
$$\mathbf{rec}_{\mathbf{x}_i} = \text{Dense}(f_{sd}(\mathbf{z}_i)),$$

335 where $\mathbf{rec}_{\mathbf{x}_i}$ denotes the reconstructed features and f_{sd} denotes the decoder
336 of STAGATE.

337 **Implementation details**

338 **Gene Symbol Unification**

339 We established a comprehensive gene reference set based on the
340 CELLxGENE dataset, encompassing 60,690 genes along with their
341 corresponding Ensemble IDs and gene symbols. This reference set
342 incorporates, but not limited to human protein-coding genes and common
343 mitochondrial genes. During the training and fine-tuning processes, each gene
344 is assigned a unique token ID. For genes not present in the training set, new
345 token IDs are dynamically allocated to ensure that no critical information is lost
346 during model adaptation.

347 **Data Preprocessing**

348 We implemented a standardized data preprocessing pipeline using the *Scanpy*
349 framework, which includes quality control filtering, library size normalization
350 (scaling total counts per cell to 1e4), and log transformation (log1p) of
351 normalized counts. For the pretraining phase, we further optimize the input by

352 filtering out zero-expression genes to reduce computational complexity,
353 followed by padding the sequences to a uniform length to facilitate efficient
354 batch processing.

355 **Differential gene expression analysis and spatial domain annotation**

356 We employed the *MAST* algorithm implemented in the *FindMarkers* function of
357 the *Seurat* package to identify differentially expressed genes for each spatial
358 domain with a 5% FDR threshold (Benjamin-Hochberg adjustment). Then,
359 Spatial domains are annotated by marker genes and comparing expression
360 spatial patterns against manual annotations.

361 **Trajectory inference**

362 After obtaining the clustering labels, we employed the PAGA algorithm to
363 depict spatial trajectory.

364 **Gene enrichment analysis**

365 We performed gene set enrichment analysis (GSEA) on the differentially
366 expressed genes sorted by adjusted p-values using *enrichGO* function in the
367 *clusterProfiler*²² package, showing the top five enriched pathways. Gene sets
368 are downloaded from the Molecular Signatures Database (MSigDB, Broad
369 Institute) including C2 (KEGG), C5 (GO BP: biological process, GO CC:
370 cellular component, GO MF: molecular function).

371 **Evaluation metrics**

372 To quantitatively assess the cell type annotation performance, we employed
373 two complementary metrics: accuracy and F1 score. The accuracy metric,
374 defined as the ratio of correctly classified cells to the total number of cells, is
375 formulated as:

$$376 \quad \text{Accuracy} = (TP + TN) / (TP + TN + FP + FN),$$

377 where *TP*, *TN*, *FP*, and *FN* denote true positives, true negatives, false
378 positives, and false negatives, respectively.

379 To provide a more comprehensive evaluation, particularly for imbalanced
380 cell type distributions, we additionally utilized the F1 score, which is the
381 harmonic mean of precision and recall

$$382 \quad F1 = 2 \cdot (Precision \cdot Recall) / (Precision + Recall),$$

383 where

384 $Precision = TP/(TP + FP)$

385 and

386 $Recall = TP/(TP + FN)$.

387 We adopted adjusted Rand index (ARI) to measure clustering

388 performance of CellPatch and STAGATE in spatial domain detection.

389 Specifically, given two sets of clustering labels, the ARI is calculated as

$$ARI = \frac{\sum_{i,j} \binom{n_{ij}}{2} - (\sum_i \binom{n_i}{2} \sum_j \binom{n_j}{2}) / \binom{n}{2}}{\frac{1}{2} (\sum_i \binom{n_i}{2} + \sum_j \binom{n_j}{2}) - (\sum_i \binom{n_i}{2} \sum_j \binom{n_j}{2}) / \binom{n}{2}},$$

390 where n_{ij} is the number of spots overlapped by cluster i and cluster j . n_i

391 and n_j are the number of spots in cluster i and j , respectively.

392 **Model Parameters**

393 **Pretrain**

394 The model is implemented using *PyTorch*. The model architecture is based on
395 a Transformer structure with 64 patch tokens and an embedding dimension of
396 32. The encoder consists of 1 cross-attention layer followed by 2 self-attention
397 layers for pathway embedding processing, with 2 attention heads. The
398 encoder projects the output features into a 10-dimensional space. The
399 decoder comprises 1 cross-attention layer and 1 self-attention layer for gene
400 embedding processing, with final projection to 1-dimensional output. We set
401 the maximum number of genes to 70,000 and utilized 64 learned pathway
402 tokens as preset embeddings.

403 For training, we used the Adam optimizer with a learning rate of 1e-4 and
404 weight decay of 5e-5. The model was trained for 50 epochs with a batch size of
405 512. Input gene expression data underwent log transformation, and random
406 masking was applied to 30% of non-zero values during training. The maximum
407 sequence lengths were set to 3,000 and 1,000 for encoding and decoding
408 respectively.

409 **Cell type annotation**

410 For the downstream cell type annotation task, we finetuned the pretrained
411 model by replacing the original decoder with three dense layers. The decoder
412 architecture consists of a first dense layer that projected the pathway

413 embeddings (dimension: $n_{pathway} \times 32$) to 512 dimensions, followed by a
414 second layer reducing to 100 dimensions, and a final classification layer
415 outputting predictions for cell types. Each of the first two layers is followed by
416 ReLU activation and dropout regularization. The model is finetuned for 50
417 epochs using a batch size of 256. The Adam optimizer was employed with a
418 learning rate of 1e-3 and weight decay of 5e-4. Gradient clipping was applied
419 with a threshold of 5 to prevent exploding gradients.

420 **CellPatch-STAGATE Integration**

421 The first layer of CellPatch is integrated into the STAGATE framework to
422 enhance spatial feature extraction. The model architecture maintains the
423 original cross-attention layers for patch-gene interactions, followed by
424 STAGATE's graph attention network with a two-layer structure to obtain
425 30-dimensional latent representations z_i . Finally, the pathway embeddings are
426 projected through a linear decoder to match the output dimension.

427 The integrated model is trained for 1,000 epochs using the Adam optimizer
428 with a learning rate of 1e-4 and weight decay of 1e-4. Gradient clipping was
429 applied with a threshold of 0.1 to ensure stable training.

430 **Data availability**

431 All data used in this study are publicly accessible. For pretraining, a total of
432 10M cells were sampled from the CELLxGENE database
433 (<https://cellxgene.cziscience.com/datasets>), which comprises 30M cells across
434 various cell types. Our model was evaluated using multiple publicly available
435 datasets, including Zheng68k dataset from the 'Fresh 68K PBMCs' section
436 (<https://support.10xgenomics.com/single-cell-gene-expression/datasets>)²³,
437 pancreatic datasets from the [scGPT foundation model processed datasets](#)
438 (<https://hemberg-lab.github.io/scRNA.seq.datasets/>)²⁴, the Myeloid dataset
439 from Gene Expression Omnibus (GEO: GSE154763)²⁵, and the Segerstolpe
440 dataset from ArrayExpress (E-MTAB-5061)²⁶. The human DLPFC dataset
441 measured by 10X Visium is available at (<http://research.libd.org/spatialLIBD/>)²⁷.
442 The STARmap dataset for mouse medial prefrontal cortex dataset is available
443 at (<https://www.starmapresources.org/data>)²⁸; Data and scripts associated
444 with this study are available at Github.
445 (<https://github.com/HanwenZhu98/CellPatch>).

446 **Code availability**

447 The source code of CellPatch is freely available on Github
448 (<https://github.com/HanwenZhu98/CellPatch>).

449 **References**

- 450 1. Schaum, N. *et al.* Single-cell transcriptomics of 20 mouse organs creates a
451 Tabula Muris: The Tabula Muris Consortium. *Nature* **562**, 367 (2018).
- 452 2. Plass, M. *et al.* Cell type atlas and lineage tree of a whole complex animal
453 by single-cell transcriptomics. *Science* **360**, eaaq1723 (2018).
- 454 3. Cao, J. *et al.* The single-cell transcriptional landscape of mammalian
455 organogenesis. *Nature* **566**, 496–502 (2019).
- 456 4. Heumos, L. *et al.* Best practices for single-cell analysis across modalities.
457 *Nat. Rev. Genet.* **24**, 550–572 (2023).
- 458 5. Cui, H. *et al.* scGPT: toward building a foundation model for single-cell
459 multi-omics using generative AI. *Nat. Methods* 1–11 (2024).
- 460 6. Regev, A. *et al.* The human cell atlas. *elife* **6**, e27041 (2017).
- 461 7. Han, X. *et al.* Pre-trained models: Past, present and future. *AI Open* **2**,
462 225–250 (2021).
- 463 8. Bommasani, R. *et al.* On the opportunities and risks of foundation models.
464 *ArXiv Prepr. ArXiv210807258* (2021).
- 465 9. Yang, F. *et al.* scBERT as a large-scale pretrained deep language model
466 for cell type annotation of single-cell RNA-seq data. *Nat. Mach. Intell.* **4**,
467 852–866 (2022).
- 468 10. Z. Cui, T. Xu, J. Wang, Y. Liao, & Y. Wang. GeneFormer: Learned Gene
469 Compression using Transformer-Based Context Modeling. in *ICASSP 2024 -*
470 *2024 IEEE International Conference on Acoustics, Speech and Signal*
471 *Processing* (ICASSP) 8035–8039 (2024).
472 doi:10.1109/ICASSP48485.2024.10448360.
- 473 11. Hao, M. *et al.* Large-scale foundation model on single-cell transcriptomics.
474 *Nat. Methods* 1–11 (2024).
- 475 12. Zhao, S., Zhang, J., Luo, Y., Wu, Y. & Nie, Z. LangCell: Language-Cell
476 Pre-training for Cell Identity Understanding. *ArXiv Prepr. ArXiv240506708*
477 (2024).
- 478 13. Han, K. *et al.* Transformer in transformer. *Adv. Neural Inf. Process. Syst.*
479 **34**, 15908–15919 (2021).

480 14. Grønbech, C. H. *et al.* scVAE: variational auto-encoders for single-cell
481 gene expression data. *Bioinformatics* **36**, 4415–4422 (2020).

482 15. Lopez, R., Regier, J., Cole, M. B., Jordan, M. I. & Yosef, N. Deep
483 generative modeling for single-cell transcriptomics. *Nat. Methods* **15**,
484 1053–1058 (2018).

485 16. Wei, C., Xie, S. M. & Ma, T. Why do pretrained language models help in
486 downstream tasks? an analysis of head and prompt tuning. *Adv. Neural Inf.*
487 *Process. Syst.* **34**, 16158–16170 (2021).

488 17. Yuan, L. *et al.* Tokens-to-token vit: Training vision transformers from
489 scratch on imagenet. in *Proceedings of the IEEE/CVF international conference*
490 *on computer vision* 558–567 (2021).

491 18. Dong, K. & Zhang, S. Deciphering spatial domains from spatially resolved
492 transcriptomics with an adaptive graph attention auto-encoder. *Nat. Commun.*
493 **13**, 1739 (2022).

494 19. Wolf, F. A. *et al.* PAGA: graph abstraction reconciles clustering with
495 trajectory inference through a topology preserving map of single cells.
496 *Genome Biol.* **20**, 59 (2019).

497 20. Subramanian, A. *et al.* Gene set enrichment analysis: A knowledge-based
498 approach for interpreting genome-wide expression profiles. *Proc. Natl. Acad.*
499 *Sci.* **102**, 15545–15550 (2005).

500 21. Mootha, V. K. *et al.* PGC-1α-responsive genes involved in oxidative
501 phosphorylation are coordinately downregulated in human diabetes. *Nat.*
502 *Genet.* **34**, 267–273 (2003).

503 22. Yu, G., Wang, L.-G., Han, Y. & He, Q.-Y. clusterProfiler: an R Package for
504 Comparing Biological Themes Among Gene Clusters. *OMICS J. Integr. Biol.*
505 **16**, 284–287 (2012).

506 23. Zheng, G. X. Y. *et al.* Massively parallel digital transcriptional profiling of
507 single cells. *Nat. Commun.* **8**, 14049 (2017).

508 24. Chen, J. *et al.* Transformer for one stop interpretable cell type annotation.
509 *Nat. Commun.* **14**, 223 (2023).

510 25. Cheng, S. *et al.* A pan-cancer single-cell transcriptional atlas of tumor
511 infiltrating myeloid cells. *Cell* **184**, 792-809. e23 (2021).

512 26. Segerstolpe, Å. *et al.* Single-cell transcriptome profiling of human
513 pancreatic islets in health and type 2 diabetes. *Cell Metab.* **24**, 593–607
514 (2016).

515 27. Maynard, K. R. *et al.* Transcriptome-scale spatial gene expression in the
516 human dorsolateral prefrontal cortex. *Nat. Neurosci.* **24**, 425–436 (2021).

517 28. Wang, X. *et al.* Three-dimensional intact-tissue sequencing of single-cell
518 transcriptional states. *Science* **361**, eaat5691 (2018).

519

520 **Acknowledgements**

521 This work was supported by the National Natural Science Foundation of China
522 (32270683 and 32470662), the Beijing Natural Science Foundation (5242006),
523 the Fundamental Research Funds for the Central Universities
524 (BMU2021YJ064 and PKU2022LCXQ027), and the National Key R&D
525 Program of China (2021YFC1712805) to H.J.W.; the National Natural Science
526 Foundation of China (61972257), the Natural Science Foundation of Shanghai
527 (20JC1413800), and the National Key R&D Program of China
528 (2018YFA0900600) to X.Z.. We gratefully acknowledge the High-performance
529 Computing Platform of Peking University for conducting the data analyses.

530 **Author information**

531 **Ethics declaration**

532 No ethical approval was required for this study. All utilized public datasets
533 were generated by other organizations that obtained ethical approval.

534 **Competing interests**

535 The authors declare no competing interests.

536 **Supplementary information**

537 Evaluation Metrics & Supplementary Figures