S3R: Spatially Smooth and Sparse Regression Reveals High-Dimensional Regulatory Networks in Spatial Transcriptomics

Xinyu Zhou[1,2], Pengtao Dang[3], Haixu Tang[1], Laura Xianlu Peng[4], Jen Jen Yeh[4], Rosie Sears[3], Brian Neelon[5], Teresa Zimmers[3], Chi Zhang[2,3, $], Sha Cao[2,3 $]

[1] Department of Computer Science, Indiana University, Bloomington, IN, 46202, USA

[2] Center for Computational Biology and Bioinformatics, Department of Medical and Molecular Genetics, Department of Biostatistics and Health Data Science, Indiana University School of Medicine, Indianapolis, IN, 46202, USA

[3] Brenden-Colson Center for Pancreatic Care, Department of Biomedical Engineering, Department of Molecular and Medical Genetics, Department of Cell, Developmental and Cancer Biology, Oregon Health and Science University, Portland, OR, 97202, USA

[4] Departments of Surgery and Pharmacology, University of North Carolina, Chapel Hill, NC, 27514, USA

[5] Departments of Department of Public Health Sciences, Medical University of South Carolina, Charleston, SC 29425, USA

[$] Corresponding author: Chi Zhang (zhangchi@ohsu.edu); Sha Cao (caosh@ohsu.edu)

Abstract

Spatial transcriptomics (ST) data demands models that recover how associations among molecular and cellular features change across tissue while contending with noise, collinearity, cell mixing, and thousands of predictors. We present Spatially Smooth Sparse Regression (S3R), a general framework that estimates location-specific coefficients linking a response feature to high-dimensional spatial predictors. S3R unites structured sparsity with a minimum-spanning-tree–guided smoothness penalty, yielding coefficient fields that are coherent within neighborhoods yet permit sharp boundaries. S3R enables large scale data analysis with an efficient implementation using a reduced MST graph, multi-GPU training, and parallel hyperparameter search. In synthetic data, S3R accurately recovers spatially varying effects, selects relevant predictors, and preserves known boundaries. Applied to human dorsolateral prefrontal cortex, S3R recapitulates layer-specific target–TF associations with concordant layer-wise correlations in matched single-cell data. In acute *Haemophilus ducreyi* skin infection, S3R converts spot-level mixtures into cell type–attributed expression fields and reveals spatial gradients; applying SVG tests to these fields increases concordance and recovers gradients missed by spot-level methods. In pancreatic ductal adenocarcinoma, S3R constructs cross–cell type, cross-gene co-variation tensors that prioritize interactions among cell types with interacting genes enriching pathways consistent with known biology. Because responses and predictors in S3R are user-defined, it could flexibly address diverse ST questions within a single, scalable, and interpretable regression framework.

Keywords: spatial transcriptomics, spatially variable relationships, high-dimensional data, MST, Adam optimizer

Introduction

Spatial transcriptomics (ST) has transformed our ability to map gene expression in intact tissues, revealing how cellular neighborhoods and tissue architecture shape molecular phenotypes [1]. By preserving spatial context that is lost in dissociated single-cell assays, ST enables questions that were previously out of reach [2]: how adjacent cell populations communicate [3, 4], how a cell's position within a layer or niche modulates its expression program [5], and how proximity among distinct cell types influences function [6, 7]. Existing computational approaches have delivered important advances, by identifying spatially variable genes and segmenting tissues into expression domains or biological niches [4, 8, 9], yet they often analyze a small set of features at a time, treat locations as discrete blocks, or rely on dimension reduction that can occlude important signals [10, 11]. What remains comparatively underexplored is not merely where spatial features are patterned, but how ***relationships*** among spatial features, such as genes, transcriptional regulators, cell-type abundances, and other covariates, change with position [12, 13]. Classical regression treats coefficients as global, assuming stationarity; mixture models allow a few discrete regimes but do not exploit coordinates and often miss the gradual transitions that characterize tissue architecture [14].

Addressing this gap in spatial transcriptomics calls for models that estimate location-specific associations while remaining stable enough to be interpretable in the presence of noise, collinearity, missing data and spatial autocorrelation, and respecting sharp boundaries where microenvironments change. High dimensionality further complicates the computation when thousands of spatial features (such as genes) are candidates to explain spatial variation, and spot-level measurements aggregate contributions from mixtures of cells in the case of non-single cell resolution ST technologies such as Visium and NanoString GeoMx DSP. A framework that could enforce parsimony for interpretability, and allow coefficients to vary smoothly where appropriate, and still permit sharp changes at microenvironmental boundaries, is therefore needed.

Here we introduce Spatially Smooth Sparse Regression (S3R), a general approach for modeling spatially varying associations in high-dimensional ST data. S3R couples structured sparsity to allow only a small subset of predictors to be retained for interpretability and statistical stability, with spatially adaptive smoothness that borrows strength within neighborhoods yet permits sharp transitions at histological boundaries. Concretely, S3R estimates location-specific coefficients for each predictor–response pair, and produces spatial fields whose magnitudes encode the strength of association and whose gradients capture continuous trends or abrupt shifts between microenvironmental niches. Because responses and predictors can be flexibly defined, such as gene expression as a function of transcription factor expression or cell-type proportions, S3R could unify several analysis goals within a single, interpretable regression framework.

We applied S3R to two different types of data: data with relatively well-defined and pure spatial structure and cellular neighborhoods such as brain and infected skin, and data with highly mixed spatial intertwining of different cell types, such as pancreatic and breast cancer tissues. We demonstrate S3R across three distinct settings that highlight complementary strengths. In human dorsolateral prefrontal cortex (DLPFC), where laminar structure is well annotated, S3R recovers spatially coherent, layer-specific associations between targets and transcriptional regulators, with concordant layer-wise correlations in matched single-cell data. In acute *Haemophilus ducreyi* skin infection, S3R converts spot-level mixtures into cell type–attributed expression fields and reveals tissue-scale gradients that are not visible to other methods. In pancreatic ductal adenocarcinoma (PDAC), S3R enables transcriptome-wide, cross–cell type, cross-gene co-variation analysis, and summarizes replicated gene-gene couplings between interacting cell type pairs. Across these applications, S3R moves beyond cataloging "where" expression is high or low to infer "how" spatial context shapes associations and yields testable hypotheses about cell type–specific regulation and intercellular organization in situ. Together, these results position S3R as a practical, general tool for dissecting spatially varying relationships in complex tissues.

Results

Methods overview

Spatially smooth and Sparse Regression (S3R) is a computational framework designed to resolve various spatially varying relationships in noisy, high-dimensional spatial transcriptomics data (**Fig. 1**). The method addresses two critical challenges in modeling the relationship among spatial features: biologically meaningful relationships may not be static across tissue regions, and the spatial features that encode such relationships are frequently masked by measurement noise, collinearity, and a high dimensional predictor set. S3R unifies sparsity-inducing regularization with spatially aware smoothing constraints. At its core, S3R models the relationship between one spatial feature, namely the response feature (**Fig. 1a**), and a set of other spatial features (**Fig. 1b**), namely the predictor features, through spatially adaptive coefficients, $\boldsymbol{\beta_k(s_i)}$ (**Fig. 1c**), which quantify the influence of each predictor feature $\boldsymbol{k}$ at location $\boldsymbol{s_i}$ onto the response feature. To ensure robustness, the framework imposes three complementary penalties: a *minimum spanning tree (MST)-based spatial smoothness term* that enforces local homogeneity in coefficients across adjacent tissue regions (**Fig. 1d**)[15] , an $\boldsymbol{L_1}$ penalty to select predictor features with non-zero effects at individual locations, and a *group* $\mathbf{L_2}$ penalty to eliminate non-predictive features globally (**Fig. 1e**). The MST, constructed from spatial coordinates using Euclidean distances, reduces computational complexity by representing tissue topology with only $\boldsymbol{n-1}$ edges while maintaining full connectivity. Optimization is achieved through the highly efficient Adam optimizer [16] using a scalable multi-GPU parallel framework [17], with hyperparameters tuned in parallel using Optuna[18]. These optimizations enable efficient training and model selection on large-scale ST datasets. On comprehensively configured simulation settings, S3R demonstrated superior performance in recovering the coefficient clustering patterns by visual inspection (**Extended Fig. 1a**) as well as Adjusted Rand Index (ARI) (**Extended Fig. 1b**), for both high dimensional and relatively lower dimensional predictor settings, and different configurations of the spatial heterogeneity. See Methods and Materials for more details.

S3R is a flexible regression framework that accommodates diverse predictor–response configurations while estimating location-specific effects $\boldsymbol{\beta}_k(\boldsymbol{s}_i)$. For instance, when the response is the expression of a target gene and the predictors are the expressions of putative transcription factors, $\boldsymbol{\beta}_k(\boldsymbol{s}_i)$ quantifies the association of transcription factor $\boldsymbol{k}$ with the target's expression at location $\boldsymbol{s}_i$. When the response is a gene's expression and the predictors are spot-level cell-type proportions, $\boldsymbol{\beta}_k(\boldsymbol{s}_i)$ describes how local variation in the abundance of cell type $\boldsymbol{k}$ relates to that gene's expression across the tissue. Conversely, when the response is the spatial abundance of a given cell type (for example, its proportion per spot) and the predictors are genome-wide gene expressions or gene-set summaries, $\boldsymbol{\beta}_k(\boldsymbol{s}_i)$ characterizes location-dependent enrichment or depletion of that cell type associated with the expression of gene or gene set $\boldsymbol{k}$. In the following, we present case studies that illustrate these settings and demonstrate the breadth of biological questions S3R can address.

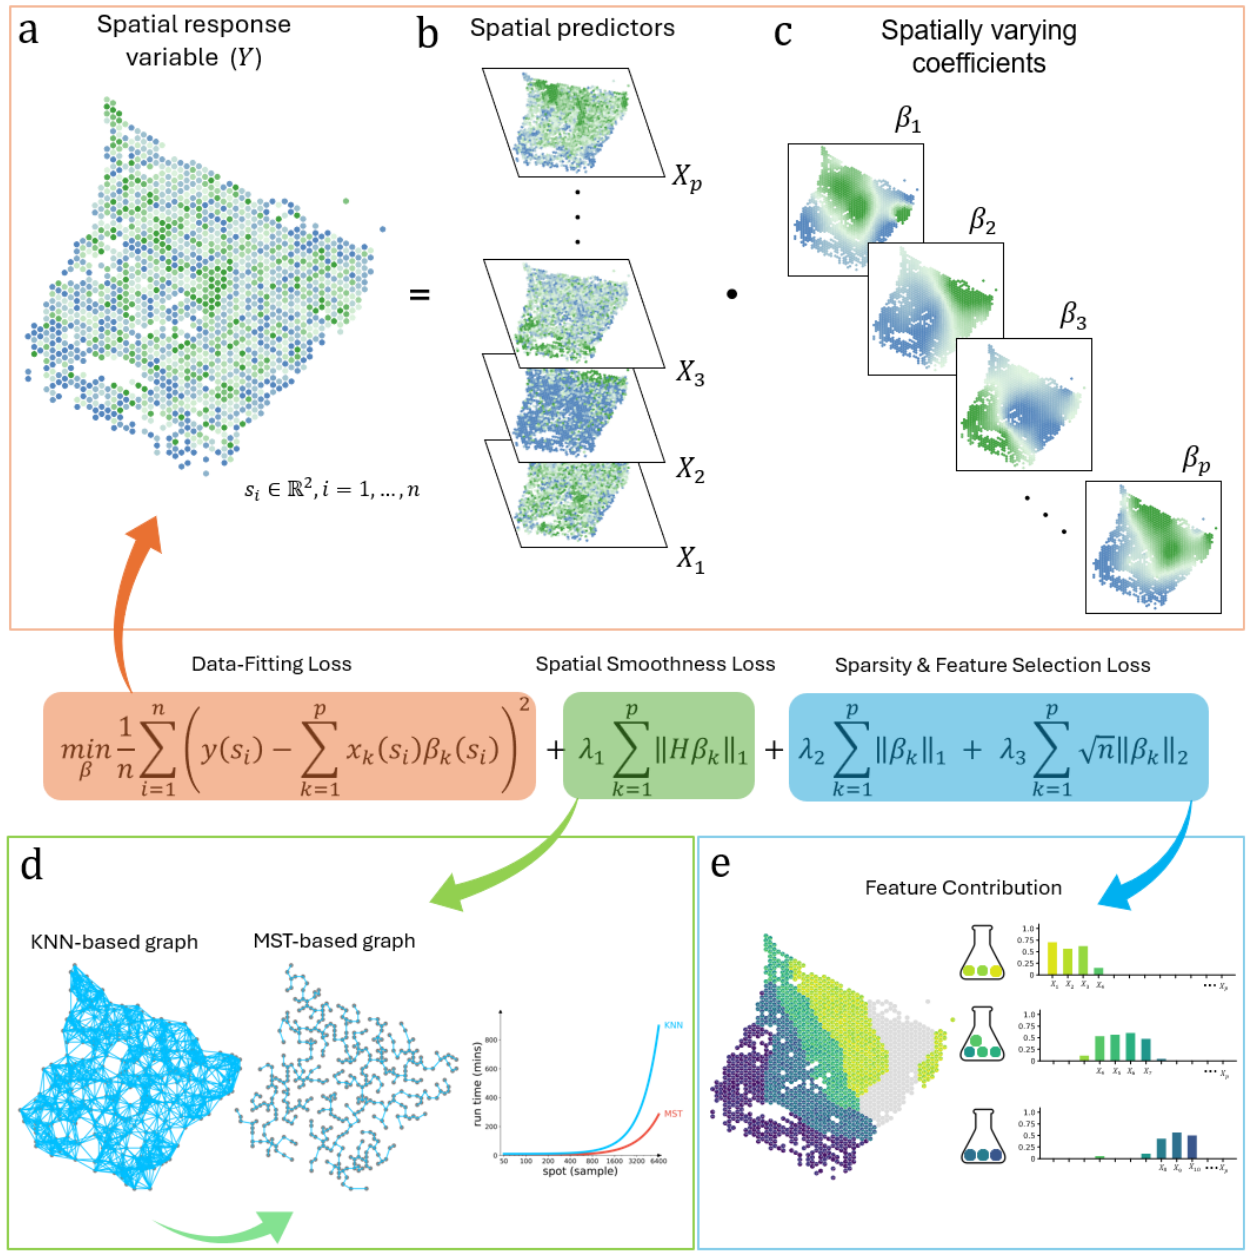

Fig. 1 | Workflow of the S3R method.

(a-b) Input to S3R includes a spatial response feature and spatial predictor profile matrices across spatial locations.

(c) The model outputs estimated spatially varying coefficients for each predictor at each spatial location .

Center, Composite loss function: The estimation is performed by minimizing a composite loss function, which integrates three terms: a data-fitting loss, a spatial smoothness loss, and sparsity/feature-selection penalties.

(d) A minimum spanning tree (MST) graph is constructed from spatial coordinates to enforce spatial smoothness in coefficients .

(e) The L_1 penalty shrinks individual coefficients to zero, while the group L_2 penalty eliminates non-predictive features globally, collectively, they select predictors with spatially varying contributions to the response.

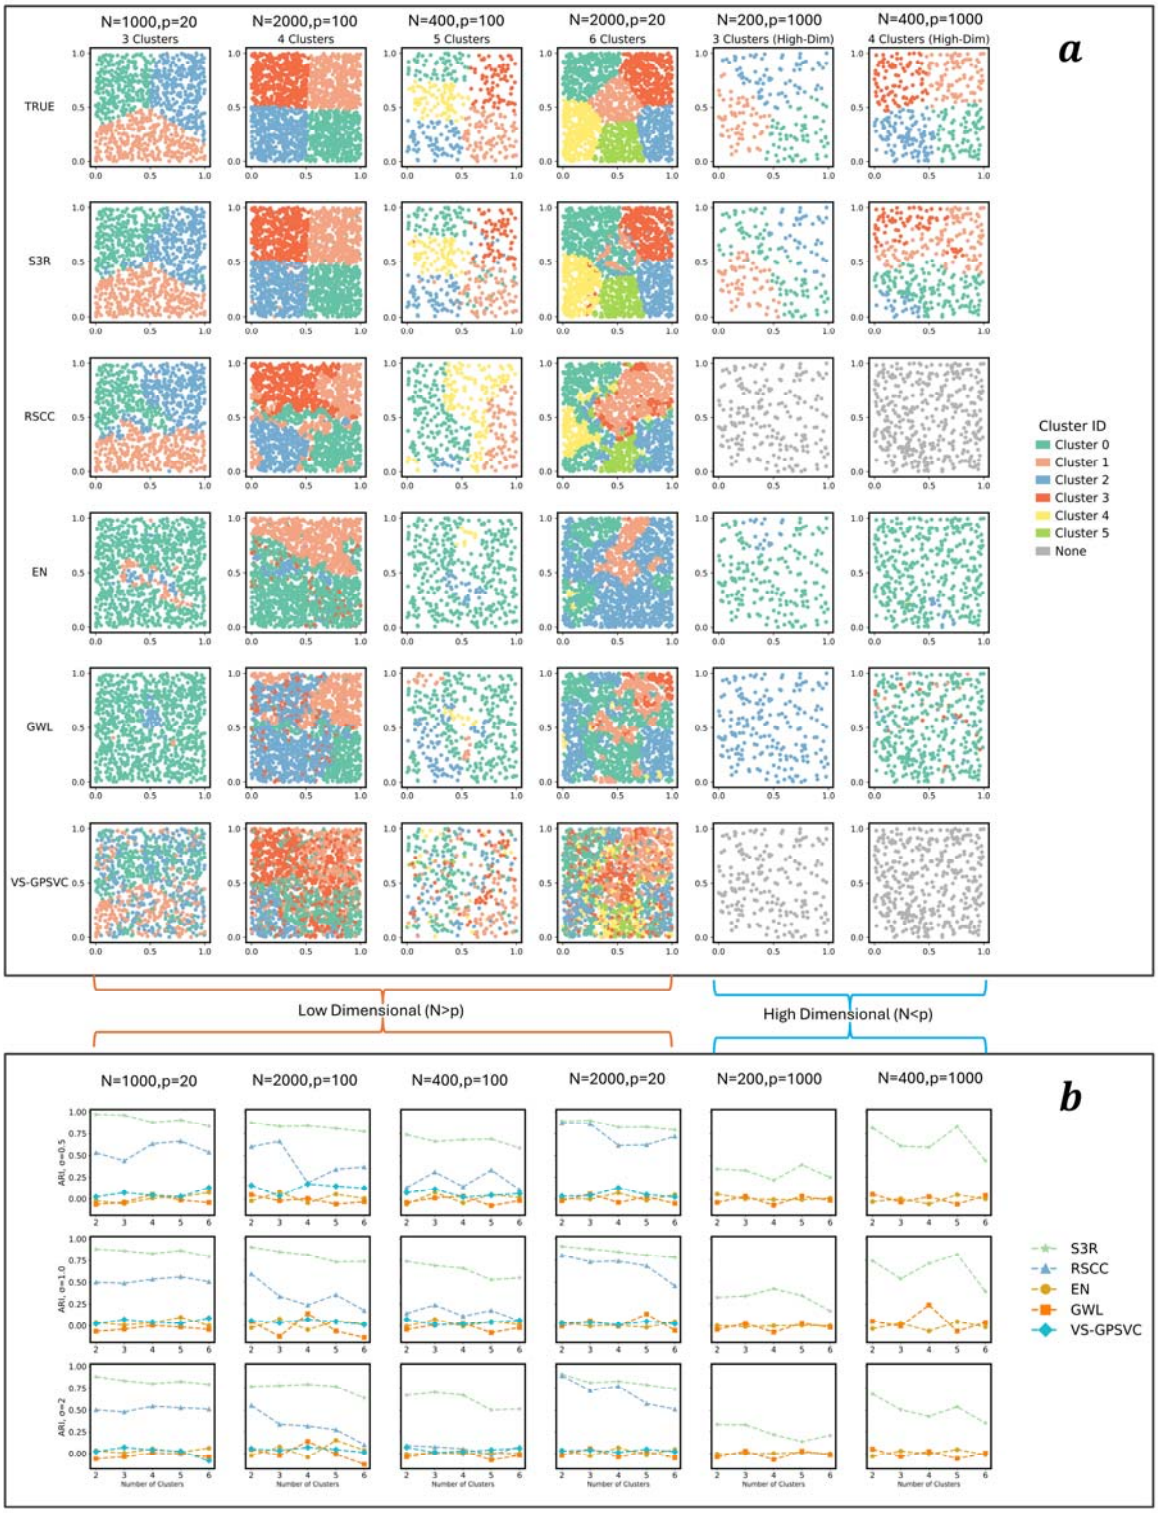

Extended Figure 1| Evaluation of S3R on simulation data. a, A visual comparison of S3R with other benchmarking methods in recovering the coefficient clustering patterns across various simulation scenarios. From left to right are different sample size (N) and feature dimensions (p). From top to bottom are ground-truth coefficient cluster patterns, and estimated coefficient clusters by different methods. x-axis and y-axis are spatial coordinates. Colors of the points indicate different clusters. **b,** Quantitative evaluation of S3R with other competing methods based on Adjusted Rand Index (ARI) of the estimated coefficient clusters with ground truth, under various simulation settings. The dashed line indicates different methods. x-axis shows different numbers of simulated coefficient cluster patterns, and y-axis shows the ARI. From left to right: different sample size (N) and feature dimensions (p).

Application and validation of S3R on recovering true target-TF regulatory relations on a DLPFC dataset

To demonstrate the ability of S3R on resolving spatially varying target–TF associations, we analyzed spatial transcriptomics data collected from human dorsolateral prefrontal cortex (DLPFC) sections from Maynard et al. [19] (**Fig. 2a**). This dataset leverages well-defined laminar architecture and annotated layer-specific gene expression covering all six cortical layers and white matter (WM). We applied S3R to study the spatially varying regulatory relationships for a few selected targeted genes on two DLPFC samples (151673 and 151507; **Fig. 2a**). For each target gene, we modeled its expression as the response and used a curated panel of transcription factors as predictors; we refer to this fixed panel as "regulators" (**Fig. 2b-d**) and reserve "predictors" for the transcriptome-wide analysis below (**Fig. 2b–f**). The candidate transcription factors include all members from three Gene Ontology gene lists, including GO:0003700 (DNA-binding transcription factor activity), GO:0006355 (regulation of transcription, DNA-templated), and GO:0006351 (transcription, DNA-templated) [20]. Regulators or predictors expressed in less than 10% of spots of the corresponding sample are removed prior to analysis.

We focused on three target genes with well-characterized white-matter-enriched expression: MBP, MOBP, and PLP1. After applying S3R, we selected regulators with non-zero location-specific coefficients, and obtained 26/16, 27/20 and 31/20 regulators for MBP, MOBP and PLP1 in the two sections respectively (selection procedure in Methods and Materials). The top four selected regulators for the three target genes in both slides are showcased in **Fig. 2b**, where the arrows connect each target gene to its top selected regulators by S3R among all the ~1,400 candidates, with edge width proportional to the spatial effect size. **Fig 2c** shows the location-specific coefficients $\boldsymbol{\beta_k(s_i)}$ for the three target genes and their top two selected regulators in the two DLPFC sections, and each figure panel corresponds to one target–regulator pair. The red dashed outlines highlight specific cortical layers where strong spatial effects per target were detected and revealed coherent laminar patterns: MBP showed associations with SOX8 and MYRF concentrated in WM across both sections; MOBP showed associations with SOX10 and MYRF prominent in WM with extension into layer 6; and PLP1 showed associations with SOX10 largely confined to WM, and with OLIG2 largely in WM extending to layer 6 (**Fig. 2c**).

For orthogonal validation, we analyzed matched single-cell RNA-seq data from the same cortical regions (**Fig. 2d**). Here, we calculated the Spearman correlation of the target gene's expression and all of the 1,400 candidate regulators' expression in a layer specific manner. It turns out that top regulators selected by S3R show strong correlation with its target gene in the

layer that is consistent with the S3R prediction (**Extended Figure 2a**). **Fig. 2d** provided a layer specific correlation between target gene and its selected top two regulators using the matched single cell RNA-seq data. The red dashed boxes indicate the layers where S3R selected regulators show the highest coefficients for the target, which coincide with the layers demonstrating the highest correlations among the regulator and target. This is consistent for all the three target genes and its top regulators across the two DLPFC samples.

To further demonstrate the S3R's ability to deal with high dimensional feature space, we next extended the analysis to all expressed genes that have non-zero expression in at least 10% of the spots (**Fig. 2e–f**). In total, about 4300 predictors were used for both sections. We here studied three genes as target genes, i.e., MBP, MOBP, and SNAP25. After applying S3R, we identified 29/16, 25/22 and 36/22 number of spatially varying and significantly associated predictors for MBP, MOBP and SNAP25, for the two slides. The estimated coefficients of the top two selected predictors, which associate with the respective response gene in a spatially varying manner, was showcased in **Fig. 2e**. The top predictors recovered by S3R again recapitulated known laminar organization: MBP was associated with PLP1 and MOG predominantly in WM; MOBP was associated with MBP and PLP1 in WM with partial extension into layer 6; and SNAP25 was associated with PCP4 and PRDM8 enriched in layer 5 and layers 3–5, respectively.

Similarly, we validated such layer-specific associations using matched single cell RNA-seq data. Again, the selected top layer specific predictors consistently show higher correlations with the target gene than the other non-selected candidate predictors (**Extended Figure 2b**). Also, layer-wise single-cell correlations peaked in the same layers highlighted by the spatial coefficient maps, corroborating the specificity of these associations estimated by S3R (**Fig. 2f**). Collectively, these results indicate that S3R recovers spatially coherent, biologically plausible target–TF associations from a curated TF panel and scales to transcriptome-wide predictors while preserving fine layer specificity across independent DLPFC sections.

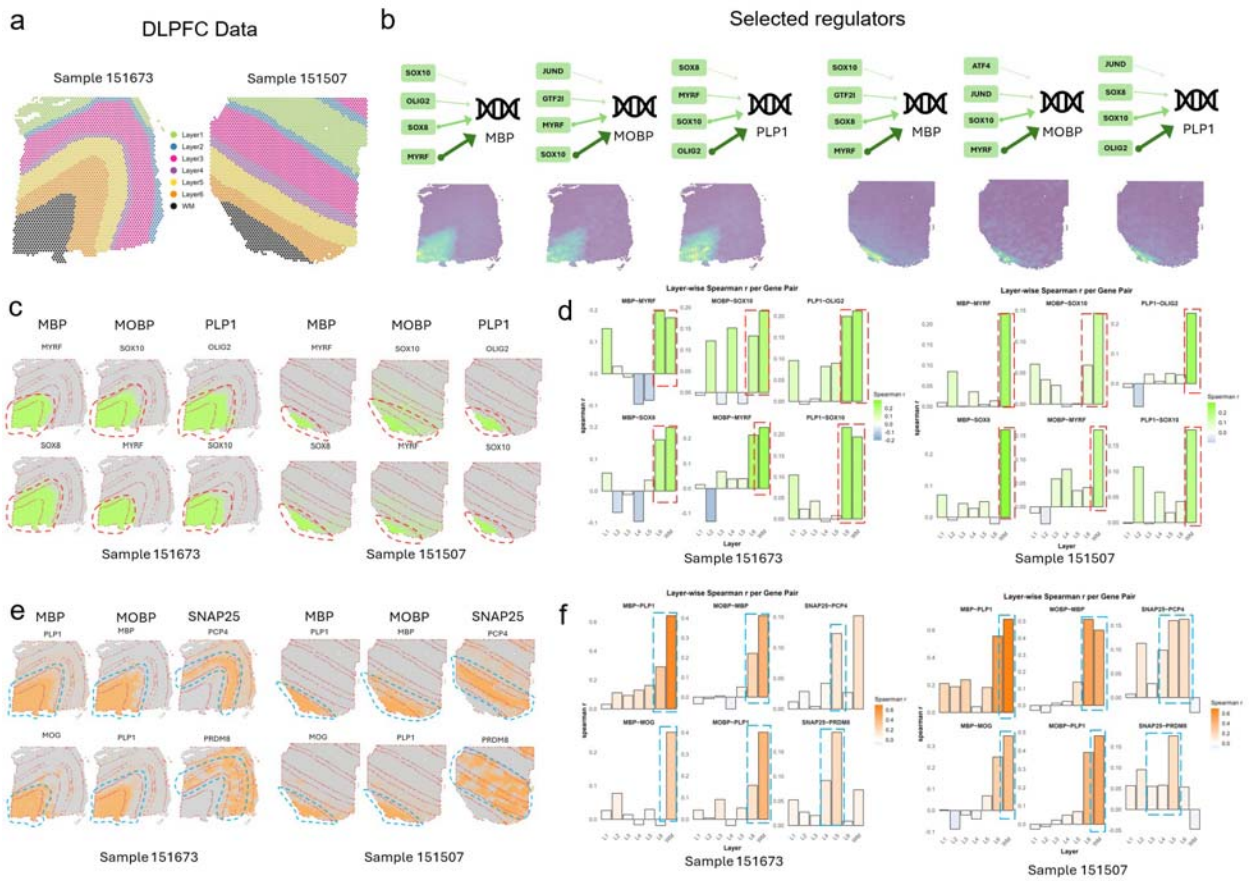

Fig. 2 | Computational evaluation of S3R on a brain DLPFC data.
(a) Two selected DLPFC sections analyzed in this study.
(b) Top: S3R selected top regulators (green box) point to targets with arrow widths proportional to the inferred effect size (coefficient magnitude); Bottom: Spatial expression profiles of the target genes in the two DLPFC sections.
(c) Spatial maps of the estimated coefficients for the top two S3R-selected regulators of three target genes, shown for both sections. Models here were fit using only a candidate set of ~1,400 putative regulators.
(d) Layer-wise correlations between the selected regulators and their targets using matched single-cell RNA-seq (x-axis: cortical layers; y-axis: correlation). Red dashed boxes mark layers with peak correlations.
(e) As in c, but the model was fit using all genes.
(f) As in d, but the model was fit using all genes.

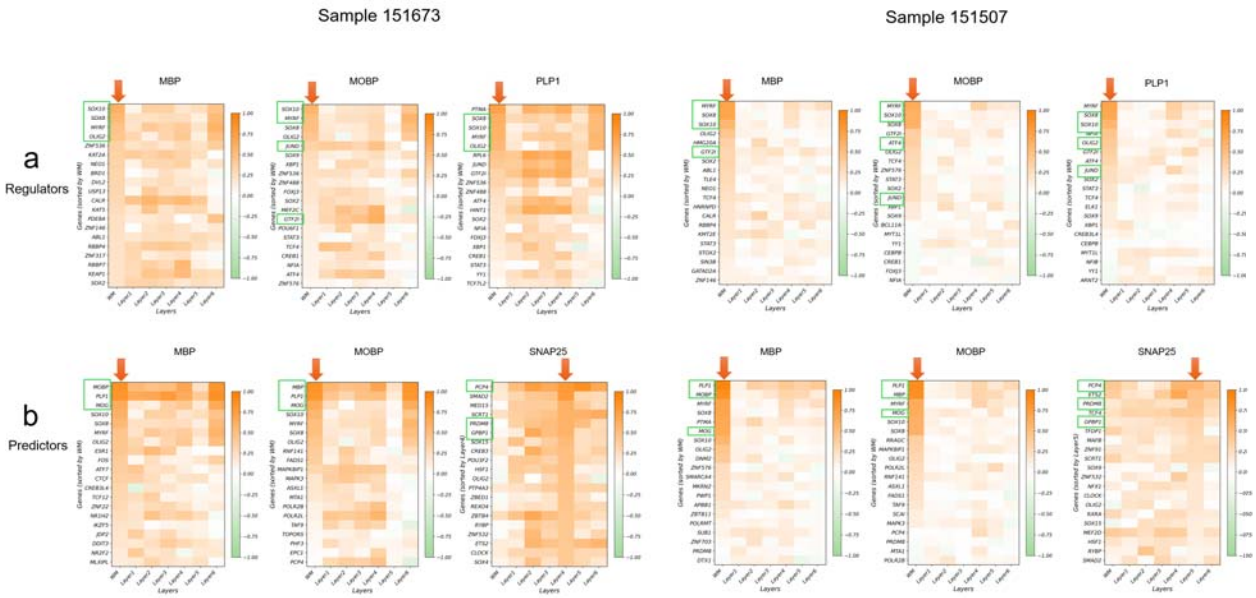

Extended Fig. 2 | Layer-wise Spearman correlations between each target gene and S3R-selected regulators/predictors for the two DLPFC sections, using matched single cell data of DLPFC. a) shows correlations with regulators and b) shows correlations with predictors. Genes highlighted with green boxes are the top candidates used in Fig. 2 (four regulators and three predictors). -axis shows different layers, and -axis shows different genes ranked by the average coefficient magnitude estimated by S3R. This figure provides a quantitative, layer-specific complement to Fig. 2d,f.

Application of S3R to recover cell type–specific expression and spatial gradients

The flexibility of S3R allows distinct spatial patterns to be interrogated by choosing the response and predictors in the regression model appropriately. To uncover cell type–specific spatial expression, we modeled, for each gene in turn, its spot-level expression as the response and the deconvolved cell-type proportions as predictors. This formulation is natural and useful for 10x Visium- or NanoString GeoMx DSP-based spatial transcriptomics data, where each spot/ROI contains multiple cells and the measured expression is a weighted average of constituent cell types. In this formulation, the location-specific coefficients quantify the association between cell type and the expression of the target gene at spatial location , yielding a cell type–attributed coefficient field for each gene. Most spatial deconvolution methods estimate only spot-level mixtures; by regressing gene expression on cell-type proportions within S3R, we recover spatially varying, cell type–attributed expression patterns for all genes. We applied this framework to our in-house 10x Visium skin dataset from acute *H. ducreyi* infection [21] (**Fig. 3**). Genes expressed in at least 10% of spots were retained and analyzed. Predictors consist of RCTD-estimated spot-level proportions [22] for eleven cell types curated from matched single-cell data: keratinocytes, fibroblasts, melanocytes, smooth-muscle cells, endothelial cells, T/NK-like cells, macrophages, dendritic cells (myeloid and plasmacytoid combined to reduce collinearity), ISG-high cells, mast cells, and B cells (**Fig. 3a**). Exact thresholds and regularization parameters are provided in Methods.

As an initial validation, we examined two genes with well-established cell type specificity: KRT14, a structural keratin predominantly expressed by keratinocytes, and HLA-DRB1, a canonical antigen-presentation gene enriched in antigen-presenting cells. S3R recovered spatially coherent coefficient fields that attributed KRT14 to keratinocytes and HLA-DRB1 to macrophage and dendritic-cell compartments, with boundaries that mirrored histological features (**Fig. 3b**). A cell type was considered attributing to a gene's expression when its coefficient field exhibited significantly non-zero values at certain spatial regions after spatial regularization (Methods). A downstream contribution analysis that ranks predictors by marginal effect confirmed keratinocytes as the dominant source attributed for KRT14 and macrophages/dendritic cells for HLA-DRB1 (**Fig. 3c**). These associations were corroborated in matched single-cell RNA-seq profiles, where expression localized to the expected cell populations (**Fig. 3d**).

A direct application of the resulting spatial and cell type specific expression from S3R is to examine the spatially varying expression gradients across different cell types. We next asked whether S3R resolves spatial gradient organization by using VEGFA as an example. We highlight VEGFA because it anchors angiogenic remodeling in acute cutaneous infection, is robustly expressed at Visium resolution, has cell type–dependent sources in skin, and thus provides a test case of S3R's ability to disentangle cell mixing and recover spatial gradients. In **Fig. 3e**, the red dashed-line in the H&E staining image marks the epidermal-dermal axis with epidermis and dermis towards the bottom/top of the image (left panel **Fig. 3e**), also called as normalized vertical coordinate. S3R attributed a prominent coefficient field to keratinocytes and fibroblasts, that varied along the dermal–epidermal axis. Fibroblast-attributed VEGFA coefficients decreased toward deeper dermis and is the highest at the epidermis level (**Fig. 3e**). The fibroblast-VEGFA coefficient showed a strong monotonic trend with vertical coordinate ($R^2 = 0.55$, $P < 1 \times 10^{-1}$□; **Fig. 3f**), whereas the raw spot-level VEGFA signal exhibited only a negligible association with the normalized vertical coordinate ($R^2 = 0.0047$, $P = 1.7 \times 10^{-}$□), illustrating how cell mixing can dilute gradients in aggregate signals.

Finally, we compared spatially varying genes (SVGs) identified by S3R with three widely used spot-level methods, namely Moran's I [23], nnSVG [23, 24], and SPARK [23, 25], on the same section (**Fig. 3g**). S3R called SVGs based on the spatial and cell type specific expression out from S3R, combined with Moran's I spatial autocorrelation (details in Methods and Materials). When applied directly to spot-level expression, cross-method concordance was modest: S3R, Moran's I, nnSVG, and SPARK called 546, 118, 994, and 489 SVGs, respectively, with only partial overlap (**Fig. 3g**, top). Overall, the consistency of the SVGs across different methods were poor; however, S3R detected SVGs overlap better with other methods (Top panel of **Fig. 3g**). For example, majority of SVGs called by nnSVG overlap with those by S3R; half SVGs detected by Moran's I and Spark overlap with those by S3R. Because heterogeneous cell mixtures can drive method-specific calls, we reapplied nnSVG to the S3R-derived cell type–attributed expression matrices. Here, we didn't apply SPARK as it only accepts count-based data as input. Removing mixture confounding in this way substantially increased concordance, yielding a larger shared set among the two approaches (**Fig. 3g**, bottom). As an example, S3R identified VEGFA attributed to fibroblasts and CXCL8 attributed to macrophages as SVGs. This means that fibroblast expressed VEGFA and macrophage expressed CXCL8 display spatial

gradients, however, the spot-level expressions of VEGFA and CXCL8 didn't display spatial gradients, and are missed by the spot-level SVG methods (**Fig. 3h**). This application illustrated how cell type–resolved modeling reveals patterns that are diluted in aggregate signals.

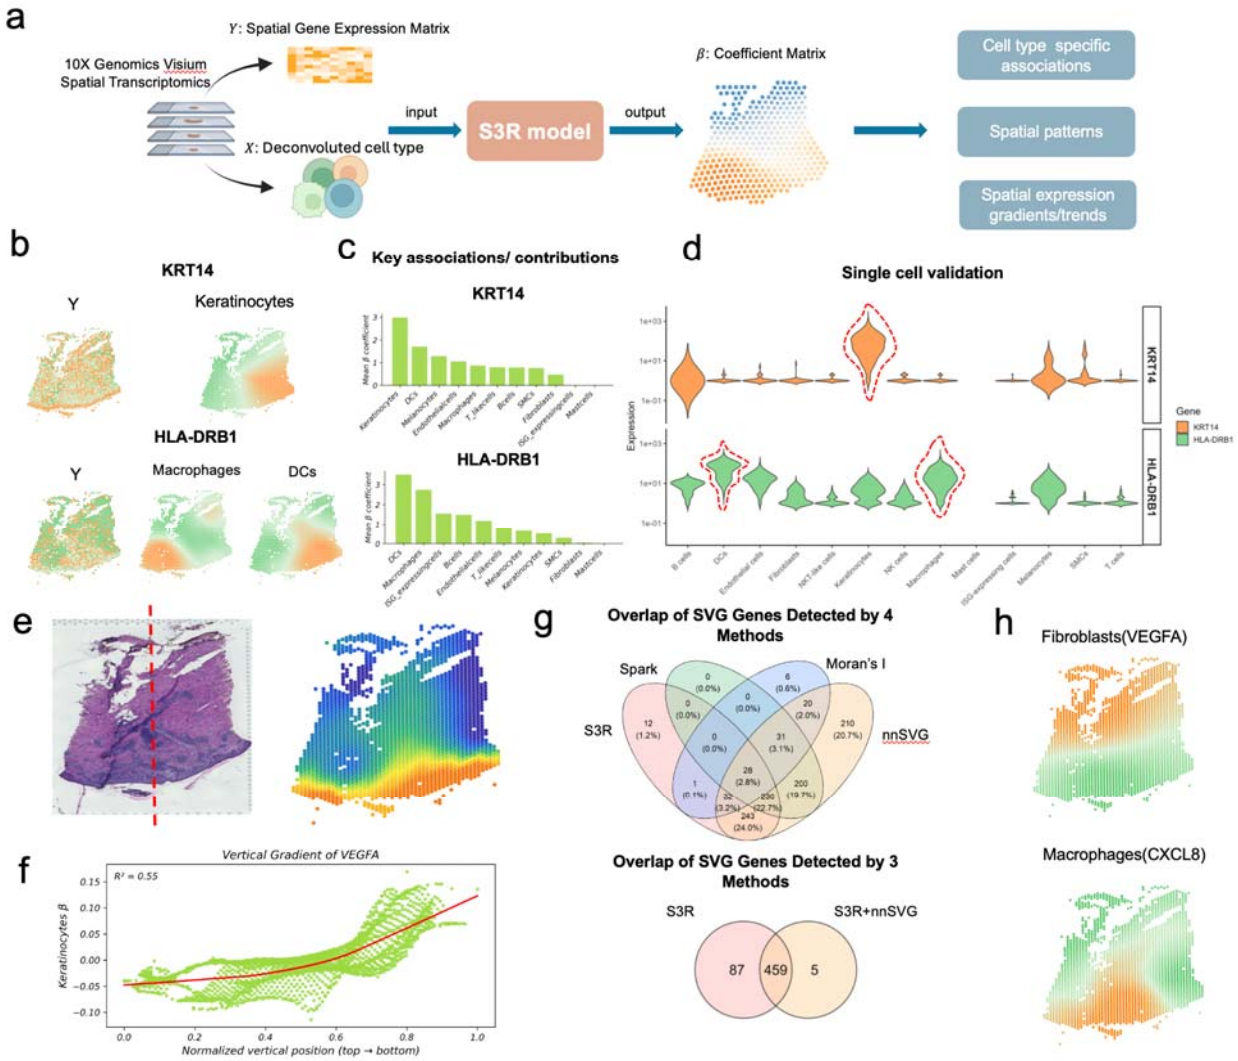

Fig. 3 | Application of S3R to spatial transcriptomics data from acute *H Ducreyi*-infected skin samples.

(a) Inputs to S3R comprise the spot-level expression of a target gene (response) and spot-level cell-type proportions deconvolved with RCTD (predictors); the model outputs a spot-by–cell type coefficient matrix that approximates cell type–attributed, spatially varying expression for the target.

(b) Spot-level expression (left) and S3R-estimated cell type–attributed expression maps (middle/right) for KRT14 (top) and HLA-DRB1 (bottom).

(c) Mean cell type–attributed expression for KRT14 (top) and HLA-DRB1 (bottom) across spots (x-axis: cell types; y-axis: mean attributed expression).
(d) Distributions of KRT14 (top) and HLA-DRB1 (bottom) expression across cell types in matched single-cell RNA-seq.
(e) H&E section of the slide (left) and the estimated spot specific expression of VEGFA by keratinocytes (right). The red dashed line indicates a normalized vertical coordinate perpendicular to the epidermis (bottom) and dermis (top) junction.
(f) Relationship between normalized vertical position (x-axis; epidermis to dermis) and fibroblast-attributed VEGFA coefficients (y-axis).
(g) Overlap of spatially varying genes (SVGs) called by S3R, SPARK, Moran's I, and nnSVG using spot-level expression (top) and using S3R-derived cell type–attributed expression (bottom).
(h) Examples of gradients detected by S3R but missed by spot-level methods: fibroblasts–VEGFA (top) and macrophages–CXCL8 (bottom).

Application of S3R to recover cross–cell type, cross-gene spatial four-way co-variation.
Aside from recovering the cell type specific spatial expression gradients shown above, we here demonstrate another important utility where we used S3R to infer cell type–attributed, spatially varying gene fields and then quantified cross–cell type, cross-gene co-variation at spot resolution. For each gene *g*, we again modeled its spot-level expression as the response and RCTD-deconvolved cell-type proportions as predictors, yielding location-specific coefficients $\boldsymbol{\beta}_{k,g}(\boldsymbol{s}_i)$ that we interpret as the cell type–attributed expression profile for gene $\mathbf{g}$ in cell type $\mathbf{k}$ at spot $\mathbf{s}_\mathbf{i}$ (**Fig. 4a**). Stacking these coefficient fields across genes produces, for each cell type, a spots-by-genes matrix with many entries shrunk to zero by S3R's sparsity penalties; additional post hoc filtering set small coefficients to zero when not significantly different from zero (see Methods). We applied this framework to our in-house 10x Visium sections from seven pancreatic ductal adenocarcinoma (PDAC) patients, which include ductal epithelium and fibroblast subtypes (permCAF or myCAF and restCAF or iCAF), immune populations, and endothelium [26].
To quantify cross–cell type gene co-variation, we considered any pair of cell types ($\mathbf{C}_1, \boldsymbol{C}_2$) and any pair of genes ($\boldsymbol{G}_1, \boldsymbol{G}_2$) and computed the Spearman correlation across spots between $\boldsymbol{\beta}_{C_1,G_1}(\mathbf{s})$ and $\boldsymbol{\beta}_{C_2,G_2}(\mathbf{s})$. Because the $\boldsymbol{\beta}$-fields are spatially smoothed, simple rank correlation approximated more elaborate spatial autocorrelation measures in practice; nonetheless, all significance assessments used sample-wise permutation controls and false-discovery-rate correction, and we retained only pairs that exceeded an effect-size threshold and replicated across sections (see Methods). The resulting quantities define a four-way tensor over genes and cell types that summarizes cross–cell type co-variation at tissue scale.
In **Fig. 4b-c**, we illustrated examples that highlight distinct spatial couplings that are consistent in two selected sections. A ductal–permCAF pairing showed coordinated spatial variation between ductal PKM and permCAF CD44, consistent with metabolic remodeling in the tumor epithelium aligned with CAF-mediated adhesion and matrix engagement (Left, **Fig. 4b-c**). A macrophage–permCAF pairing exhibited co-variation between macrophage SDC2 and permCAF CXCL8, in line with chemokine-rich stromal niches (Middle, **Fig. 4b-c**). And a

restCAF–macrophage pairing displayed endomucin (EMCN) and LAMTOR4 co-variation suggestive of vascular–immune remodeling in CAF-adjacent regions (Right, **Fig. 4b-c**). In each case, the S3R coefficient maps (top row within each figure panel) localized a subset of the high-expression spots seen in the raw expression maps (bottom row within each figure panel), as expected when spot-level signals arise from mixtures and only a fraction of spots exhibit the gene in the cell type of interest. Spatial distributions of the corresponding cell types are provided in **Extended Fig. 4**.

We summarized cross–cell type coupling across patients by binarizing section-level associations into seven gene–gene–cell–cell tensors (1 = significant, 0 = not significant under the replicate-wise thresholds) for the seven sections and then aggregating across seven sections. The resulting average interaction strengths, displayed as a bubble plot (**Fig. 4d**), indicated prominent coupling between ductal cells and permCAFs, followed by restCAFs, monocytes/macrophages, and endothelium; macrophages coupled most strongly with ductal cells and permCAFs; and B cells showed their strongest coupling with T cells. We catalogued the top, consistently replicated gene pairs for key cell-type pairs (**Fig. 4e**). Pathway enrichment on genes participating in these cross–cell type couplings revealed kinase signaling, receptor-mediated endocytosis, and macromolecule biosynthesis in ductal–permCAF pairs (**Fig. 4f**, left), consistent with reports that bidirectional malignant cell–fibroblast crosstalk programs myCAF molecular/functional heterogeneity and promotes metastasis [27], that macropinocytosis sustains the myCAF state under glutamine stress to support ECM deposition and PDAC progression [28], and that TGF-β–driven myCAF differentiation rewires CAF metabolism toward glycolysis to supply tumor cells with metabolic intermediates [29-31]. Enrichment for cell migration/adhesion and cytokine production pathways in myCAF–macrophage pairs aligns with studies showing CAF programs that recruit and polarize myeloid cells and engage macrophage–CAF chemokine crosstalk [32, 33] (**Fig. 4f**, middle). Finally, enrichment for intracellular protein transport and signal transduction in B–T cell pairs is consistent with coordinated B–T activity within tumor-associated tertiary lymphoid structures linked to antitumor immunity and improved outcomes (**Fig. 4f**, right).

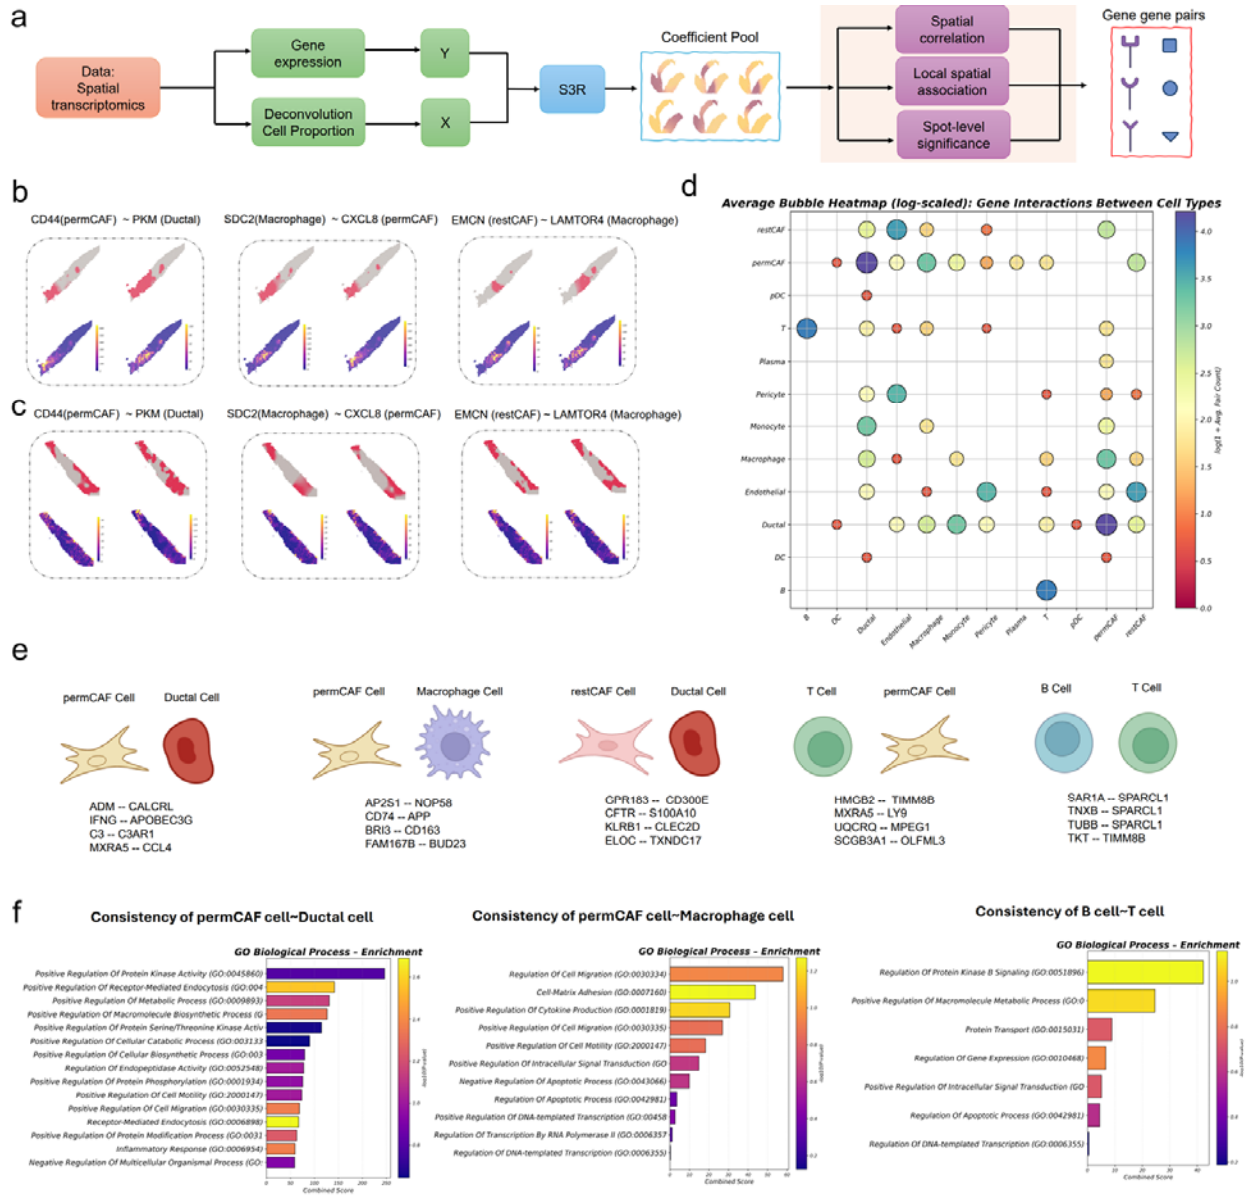

Fig. 4 | Application of S3R to PDAC spatial transcriptomics data.

(a) Inputs to S3R comprise the spot-level expression of a target gene (response) and spot-level cell-type proportions deconvolved with RCTD (predictors); the model outputs a spot-by–cell type coefficient matrix for each gene.

(b) Coefficient maps (top) for cell type pairs exhibiting co-variation between gene-gene pairs, and original spot-level expression of the same gene pairs (bottom), for three four-way associations in section 1.

(c) Same as (b), but in section 2.

(d) Bubble heatmap showing average interaction frequency/strength between cell-type pairs across seven PDAC sections. Bubble size and color represent interaction strength.

(e) Representative top gene pairs for key cell type pairs (e.g., permCAFs–ductal, permCAFs–macrophages, B–T cells).

(f) GO biological-process terms enriched by interacting genes for key cell type pairs.

Discussion

Here, we propose a spatially smooth and sparse regression model to fill a gap in spatial transcriptomics data analysis by estimating location-specific associations between a response feature and a high-dimensional set of candidate predictors while enforcing parsimony and spatial coherence. Conceptually, S3R treats "what relates to what" as a spatially varying problem: coefficients are learned as fields over the tissue, borrowing strength within neighborhoods yet permitting sharp transitions at histological boundaries. Methodologically, the combination of sparsity ($\boldsymbol{L_1}$ and group-$\boldsymbol{L_2}$) with a minimum-spanning-tree–guided smoothness term yields interpretable models that remain stable in the presence of noise, missing data and collinearity, and computationally tractable at the scales typical of Visium and related platforms. A multi-GPU implementation and parallel hyperparameter tuning enable efficient training on large tissues and large predictor sets.

A central strength of S3R is its ability to operate in high-dimensional regimes without having to collapse to pre-screened, low-rank surrogates. In simulations, S3R recovered spatially varying coefficients, achieved accurate variable selection, and reproduced known spatial partitions. S3R's regression formulation also makes it highly flexible to address distinct questions, because responses and predictors are user-defined. With a target gene expression as the response and TF expression as predictors, S3R estimates spatially varying target–TF associations. With cell-type abundance as the response and genome-wide genes (or gene-set scores) as predictors, S3R can characterize location-dependent enrichment or depletion of cell types. With gene expression as the response and deconvolved cell-type proportions as predictors such as in Visium platform, S3R yields cell type–attributed expression fields for every gene, enabling gradient analyses and improved detection of spatially varying genes once cell mixing is mitigated. In addition, by correlating S3R coefficient fields across cell types and genes, one can construct cross–cell type, cross-gene co-variation tensors that prioritize candidate interactions for follow-up. This analysis complements existing computational approaches to cell–cell communication, where single-cell RNA-seq loses spatial coordinates and therefore cannot directly assess proximity-dependent effects [34-36]**,** and spot-level spatial methods often lack cell-type specificity when mixtures are present. Operating on S3R's cell type–attributed coefficient fields resolves both issues, by preserving spatial localization and attributing signal to specific cell types and genes. Importantly, it is not restricted to ligand–receptor pairs but to alternative mechanisms encompassing all possible genes.

In summary, S3R provides a general, scalable approach for modeling spatially varying relationships in high-dimensional spatial transcriptomics data. By uniting structured sparsity with adaptive spatial smoothness and by allowing analysts to choose biologically meaningful responses and predictors, S3R examines how spatial context shapes associations. We anticipate S3R will serve as a practical tool for a broad class of spatial analyses and as a bridge from descriptive maps to mechanistic investigation of tissue microenvironments. Despite these promising results, some limitations remain. S3R estimates associations, not causal effects; interpreting coefficient fields mechanistically requires orthogonal evidence such as spatial

proteomics and perturbations experiments. The accuracy of cell type–attributed fields depends on the quality of deconvolution, and uncertainty propagation from deconvolution to S3R would make downstream inferences more rigorous. Like any smoothing approach, there is a bias–sharpness trade-off. While the MST-guided term preserves abrupt changes, automatic boundary detection and adaptive, data-driven neighborhoods could further improve edge preservation. Finally, formal inference on the coefficient fields and integration with additional modalities are represent future extensions.

Methods

Overview of the S3R model

A key question in spatial transcriptomics is how to model relationships among spatial features. In particular, we ask which features, such as gene expression or cell-type abundance, associate with a spatial phenotype or another gene's expression, and how those effects vary across a tissue. Traditional regression typically ignores spatial structure, whereas classical spatial models were not designed for the high dimensionality of transcriptomic predictors, leading to a methodological gap at the intersection of spatial statistics and high-dimensional regression [1, 37, 38]. One line of work addresses spatial homogeneity by allowing regression coefficients to vary over space, treating them as smooth functions. Spatially varying-coefficient models often use Gaussian-process fields or conditional autoregressive priors to encourage spatial coherence, and related non-Bayesian approaches include geographically weighted regression [37, 39-41]. These models capture gradual spatial trends in covariate effects, for example, a gene's influence from a tumor core to its periphery, but purely smooth formulations have important limitations for ST analysis. In practice, tissues contain discrete microenvironments where effects shift abruptly; modeling such effects typically requires sparsity- and edge-preserving which smooth-only models do not readily provide.

Another class of methods addresses spatial heterogeneity by imposing penalized regression frameworks that combine sparsity and spatial structure. A prominent example is the fused lasso, which adds an L_1 penalty not only on the regression coefficients (as in the ordinary lasso) but also on differences between coefficients in neighboring spatial locations [42]. This encourages solutions that are sparse and also locally constant in space [42]. To perform variable selection in high dimensions, the group lasso is a useful regularization approach that selects or drops predefined groups of features as a whole [43]. In a spatial context, one can define each gene's coefficient profile across all locations as a group. A group-lasso penalty then either retains that gene by allowing its coefficients to be nonzero at some locations, or sets the entire coefficient surface to zero, thereby performing gene selection. Recent methods in high-dimensional spatial regression have begun to integrate these ideas. Notably, Li and Sang (2019) proposed a spatial homogeneity pursuit method (SCC) that encourages spatial similarity in regression coefficients between locations connected by an Minimum Spanning Tree, to automatically detect cluster patterns and change points in coefficients [44]. Building on this concept, Zhong et al. (2023) introduced the Regularized Spatially Clustered Coefficient (RSCC) model to address large spatial datasets with many predictors [45]. RSCC combines a MST-guided graph penalty with a group lasso penalty to simultaneously achieve variable selection and spatially clustered coefficient patterns [45]. These developments underscore the

value of penalties that accommodate abrupt jumps in coefficients while selecting relevant variables. However, their primary focus is on partitioning the space into homogeneous clusters; fully leveraging continuous gradients in the estimated coefficients, such as in the case of disease microenvironment of spatial transcriptomics data, remains challenging. Moreover, while the group-lasso component in RSCC promotes group-level interpretability but does not yield location-specific sparsity, that is, once a predictor (or group) is selected by RSCC, it tends to be retained across all locations. An individualized and hence location-specific sparsity would further enhance biological interpretability, but the resulting computational burden demands scalable optimization. Overall, there remains a need for methods that adaptively balance smooth and discontinuous spatial variation in covariate effects, enable spot-level variable selection, and remain computationally feasible in high dimensions. Here we present Spatially Smooth and Sparse Regression (S3R) to address these challenges with a new regularization scheme that *jointly* penalizes redundant coefficient values and large differences between neighboring spatial coefficients.

The overall model framework of S3R

Suppose a set of high dimensional spatial variables $(X(s_i), y(s_i), i = 1, \dots, n)$, are observed at locations $s_1, \dots, s_n \in \mathbb{R}^2$, where the response variable $y(s_i)$ is assumed to be spatially correlated, and $X(s_i) = \left(x_1(s_i), \dots, x_p(s_i)\right)^T$ is the *p*-dimensional vector of explanatory variables for the observation located at s_i. Consider a spatial regression model with spatially varying regression coefficients

$$y(s_i) = \sum_{k=1}^{p} x_k(s_i)\beta_k(s_i) + \epsilon(s_i) \qquad (1)$$

In our spatial transcriptomics dataset, we could observe only one sample at each location, making the model (1) ill-posed if without any restrictions on $\beta_k(s_i)$. For example, with only one spatial realization $(X(s_i), y(s_i), i = 1, \dots, n)$ at *n* observed locations, the regression model (1) can be written in the matrix form

$$y = \boldsymbol{X}\beta + \epsilon.$$

Here, we stack all the coefficients into a vector $\beta \in R^{pn}, \beta = \left(\beta_1(s_1), \dots, \beta_1(s_n), \dots, \beta_p(s_1), \dots, \beta_p(s_n)\right)^T$ and the design matrix $\boldsymbol{X} = [diag(\boldsymbol{x}_1), \dots, diag(\boldsymbol{x}_p)]$ is an $n \times np$ matrix with $\boldsymbol{X}_k = \left(x_k(s_1), \ , x_k(s_n)\right)^T$. Clearly, this regression problem needs to be regularized since there are more variables than observations. For spatial problems, it is expected that association between a response variable and explanatory variables at nearby locations is highly likely to be homogeneous. In addition, the number of candidate explanatory variables can be large, and we need to select the ones that are truly associated with the response variable. This motivates us to assign a few regularization functions for $\boldsymbol{\beta}$ reflecting such spatial homogeneity and feature sparsity patterns.

Specifically, we propose to estimate $\boldsymbol{\beta}$ by minimizing the following objective function

$$\frac{1}{n}\sum_{i=1}^{n}\left(y(s_i) - \sum_{k=1}^{p} x_k(s_i)\beta_k(s_i)\right)^2 + \lambda_1 \sum_{k=1}^{p}\sum_{\{(i,j)\in E\}}\left\|\beta_k(s_i) - \beta_k(s_j)\right\|_1 + \lambda_2 \sum_{k=1}^{p}\|\beta_k\|_1 + \lambda_3 \sum_{k=1}^{p}\sqrt{n}\|\beta_k\|_2 \qquad (2)$$

Here, $\boldsymbol{\beta}_k = \left(\beta_k(s_1), \ldots, \beta_k(s_n)\right)^T$, k = 1, ..., p, and E is the edge set of a graph consisting of n vertices, where each vertex corresponds to one observed spatial location, and each edge connects two connected locations. $\lambda_1, \lambda_2, \lambda_3$ are tuning parameters determining the strength of penalization of spatial smoothness, individual sparsity and group sparsity. Specifically, the term associated with λ_1 is to encourage homogeneity between two regression coefficients if their corresponding locations s_i and s_j are connected by an edge in $\boldsymbol{E}$. The term associated with λ_2 induces element-wise sparsity such that each location may have a different non-zero coefficients; λ_3 promotes group sparsity at the predictor level such that one predictor may be removed globally. For simplicity, we reformulate (2) as follows:

$$\frac{1}{n}\sum_{i=1}^{n}\left(y(s_i) - \sum_{k=1}^{p} x_k(s_i)\beta_k(s_i)\right)^2 + \lambda_1 \sum_{k=1}^{p}\|\boldsymbol{H}\beta_k\|_1 + \lambda_2 \sum_{k=1}^{p}\|\beta_k\|_1 + \lambda_3 \sum_{k=1}^{p}\sqrt{n}\|\beta_k\|_2 \quad (3)$$

where $\boldsymbol{H}$ is an $m \times n$ matrix constructed from the edge set E with m edges, similar to [44]. For an edge connecting two locations s_i and s_j, we represent the penalty term $\left\|\beta_k(s_i) - \beta_k(s_j)\right\|_1$ as$\|H_m\beta_k\|_1$, where $\boldsymbol{H}_m$ is a row vector of $\boldsymbol{H}$ and contains only two nonzero elements, 1 at the i th index and −1 at the j th.

The spatial smoothness control

To construct the edge set $\boldsymbol{E}$, for most methods, the graphs are constructed by *k*-nearest neighbors or neighbors within a certain radius such that all points are connected, in which case, the number of edges is greater than the number of nodes. While many algorithms have been developed to solve the generalized lasso problem, for an arbitrary $\boldsymbol{H}$ with $m > n$ for large m and n, these algorithms are computationally very costly. An appropriate choice for E that balances model accuracy and computational efficiency should only include coordinate pairs close to each other, should lead to connectivity of all data points, and should have no redundant pairs. One choice of E that satisfies all three of these criteria is the edge set of a Minimum Spanning Tree (MST), similar to SCC and RSCC [44, 45].

Given an undirected graph $G = (V, E_0)$ with a weight function $d(e)$ that assigns a weight to each edge e in an edge set E_0, an MST is defined as the subgraph $T = (V, E), E \subseteq E_0$ that connects all vertices without any cycles and minimizes $\sum_{e \in E} d(e)$.It is known that an MST has $|V|$ vertices and $|V - 1|$ edges. It is known that for any cut of a given connected graph G, the minimum-weight edge that crosses the cut is in the MST for G. In our implementation, taking a given set of n spatial locations as vertices and the Euclidean distances between locations as the edge weights, we by default construct a *k*-nearest neighbors (KNN) graph (e.g., $k = 7$) for small-to-moderate datasets ($n \leq 1000$), while allowing users to choose either KNN or an MST. For larger datasets ($n > 1000$), we by default construct a Minimum Spanning Tree (MST), $T = (V, E)$, consisting of $n - 1$ edges with all n locations connected. By the construction, T provides a connected, acyclic graph that compactly represents the spatial topology of the n observed points and is thus an appropriate choice satisfying the aforementioned desired properties for edge set selections. In this case, $\boldsymbol{H}$ is a full row rank matrix encoding the $n - 1$ penalty terms on the coefficients β_k .

The algorithm implementation

We optimized the algorithm using Adam [16], combined with cosine/step-wise learning rate adjustments, gradient clipping, and early stopping to ensure stable convergence during training. The entire implementation uses batch mode and fully utilizes GPU acceleration, allowing the model to run on large datasets. Specifically, to allow efficient use of GPU: 1). When computing the first loss term, we performed element-by-element multiplications and then sum the results by feature, avoiding the time-consuming computation of large matrices; 2). When computing the graph-fused lasso loss, we pre-prepare the edge indices and perform the computation on the GPU using an aggregation/difference operation, which results in a computational complexity of $O(mp)$, linear in the number of edges m and the number of features p; 3). When computing the group sparsity term, we rely on PyTorch autograd's subgradient implementation and add a small ϵ stabilization term to prevent instability near zero.

Model and hyperparameter selection

Our model has three hyperparameters: λ_1 that controls the spatial smoothness of the regression coefficients; λ_2 that: controls sparsity of individual parameters; and λ_3 that controls group-level sparsity. These hyperparameters are selected by cross validation. Instead of randomly splitting the spots into training and testing sets, we divided the spatial region into several spatially coherent blocks, using some blocks for training and the rest for validation. This ensures that the model can generalize in space. Because we need to tune $\lambda_1, \lambda_2, \lambda_3$ together, the number of possible combinations is very large. If we use brute-force grid search or pure random search, the computation would be too expensive. To solve this, we use Optuna [18]., an automatic hyperparameter optimization tool. It searches hyperparameter combinations in a random + intelligent way within given ranges. Specifically, $\lambda_1, \lambda_2, \lambda_3$ are sampled on a log scale, namely, $loguniform(1e-6, 1e2))$. In Optuna, each attempt of one hyperparameter combination is called a trial. We start many trials at the same time, each trial corresponds to a different hyperparameter set. These trials are assigned to different GPUs to run in parallel, so we can test many combinations at the same time and make the search much faster.

Post-hoc significance evaluation of the coefficients

Inference of high dimensional regression is challenging, and even more in our case with three different levels of penalty functions. We implemented a post-hoc variable selection procedure to further select informative predictors with significantly non-zero coefficients. We borrow the idea in linear regression with lasso, where Lasso will be run first, and then the retained variables, much less than the original set, will be fit into the linear regression model, and inference of coefficients could be conducted and those non-significantly different from zero will be removed. Here, to achieve spatially varying variable selection, we conduct a model refitting for each spot using linear regression by using the spot and its neighbouring spots, by assuming that the neighboring spots share the same coefficient scales. To speed up the computation, in practice, we dissect the spatial map into multiple clusters for the linear regression model fitting. A coefficient was considered significant for the spot, if the associated p-value for the predictor in the linear regression is below the threshold (adjusted $p < 0.05$).

Deducing cell type specific and spot-level expressions

We applied RCTD for spatial deconvolution [22]. In terms of single cell expression reference and annotation required by RCTD, for the skin infection dataset, we utilized matching single cell RNA-seq data from the same study [21]; for the PDAC dataset, we utilized matching single cell RNA-seq data from the same group [26]. By regressing a gene's expression on the cell type proportions using S3R, we deduced cell type–specific and spot-level expression. The resulting coefficient field $\beta_{g,c}(s_i)$ quantifies how strongly a given cell type explains the expression of the target gene at each spatial location. Here, we applied a second post-hoc parameter shrinkage additional to the procedure mentioned above: 1) only non-negative coefficients are retained and all negative-coefficients are shrunken to zero; 2) positive coefficients are only allowed in spots where the cell type has non-zero abundance. This two-step selection procedure is implemented because the cell types could only contribute to a gene's expression positively; and if the cell type is not present, it should not contribute to the expression of the genes. In this way, the coefficients provide two layers of information: (i) cell type specificity, because predictors correspond to deconvolved cell-type proportions, and (ii) spot-level resolution, because coefficients vary across tissue locations. The resulting cell type by spots coefficient matrices could be utilized for downstream analysis such as spatially varying genes detection, and cross-cell and cross-gene interaction.

Detecting spatially varying genes (SVG)

S3R: **S3R** adopts a multi-step procedure in detecting spatially variable genes. Based on S3R output of the coefficients by regressing a gene's expression on the cell type proportions, we first conduct a post-hoc variable selection for all the spots, which will further shrink many coefficients to be zero. For each gene, we looked at its spot by cell type coefficient matrix. Then we used Moran's I to test spatial variation. In other words, for each gene–cell type pair, we checked if its coefficients showed spatial structure.

Moran's I: For each gene, Moran's I is computed on expression profile across spots, and significance is assessed by permutation. A significantly positive Moran's I indicates the presence of non-random spatial structure.

nnSVG: nnSVG (**n**earest-**n**eighbor **S**patially **V**arying **G**enes) models gene expression using a nearest-neighbor Gaussian process approximation. Expression is decomposed into non-spatial and spatially structured components. A gene is designated as an SVG if the estimated spatial variance is significantly greater than zero. In our analysis, nnSVG was implemented using the nnSVG R package, with expression counts and spatial coordinates encapsulated in a SpatialExperiment object. We used default settings with multi-threading enabled (n_threads = detectCores() – 1), and retained genes with adjusted $p < 0.05$ as SVGs.

Spark: Spark (**S**patial **PA**ttern **R**ecognition via **K**ernels) employs generalized linear spatial models with multiple kernel functions to capture diverse forms of spatial dependence, including local hotspots, broad gradients, and periodic structures. For each gene, kernel-specific spatial effects are jointly tested, and p-values are combined to robustly detect spatially variable expression across multiple spatial scales. We applied SPARK using the official R implementation, providing raw count matrices and spot coordinates as input. Variance components were estimated with spark.vc and significance was assessed with spark.test under default settings. Genes with adjusted $p < 0.05$ were called as SVGs.

Datasets used.

Both publicly available and in-house generated spatial transcriptomics datasets were used for benchmarking and case studies in this work. DLPFC dataset was obtained publicly from Maynard et al [19]; the H. Ducreyi infected skin dataset was an in-house dataset previously reported at Brothwell et al [21]; the UNC PDAC dataset was another in-house dataset previously reported at [26].

Computational resources

All computations were performed on the Big Red 200 supercomputer at Indiana University, an HPE Cray EX system equipped with GPU-accelerated nodes (each with four NVIDIA A100 GPUs, AMD EPYC 7713 processors, and 512 GB memory). For a dataset of size N = 3050 (spatial locations) and p = 10,986 (genes), the complete pipeline required about 4 hours on a single GPU node (using 4×A100 GPUs, 16 CPU cores, and 32 GB memory per task). This runtime covers both the hyperparameter optimization using Optuna and the model training. Once the optimal hyperparameters were determined, the final estimation of the coefficient matrix could be completed within only a few minutes.

Code availability

The full code for the model and prospective analyses is available via GitHub at https://github.com/ZhouXY199502/S3R.

References

1. Moses, L. and L. Pachter, *Museum of spatial transcriptomics.* Nat Methods, 2022. **19**(5): p. 534-546.
2. Marx, V., *Method of the Year: spatially resolved transcriptomics.* Nature methods, 2021. **18**(1): p. 9-14.
3. Cang, Z., et al., *Screening cell–cell communication in spatial transcriptomics via collective optimal transport.* Nature methods, 2023. **20**(2): p. 218-228.
4. Dries, R., et al., *Giotto: a toolbox for integrative analysis and visualization of spatial expression data.* Genome biology, 2021. **22**(1): p. 78.
5. Maynard, K.R., et al., *Transcriptome-scale spatial gene expression in the human dorsolateral prefrontal cortex.* Nature neuroscience, 2021. **24**(3): p. 425-436.
6. Ji, A.L., et al., *Multimodal analysis of composition and spatial architecture in human squamous cell carcinoma.* cell, 2020. **182**(2): p. 497-514. e22.
7. Arora, R., et al., *Spatial transcriptomics reveals distinct and conserved tumor core and edge architectures that predict survival and targeted therapy response.* Nature communications, 2023. **14**(1): p. 5029.
8. Svensson, V., S.A. Teichmann, and O. Stegle, *SpatialDE: identification of spatially variable genes.* Nature methods, 2018. **15**(5): p. 343-346.
9. Zhu, J., S. Sun, and X. Zhou, *SPARK-X: non-parametric modeling enables scalable and robust detection of spatial expression patterns for large spatial transcriptomic studies.* Genome biology, 2021. **22**(1): p. 184.
10. Fischer, D.S., A.C. Schaar, and F.J. Theis, *Modeling intercellular communication in tissues using spatial graphs of cells.* Nature Biotechnology, 2023. **41**(3): p. 332-336.
11. Yuan, Y. and Z. Bar-Joseph, *GCNG: graph convolutional networks for inferring gene interaction from spatial transcriptomics data.* Genome biology, 2020. **21**(1): p. 300.
12. Shang, L. and X. Zhou, *Spatially aware dimension reduction for spatial transcriptomics.* Nature Communications, 2022. **13**(1): p. 7203.

13. Ståhl, P.L., et al., *Visualization and analysis of gene expression in tissue sections by spatial transcriptomics.* Science, 2016. **353**(6294): p. 78-82.
14. Chang, W., et al. *Spatially and Robustly Hybrid Mixture Regression Model for Inference of Spatial Dependence*. in *2021 IEEE International Conference on Data Mining (ICDM)*. 2021. IEEE.
15. Kruskal, J.B., *On the shortest spanning subtree of a graph and the traveling salesman problem.* Proceedings of the American Mathematical society, 1956. **7**(1): p. 48-50.
16. Kingma, D.P. and J. Ba, *Adam: A method for stochastic optimization.* arXiv preprint arXiv:1412.6980, 2014.
17. Li, S., et al., *Pytorch distributed: Experiences on accelerating data parallel training.* arXiv preprint arXiv:2006.15704, 2020.
18. Akiba, T., et al. *Optuna: A next-generation hyperparameter optimization framework*. in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2019.
19. Maynard, K.R., et al., *Transcriptome-scale spatial gene expression in the human dorsolateral prefrontal cortex.* Nature Neuroscience, 2021. **24**(3): p. 425-436.
20. Velthuijs, N., et al., *Integration of transcription coregulator complexes with sequence-specific DNA-binding factor interactomes.* Biochim Biophys Acta Gene Regul Mech, 2021. **1864**(10): p. 194749.
21. Brothwell, J.A., et al., *A high-resolution view of the immune and stromal cell response to Haemophilus ducreyi infection in human volunteers.* mBio, 2025. **16**(3): p. e0388524.
22. Cable, D.M., et al., *Robust decomposition of cell type mixtures in spatial transcriptomics.* Nat Biotechnol, 2022. **40**(4): p. 517-526.
23. Chen, C., H.J. Kim, and P. Yang, *Evaluating spatially variable gene detection methods for spatial transcriptomics data.* Genome Biol, 2024. **25**(1): p. 18.
24. Weber, L.M., et al., *nnSVG for the scalable identification of spatially variable genes using nearest-neighbor Gaussian processes.* Nature Communications, 2023. **14**(1): p. 4059.
25. Sun, S., J. Zhu, and X. Zhou, *Statistical analysis of spatial expression patterns for spatially resolved transcriptomic studies.* Nat Methods, 2020. **17**(2): p. 193-200.
26. Peng, X.L., et al., *Determination of permissive and restraining cancer-associated fibroblast (DeCAF) subtypes.* bioRxiv, 2024.
27. Mucciolo, G., et al., *EGFR-activated myofibroblasts promote metastasis of pancreatic cancer.* Cancer Cell, 2024. **42**(1): p. 101-118.e11.
28. Zhang, Y., et al., *Macropinocytosis maintains CAF subtype identity under metabolic stress in pancreatic cancer.* Cancer Cell.
29. Öhlund, D., et al., *Distinct populations of inflammatory fibroblasts and myofibroblasts in pancreatic cancer.* Journal of Experimental Medicine, 2017. **214**(3): p. 579-596.
30. Zhang, D., et al., *Metabolic reprogramming of cancer-associated fibroblasts by IDH3α downregulation.* Cell reports, 2015. **10**(8): p. 1335-1348.
31. Zhao, H., et al., *Tumor microenvironment derived exosomes pleiotropically modulate cancer cell metabolism.* elife, 2016. **5**: p. e10250.
32. Raymant, M., et al., *Macrophage-fibroblast JAK/STAT dependent crosstalk promotes liver metastatic outgrowth in pancreatic cancer.* Nature Communications, 2024. **15**(1): p. 3593.
33. Boelaars, K., et al., *Pancreatic cancer-associated fibroblasts modulate macrophage differentiation via sialic acid-Siglec interactions.* Communications Biology, 2024. **7**(1): p. 430.
34. Jin, S., et al., *Inference and analysis of cell-cell communication using CellChat.* Nat Commun, 2021. **12**(1): p. 1088.

35. Efremova, M., et al., *CellPhoneDB: inferring cell-cell communication from combined expression of multi-subunit ligand-receptor complexes.* Nat Protoc, 2020. **15**(4): p. 1484-1506.
36. Wilk, A.J., et al., *Comparative analysis of cell–cell communication at single-cell resolution.* Nature Biotechnology, 2024. **42**(3): p. 470-483.
37. Fotheringham, A.S., C. Brunsdon, and M. Charlton, *Geographically weighted regression.* The Sage handbook of spatial analysis, 2009. **1**: p. 243-254.
38. Schabenberger, O. and C.A. Gotway, *Statistical methods for spatial data analysis*. 2017: Chapman and Hall/CRC.
39. Williams, C.K. and C.E. Rasmussen, *Gaussian processes for machine learning*. Vol. 2. 2006: MIT press Cambridge, MA.
40. Gelfand, A.E., et al., *Spatial modeling with spatially varying coefficient processes.* Journal of the American Statistical Association, 2003. **98**(462): p. 387-396.
41. Besag, J., *Spatial interaction and the statistical analysis of lattice systems.* Journal of the Royal Statistical Society: Series B (Methodological), 1974. **36**(2): p. 192-225.
42. Tibshirani, R., et al., *Sparsity and smoothness via the fused lasso.* Journal of the Royal Statistical Society Series B: Statistical Methodology, 2005. **67**(1): p. 91-108.
43. Beer, J.C., et al., *Incorporating prior information with fused sparse group lasso: Application to prediction of clinical measures from neuroimages.* Biometrics, 2019. **75**(4): p. 1299-1309.
44. Li, F. and H. Sang, *Spatial Homogeneity Pursuit of Regression Coefficients for Large Datasets.* Journal of the American Statistical Association, 2018. **114**: p. 1-37.
45. Zhong, Y., et al., *Sparse spatially clustered coefficient model via adaptive regularization.* Comput Stat Data Anal, 2023. **177**: p. 107581.

Acknowledgments

This work was supported by research grants IIS-2145314 and DBI-2047631 from the National Science Foundation; and R35GM155028 from the National Institutes of Health; and RSG-24-1321371-01-CDP from the American Cancer Society; and the OHSU Brenden-Colson Center for Pancreatic Care.

Author contributions

C.S. conceived the methodological framework. Z.C. designed all the real-world data applications. Z.X. designed the algorithm, simulation validation, and conducted the case studies. D.P. provided key guidance on algorithm implementation. N.B. and T.H. contributed to methodological guidance and discussion interpretation. P.L.X. and J.J.Y. provided well-annotated PDAC spatial and single cell RNA-seq data. P.L.X., J.J.Y., S.R and Z.T. contributed to applications of methods to PDAC. Z.C., Z.X., C.S. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Supplementary

Evaluation of S3R on simulation data

To rigorously validate the efficacy of the S3R method, we conducted a systematic evaluation using synthetic datasets. Ground-truth coefficient matrices were generated with 3–5 distinct spatial clusters and varying feature sparsity levels to reflect biologically plausible regulatory networks. Additive Gaussian noise (σ=0.5,1,2) was introduced to simulate technical variability, and feature dimensionalities spanned 20, 100, and 1000, representing low to high- dimensional feature spaces. A visual comparison of the clustering accuracy between S3R and four other competing methods, including RSCC, EN, GWL and VS-GPSVC was shown in **Extended Fig. 1a**. In this simulation scenario for the left four columns, the predictor matrix $\boldsymbol{X}$ is of dimension 100×20, depicting a relatively low-dimensional feature setting; the predictor matrix $\boldsymbol{X}$ is of dimension 100×2000 for the right two columns, depicting a high dimensional setting. We introduced different number of coefficient clusters that exist among the response and the predictors, such as 3,4,5,6. The detected clusters by different methods were colored differently. The spatial map on the first row shows the true coefficient matrix. The map in the second row represents the coefficient clusters given by our proposed S3R algorithm, which shows a high degree of similarity with the true spatial dissection, similar for the other methods shown in the rest of the rows.

A more detailed quantitative evaluation is presented in Extended **Fig. 1b**, where we report the Adjusted Rand Index (ARI) between the true coefficient clusters, and the coefficient clusters given by different methods, obtained under five cluster scenarios (3–6 clusters) and three noise levels (σ = 0.5, 1, 2). The first four columns represent the low-dimensional settings (N × p = 100 × 20 to 2000 × 20), while the last two illustrate the high-dimensional cases (N × p = 200 × 1000 and 400 × 1000). Across all panels, ARI scores decline as the Gaussian noise increases, yet S3R consistently attains the highest values, outperforming RSCC, EN, GWL, and VS-GPSVC in every configuration. Notably, the gap between S3R and the competing methods widens in the high-dimensional regime, underscoring the effectiveness of its sparsity and spatial smoothness constraints when p far exceeds N.

Runtime comparison of S3R and competing methods

To assess computational scalability, we compared the runtime of S3R with RSCC, Elastic Net (EN), Graph-guided Weighted Lasso (GWL), and VS-GPSVC across datasets with varying sample sizes (N = 200–2000) and feature dimensions (P = 20–1000). As shown in Supplementary **Fig. 1**, all methods perform similarly in low-dimensional settings (e.g., N = 200, P = 20), although GWL consistently incurs higher runtime. With increasing dimensionality (e.g., N = 400, P = 100; N = 1000, P = 100), VS-GPSVC becomes substantially slower. In high-dimensional regimes (e.g., N = 200–2000, P = 1000), S3R maintains practical runtime while several competing methods either slow down considerably or fail to complete (red “X” marks). Across all scenarios, S3R achieves favorable runtime relative to accuracy. All runtime experiments were conducted on a personal PC.

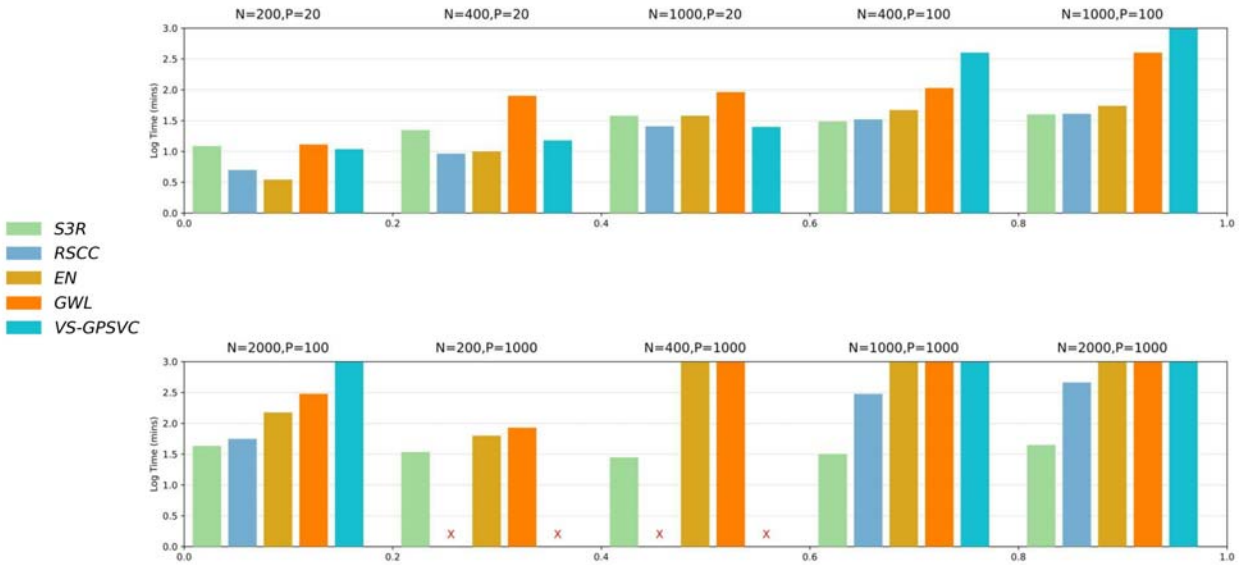

Supplementary Figure 1 | The runtime comparison of S3R and other methods.