FlashDeconv enables atlas-scale, multi-resolution spatial deconvolution via structure-preserving sketching

Chen Yang[1], Xianyang Zhang[1*], Jun Chen[2*]

[1*]Department of Statistics, Texas A&M University, College Station, Texas, 77843, USA.
[2]Division of Computational Biology, Department of Quantitative Health Sciences, Mayo Clinic, Rochester, Minnesota, 55905, USA.

*Corresponding author(s). E-mail(s): zhangxiany@stat.tamu.edu; Chen.Jun2@mayo.edu;
Contributing authors: cafferychen777@tamu.edu;

Abstract

Coarsening Visium HD resolution from 8 to 64 μm can flip cell-type co-localization from negative to positive ($\boldsymbol{r = -0.12 \rightarrow +0.80}$), yet investigators are routinely forced to coarsen because current deconvolution methods cannot scale to million-bin datasets. Here we introduce FlashDeconv, which combines leverage-score importance sampling with sparse spatial regularization to match top-tier Bayesian accuracy while processing 1.6 million bins in 153 seconds on a standard laptop. Systematic multi-resolution analysis of Visium HD mouse intestine reveals a tissue-specific *resolution horizon* (8–16 μm)—the scale at which this sign inversion occurs—validated by Xenium ground truth. Below this horizon, FlashDeconv provides the first sequencing-based quantification of Tuft cell chemosensory niches (15.3-fold stem cell enrichment). In a 1.6-million-bin human colorectal cancer cohort, FlashDeconv further uncovers neutrophil inflammatory microdomains in low-UMI regions that classification-based methods discard, recovering spatially organized biology from measurements previously considered uninformative.

Keywords: Spatial transcriptomics, Cell type deconvolution, Visium HD, Atlas-scale analysis, Randomized sketching

Introduction

Spatial transcriptomics (ST) technologies [1], such as Stereo-seq [2], Xenium [3], and Visium HD, are transforming our understanding of tissue architecture by mapping gene expression at subcellular resolution across increasingly large fields of view. As these technologies scale to atlas-level datasets comprising millions of spots or cells, the computational burden of analyzing such data has become a critical bottleneck. A fundamental task in ST analysis is cell type deconvolution, which infers the proportional composition of cell types within each spatial measurement unit by leveraging single-cell RNA sequencing (scRNA-seq) references.

Existing deconvolution methods—now numbering over seventy distinct approaches [4]—can be broadly distinguished by their modeling paradigms: generative probabilistic modeling versus constrained regression. Generative methods such as Cell2Location [5] and Stereoscope [6] achieve high accuracy by explicitly modeling count data distributions (e.g., Negative Binomial) via variational inference or maximum a posteriori (MAP) estimation, providing rigorous uncertainty quantification that is particularly valuable for small-scale studies requiring careful statistical inference. However, these methods require extensive iterative training, with runtimes scaling from hours to days for million-scale datasets. In contrast, regression-based approaches formulate the problem as constrained minimization, offering faster inference. This category includes the baseline NNLS, weighted variants such as MuSiC [7] and SpatialDWLS [8], likelihood-based regression (RCTD [9]), NMF decomposition (SPOTlight [10]), and deep-learning alignment (Tangram [11]). Benchmarks indicate that many methods fail to outperform NNLS consistently; only Cell2Location, RCTD, and MuSiC exceed baseline performance across all evaluation metrics [12]. However, most regression approaches treat spatial spots as independent observations, ignoring local tissue continuity. Methods that do incorporate spatial structure, such as CARD [13] with its conditional autoregressive (CAR) prior, typically rely on dense covariance matrices ($O(N^2)$ memory), precluding analysis of emerging high-resolution platforms with millions of spots.

Here, we present FlashDeconv, a deconvolution framework that achieves accuracy, spatial awareness, and linear scalability simultaneously. FlashDeconv leverages Randomized Numerical Linear Algebra (RandNLA) to compress the high-dimensional gene expression space ($G \approx 20,000$) into a compact "sketch" ($d \approx 512$) using a structure-preserving sketching technique. Unlike standard random projections that risk losing signals from rare cell types, our approach utilizes leverage-score importance sampling to effectively preserve biological heterogeneity. Here, "structure" refers to the geometric organization of cell types in gene expression space, which leverage scores quantify by measuring each gene's contribution to distinguishing cell type signatures. Combined with a Log-CPM data representation and a sparse graph Laplacian regularizer that scales as $O(N)$ rather than $O(N^2)$, FlashDeconv achieves accuracy comparable to top-tier Bayesian methods on synthetic benchmarks while reducing runtime by orders of magnitude, enabling atlas-scale analysis on standard commodity hardware.

A critical challenge in deconvolution is that traditional dimension reduction methods—such as Principal Component Analysis (PCA) or Highly Variable Genes (HVG) selection—rely on *variance* as a proxy for biological information [14, 15]. Intuitively,

variance measures how "loud" a gene is across the dataset—a quantity naturally dominated by abundant cell types. In contrast, leverage scores measure *distinctiveness*: whether a gene defines a unique direction in the transcriptomic space, regardless of how many cells express it. This distinction is biologically consequential: rare populations—including cancer stem cells comprising <1% of tumors [16, 17] and vascular endothelial cells that orchestrate tissue-specific stem cell niches [18]—can carry disproportionate functional significance in tissue homeostasis and pathology, yet their markers are systematically underrepresented by variance-based feature selection [19, 20]. The need to decouple biological importance from population frequency has motivated alternative approaches, including Gini-index-based gene selection [21] and deviance-based feature ranking [22]. We validate this variance-leverage decoupling through systematic experiments (Section 2), demonstrating that leverage scores identify biologically meaningful markers independently of cell type abundance. Applied to atlas-scale Visium HD data, FlashDeconv further enables systematic multi-resolution analysis, characterizing a tissue-specific "resolution horizon" for spatial information loss and revealing cellular architecture inaccessible at conventional spatial scales.

Taken together, FlashDeconv makes atlas-scale, multi-resolution spatial deconvolution practical on commodity hardware, supporting both million-bin Visium HD mapping and data-driven characterization of tissue-specific resolution horizons. In the following Results, we first describe the FlashDeconv framework and leverage-weighted sketching, then benchmark accuracy and linear scalability, and finally demonstrate atlas-scale and multi-resolution applications on Visium HD data.

Results

The FlashDeconv framework

FlashDeconv formulates spatial deconvolution as a constrained optimization problem in a compressed feature space (Fig. 1). The framework consists of three key design choices:

First, to address the extreme sparsity and mean-variance dependency of ST data, we employ a Log-CPM (Counts Per Million) transformation. While Pearson residuals are statistically principled for negative binomial count data [23], Log-CPM offers specific advantages for L_2-based sketching: its bounded norm prevents high-expression genes from dominating the sketch space, while its logarithmic compression stabilizes the extreme dynamic range of spatial count data. This represents an engineering trade-off—sacrificing some statistical optimality for compatibility with randomized compression (Supplementary Note 2).

Second, we tackle the computational redundancy of gene expression via structure-preserving randomized sketching. Instead of solving the regression on all G genes, we project the data into a lower-dimensional subspace using a sparse CountSketch matrix Ω. Unlike PCA, which maximizes explained variance and may obscure cell types that contribute little to global variance, our sketching matrix satisfies the Johnson-Lindenstrauss property [24], guaranteeing that Euclidean distances between cell type signatures are maintained in the compressed space with high probability. Crucially,

the projection is weighted by statistical leverage scores [25, 26] derived from the single-cell reference. This ensures that marker genes defining transcriptomically distinct cell types—which often have high leverage despite low variance—are preserved with high probability, avoiding the signal loss that can occur with uniform subsampling or variance-based feature selection. For example, in liver scRNA-seq data, Central Vein Endothelial cells constitute only 2% of cells, yet their marker gene *Rspo3* achieves the highest leverage score (0.0105) among all genes due to its unique expression pattern that sharply distinguishes this population from all others. In contrast, Hepatocyte markers like *Egfr* exhibit low leverage scores (0.0007) despite high specificity, because Hepatocytes share metabolic gene programs with Cholangiocytes, reducing their geometric distinctiveness in the reference space. This illustrates that leverage captures *functional distinctiveness* rather than abundance: Hepatocytes are abundant but transcriptionally overlap with other lineages, while Central Vein Endothelial cells are rare but transcriptomically unique. Standard variance-based methods (PCA, HVG) show no statistical difference between functionally distinct and overlapping markers (Mann-Whitney U test, $p = 0.97$), whereas leverage-score ranking successfully identifies the structurally informative genes (Mann-Whitney U test, $p = 0.0062$). We validate this principle systematically in Section 2. Leverage-based weighting also addresses a second, algorithmic challenge: during the hashing step of CountSketch, high-abundance genes can dominate compressed dimensions and overwhelm rare cell markers assigned to the same hash bucket. By scaling genes in proportion to their leverage scores before compression ($w_g = \sqrt{p_g \cdot G}$, where p_g is the normalized leverage probability), FlashDeconv ensures that rare cell markers receive preferential amplification relative to their baseline magnitude, maintaining adequate signal strength despite random hash collisions [27] (see Methods for formal derivation and empirical validation across six tissues).

Third, we incorporate spatial information using a graph Laplacian regularizer. Unlike CARD, which constructs and operates on a dense $N \times N$ spatial kernel matrix with $O(N^2)$ memory and $O(N^2 \cdot K)$ per-iteration complexity (where K is the number of cell types), our approach imposes spatial smoothness via a sparse penalty term $\mathrm{Tr}(\beta^T L \beta)$—a formulation originally developed for structured genomic data [28] and shared with recent spatial-aware methods such as SDePER [29] and FAST [30]—where L is the graph Laplacian of the spatial neighbor network. This regularization induces local smoothness that is mathematically equivalent to a Gaussian Markov Random Field (GMRF) [31] but requires only sparse matrix operations. Because each spot connects to only k neighbors (default $k = 6$), the spatial term scales as $O(N \cdot k)$ rather than $O(N^2)$, yielding linear time complexity. This formulation allows FlashDeconv to model spatial autocorrelation across millions of spots using a fast Block Coordinate Descent (BCD) algorithm, a scale where dense matrix inversions are computationally prohibitive. Biologically, this regularization encodes the intuition that tissue composition varies smoothly: for a spot with sparse sequencing depth, the algorithm borrows strength from its spatial neighbors, allowing coherent tissue structures to emerge from noisy measurements.

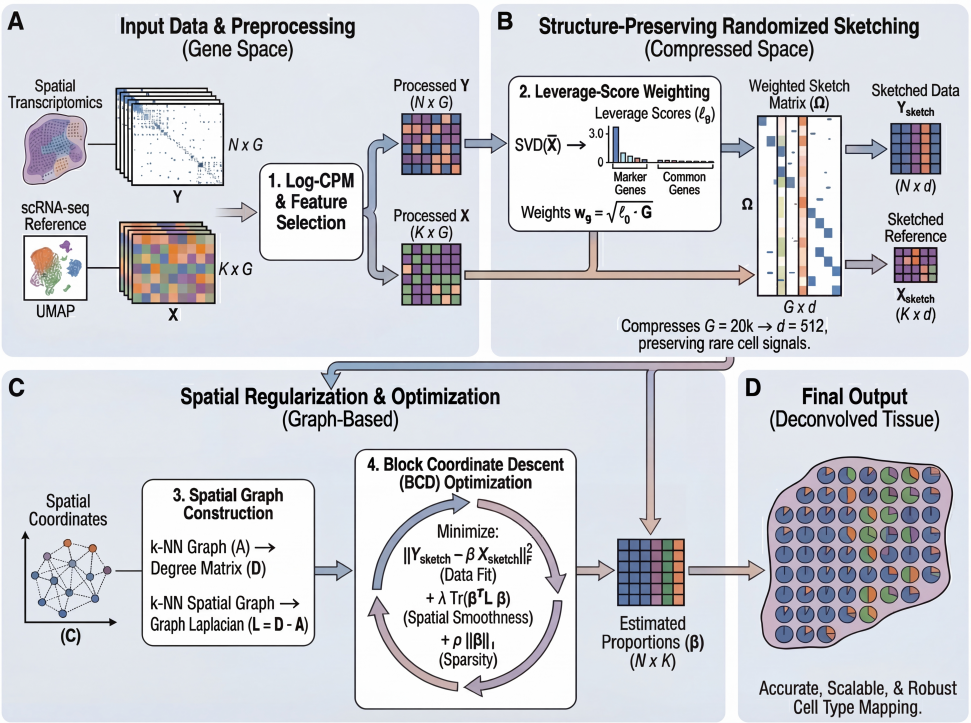

Fig. 1 Overview of the FlashDeconv framework. (a) Input data and preprocessing. Spatial transcriptomics data Y $(N \times G)$ and scRNA-seq reference signatures X $(K \times G)$ are normalized using Log-CPM transformation. Gene selection combines HVGs from Y with cell-type markers from X (whole transcriptome), yielding the union set $\mathcal{G}$. **(b)** Structure-preserving sketching. Leverage scores ℓ_g are computed from the SVD of the reference matrix to identify discriminative genes. A weighted sketch matrix Ω $(G \times d)$ compresses the gene space from $G \approx 20{,}000$ to $d = 512$ dimensions while preserving rare cell type signals. **(c)** Spatial regularization and optimization. A k-nearest neighbor graph is constructed from spatial coordinates, and the graph Laplacian L enforces spatial smoothness. The optimization problem minimizes reconstruction error with spatial and sparsity regularization via block coordinate descent (BCD). **(d)** Final output. FlashDeconv produces cell type proportion estimates β $(N \times K)$ for each spatial location, enabling accurate, scalable, and robust cell type mapping across the tissue.

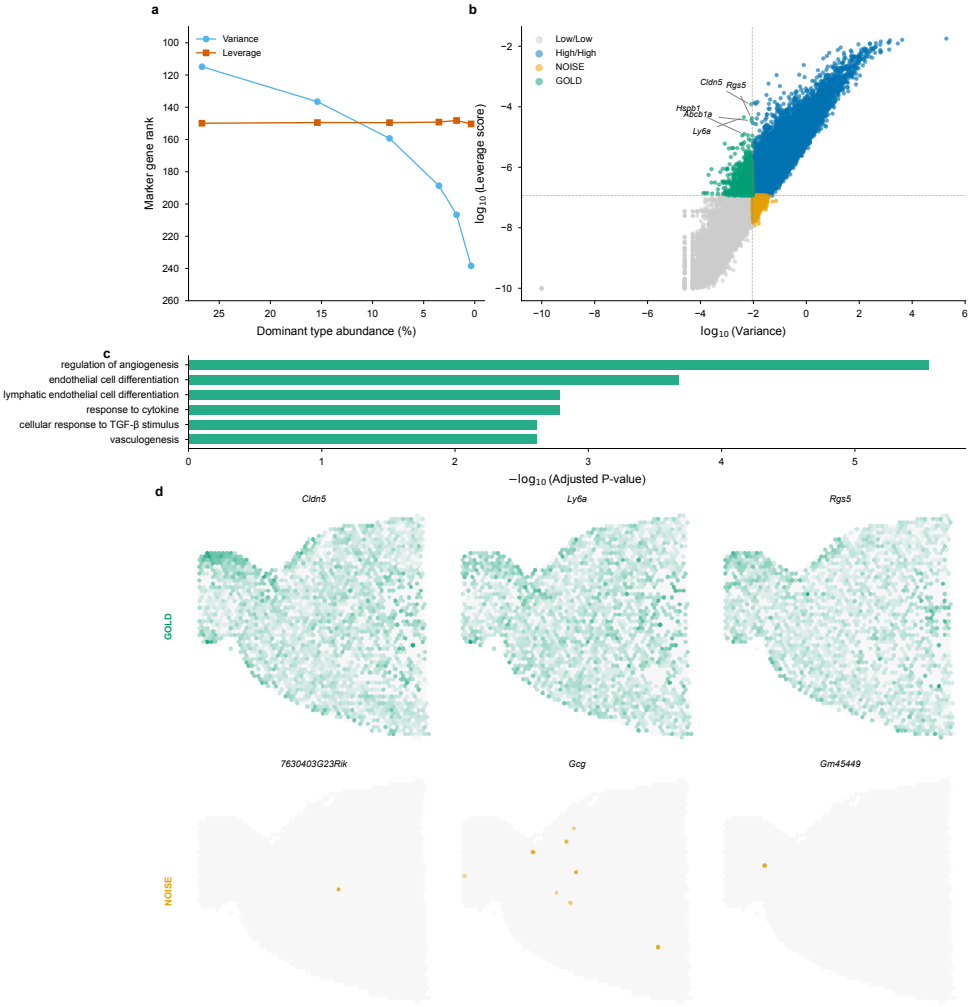

Fig. 2 Leverage scores decouple biological identity from population abundance. (a) Abundance invariance test. Ranking stability of oligodendrocyte markers (top 20 genes by expression in oligodendrocytes) in mouse brain scRNA-seq (31,053 genes) as cell population is downsampled from 26.7% to 0.4%. Rank denotes average position when all genes are sorted by score (rank 1 = highest). Variance-based ranking (blue) degrades from rank 115 to 240 as abundance decreases—a two-fold deterioration. Leverage-score ranking (red) remains stable at rank ∼150 regardless of population size, demonstrating true decoupling of biological identity from numerical prevalence. **(b)** The variance-leverage plane. Classification of 31,053 genes by variance (x-axis) and leverage score (y-axis). Four quadrants emerge: structurally informative "GOLD" genes (green, low variance/high leverage) include vascular markers (*Cldn5*, *Rgs5*, *Ly6a*, *Abcb1a*, *Hspb1*) that define rare anatomical structures; variance-dominated "NOISE" genes (red, high variance/low leverage) contain 35% unannotated *Gm*-series transcripts compared to only 6% in the GOLD set, indicating that high variance alone does not ensure cell-type discriminative power. **(c)** Functional enrichment analysis. GO Biological Process enrichment reveals GOLD genes are significantly enriched for regulation of angiogenesis (FDR-adjusted $p = 2.8 \times 10^{-6}$), endothelial cell differentiation (FDR-adjusted $p = 2.1 \times 10^{-4}$), vasculogenesis, and blood vessel morphogenesis. NOISE genes show zero significant GO terms at FDR-adjusted $p < 0.05$. Genome-wide cell type specificity analysis further confirms that GOLD genes systematically target rare populations (median 0.27% abundance) versus NOISE genes (0.51%; $p = 3.25 \times 10^{-25}$), with Endothelial cells as the top target—validating leverage as an unsupervised metric for biological distinctiveness. **(d)** Spatial verification on Visium. Top row: GOLD genes (*Cldn5*, *Ly6a*, *Rgs5*) reconstruct clear vascular anatomical structures on mouse brain Visium sections (spatial structure score = 1.33). Bottom row: NOISE genes exhibit random, speckle-like distribution patterns (structure score = 0.87; Mann-Whitney $p = 5.6 \times 10^{-5}$). This visual contrast demonstrates that leverage selects for genuine biological structure rather than technical variation.

Leverage decouples biological identity from population abundance

A fundamental premise of FlashDeconv is that statistical variance conflates biological signal with population frequency, systematically disadvantaging rare cell types. To systematically test this claim, we designed a series of experiments that trace the evidence from mathematical principle through molecular function to spatial phenotype.

Abundance invariance test.

We performed an *in silico* stress test using the mouse brain scRNA-seq reference (40,532 cells, 59 cell types, 31,053 genes). Starting from the native cell type distribution, we artificially downsampled the dominant oligodendrocyte population from 26.7% to 0.4%—a 67-fold reduction—while keeping all other populations constant. At each downsampling level, we identified the top 20 marker genes (by expression level in oligodendrocytes), computed both variance-based (normalized dispersion [32]) and leverage-based rankings across all genes, and tracked how the average marker rank changed (Fig. 2a). The results reveal a striking divergence: variance-based ranking degraded linearly with abundance, with marker rank dropping from 115 to 240 as the population shrank—a two-fold deterioration. In contrast, leverage-score ranking remained stable at rank ∼150 throughout the entire range, demonstrating that leverage successfully decouples a gene's discriminative power from its parent cell type's numerical prevalence. Quantitatively, the stability can be attributed to leverage's reduced coupling to mean expression: while gene variance correlates almost perfectly with mean expression (Spearman $\rho = 0.998$), leverage scores show substantially greater spread around this relationship ($\rho = 0.965$), enabling identification of low-abundance but high-specificity markers (Supplementary Fig. S5).

Gene quadrant analysis reveals systematic bias.

To understand the broader implications of this decoupling, we mapped all 31,053 genes into a variance-leverage coordinate system (Fig. 2b). Four distinct quadrants emerged, defined by median log-transformed values to ensure equal partitioning without arbitrary threshold selection. The "GOLD" quadrant (low variance, high leverage) contains genes that standard HVG selection would discard but that carry high discriminative power—including classic vascular markers such as *Cldn5* (claudin-5, tight junction protein), *Rgs5* (pericyte marker), *Ly6a* (stem cell antigen), *Abcb1a* (blood-brain barrier transporter), and *Hspb1* (heat shock protein). These genes define rare but anatomically critical cell populations: brain endothelial cells and pericytes constitute <3% of brain tissue yet form the blood-brain barrier essential for neural function. Conversely, the "NOISE" quadrant (high variance, low leverage) contains genes with high variance but low cell-type specificity: 35% are unannotated *Gm*-series transcripts (predicted genes with unknown function) compared to only 6% in the GOLD set—a 6-fold enrichment indicating that variance-based selection systematically prioritizes genes lacking discriminative power for cell type identification.

Functional enrichment provides molecular validation.

To confirm that the GOLD/NOISE distinction reflects genuine biological differences rather than statistical coincidence, we performed Gene Ontology enrichment analysis using Enrichr [33, 34] with Benjamini-Hochberg FDR correction on both gene sets (Fig. 2c). GOLD genes showed highly significant enrichment for vascular biology: *regulation of angiogenesis* (FDR-adjusted $p = 2.8 \times 10^{-6}$), *endothelial cell differentiation* (FDR-adjusted $p = 2.1 \times 10^{-4}$), *vasculogenesis* (FDR-adjusted $p = 2.4 \times 10^{-3}$), and *blood vessel morphogenesis*. This coherent functional signature—blood vessel development in brain tissue—confirms that leverage identifies genes with coordinated biological roles. In stark contrast, NOISE genes yielded *zero* significant GO terms at FDR-adjusted $p < 0.05$, consistent with their enrichment for unannotated transcripts and suggesting that high variance in the absence of high leverage indicates expression variability without coordinated cell-type specificity.

Genome-wide specificity analysis rules out selection bias.

To verify that the GOLD/NOISE distinction reflects intrinsic biological structure rather than post-hoc cherry-picking, we performed a genome-wide cell type specificity analysis across all 59 cell types in the reference (Supplementary Fig. S10). For every gene in the genome, we identified its primary cell type target based on expression specificity, without any manual pre-selection. Strikingly, GOLD genes preferentially mark significantly rarer cell populations (median target abundance 0.27%) compared to NOISE genes (median 0.51%; Mann-Whitney U test $p = 3.25 \times 10^{-25}$). Crucially, the top cell type targeted by the GOLD gene set was Endothelial cells (0.1% abundance)—consistent with the independent GO enrichment for angiogenesis. This confirms that the variance-leverage plane objectively separates markers of rare anatomical structures from genes lacking cell-type discriminative power, without requiring prior knowledge of cell type identity.

Spatial visualization confirms anatomical relevance.

The ultimate test of a gene's biological importance is whether it corresponds to recognizable tissue structure. We visualized representative GOLD and NOISE genes on mouse brain Visium sections (Fig. 2d). GOLD genes (*Cldn5*, *Ly6a*, *Rgs5*) reconstructed clear vascular anatomical patterns—bright spots tracing blood vessel distributions consistent with known brain vasculature. To quantify this, we computed a spatial structure score measuring the ratio of global to local variance (higher values indicate organized spatial patterns). GOLD genes achieved a structure score of 1.33 ± 0.23, significantly exceeding NOISE genes at 0.87 ± 0.54 (Mann-Whitney $p = 5.6 \times 10^{-5}$). Visually, NOISE genes exhibited random, speckle-like "salt-and-pepper" distributions lacking coherent spatial structure. This spatial validation closes the evidence loop: genes identified as important by leverage scores correspond to real anatomical structures, while genes prioritized by variance alone do not.

Downstream impact on deconvolution accuracy.

Finally, we tested whether these theoretical insights translate to practical performance gains through systematic ablation experiments (Supplementary Fig. S7).

We compared four dimensionality reduction approaches within the FlashDeconv framework—Leverage-Score Sketching (LSS), Uniform CountSketch, PCA, and HVG selection—using identical downstream optimization across 54 Silver Standard datasets (6 tissues $\times$ 9 abundance scenarios, excluding 2 replicate conditions). LSS achieved rare cell type detection of $r = 0.35$, compared to Uniform ($r = 0.14$, +147%), PCA ($r = 0.16$, +124%), and HVG ($r = 0.12$, +197%). A stress test revealed the critical distinction: when cell type abundance was artificially reduced to 0.17%, LSS maintained detection ($r = 0.62$) while PCA ($r = 0.02$) and Uniform ($r = 0.08$) failed completely. At 0.03% abundance, LSS remained positive ($r = 0.33$) while alternatives showed negative correlation, indicating complete signal loss. These results demonstrate that leverage-based sketching preferentially preserves rare cell type signals across diverse tissue contexts.

Cross-tissue generalization.

To confirm that the GOLD/NOISE framework is not specific to brain tissue, we repeated the quadrant analysis on mouse kidney scRNA-seq data from the Spotless benchmark (7,501 cells, 16 cell types). Podocytes—glomerular epithelial cells critical for kidney filtration—constitute only 0.17% of the reference population, representing an extreme test case for rare cell detection. All five canonical podocyte markers (*Nphs1*, *Nphs2*, *Podxl*, *Synpo*, *Wt1*) appeared in the GOLD quadrant, while markers of abundant tubular cell types (60% of cells) fell in high-variance regions (Supplementary Fig. S12). This cross-tissue consistency confirms that the variance-leverage separation reflects a general property of gene expression structure rather than a tissue-specific artifact.

Together, these experiments establish that leverage scores provide a principled, abundance-independent measure of biological informativeness. Critically, the ablation comparison isolates the effect of leverage weighting during compression: leverage-weighted and uniform sketching operate on identical gene sets (HVG $\cup$ markers), yet leverage weighting improves rare cell detection by 147%—demonstrating that the benefit arises from preserving marker signals during the G-to-512 projection, not from gene selection per se (see Methods for two-stage design rationale). By incorporating leverage-weighted importance sampling into the sketching process, FlashDeconv corrects a systematic bias inherent to variance-based feature selection, enabling accurate detection of rare but biologically critical cell populations.

FlashDeconv achieves competitive accuracy on synthetic benchmarks

We systematically evaluated FlashDeconv using the *Spotless* benchmark suite [12], a reproducible pipeline for benchmarking cell type deconvolution methods. Spotless provides two types of ground-truth data: (1) Silver Standards, synthetic "pseudo-spots" generated by computationally mixing single-cell transcriptomes with known cell type proportions, and (2) Gold Standards, real spatial transcriptomics data (seq-FISH+, STARMap) where ground-truth composition is derived from co-registered single-molecule FISH imaging.

We benchmarked FlashDeconv on all 56 publicly available Silver Standard datasets from Spotless (6 tissues $\times$ 9 abundance patterns, plus 2 replicate conditions). Across these datasets representing various tissue types and abundance scenarios (e.g., dominant cell types, rare cell types, uniform distributions), FlashDeconv achieved a mean Pearson correlation of 0.944 and median RMSE of 0.065 (Fig. 3a,b).

Notably, FlashDeconv outperformed the baseline NNLS method (Pearson ≈ 0.90) and matched or exceeded the performance of computationally intensive methods such as Cell2Location (RMSE $\sim 0.05 - 0.08$) and RCTD (RMSE $\sim 0.06 - 0.09$). We specifically investigated the "Rare Cell Type" scenarios, where one cell type is 5-15 times less abundant than others and standard random projections often fail. FlashDeconv achieved mean AUPR $= 0.960 \pm 0.036$ (s.d., $n = 56$ datasets) on rare cell type detection, matching the performance of top-tier probabilistic methods (RCTD, Cell2Location: AUPR 0.95–0.98) and confirming that our leverage-score-based sketching successfully preserves signals from low-abundance populations.

To validate performance on platform effects, we utilized the "Gold Standard" seqFISH+ and STARMap datasets. On the STARMap dataset (108 spots), FlashDeconv ranked #1 out of 13 methods, demonstrating its robustness on data with realistic noise structures. On the extremely small seqFISH+ datasets (<10 spots per FOV), FlashDeconv maintained competitive performance, although the benefits of sketching and spatial regularization are naturally less pronounced in micro-scale regimes (Supplementary Fig. S8).

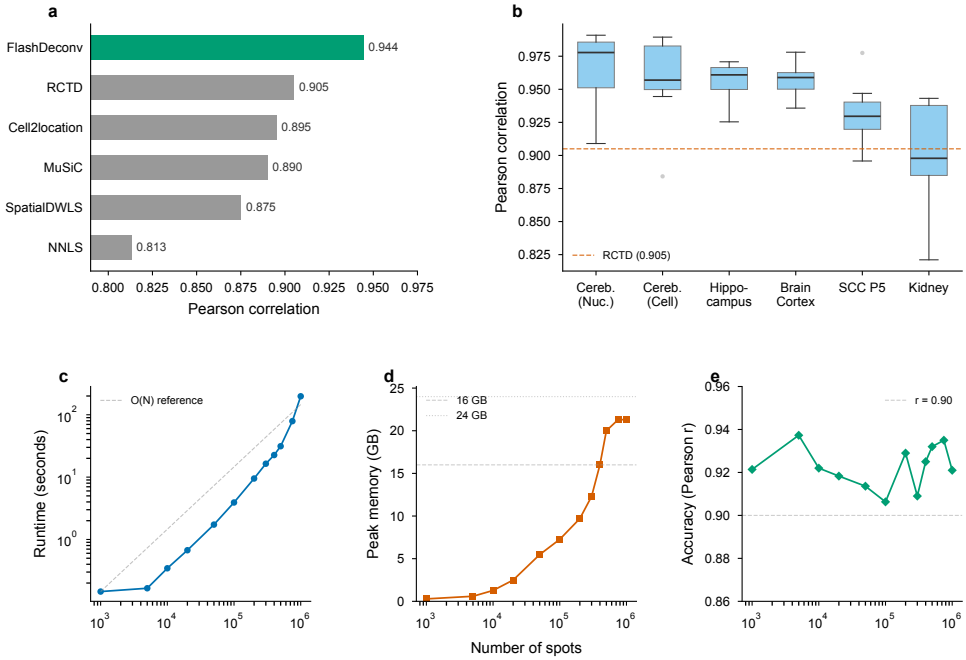

Fig. 3 FlashDeconv achieves competitive accuracy with linear scalability. (a) Comparison with published methods on 56 Silver Standard datasets. FlashDeconv achieves mean Pearson = 0.944, compared to RCTD (0.905) and Cell2Location (0.895). Note that synthetic performance does not always translate to real-data ranking (see Section 2). **(b)** Performance stratified by tissue type. FlashDeconv maintains high accuracy across brain tissues (Pearson > 0.95 for cortex, cerebellum, hippocampus), with lower performance on kidney (Pearson = 0.90) due to highly correlated cell type signatures. **(c)** Runtime approaches linear scaling for large datasets, converging toward the $O(N)$ reference line (dashed) as dataset size increases. FlashDeconv processes 100K spots in under 4 seconds and 1M spots in approximately 3 minutes. **(d)** Peak memory usage scales linearly up to ∼400K spots, remaining within 24 GB for million-scale datasets—well within commodity hardware limits. **(e)** Accuracy (Pearson correlation with ground truth) is preserved across all scales, maintaining $r > 0.90$ throughout, demonstrating that the sketching approximation achieves substantial computational savings without sacrificing deconvolution quality. Results are highly reproducible: pairwise correlations across 10 runs with different random seeds exceed $r > 0.98$ (Supplementary Fig. S6).

Real-world validation on Visium case studies

To assess performance on real Visium data with biologically grounded ground truth, we evaluated FlashDeconv on two case studies from the Spotless benchmark: mouse liver tissue sections and a melanoma tumor dataset. These case studies provide direct comparison with 12 competing methods on actual spatial transcriptomics data rather than simulations.

The liver dataset consists of four Visium slides (5,762 total spots) with zonation annotations and matched snRNA-seq reference (133,779 cells, 9 cell types). Following the Spotless protocol, we evaluated two metrics: Jensen-Shannon Divergence (JSD) against snRNA-seq-derived proportions, and Area Under the Precision-Recall curve

(AUPR) for detecting portal and central vein endothelial cells based on spatial zonation patterns. FlashDeconv achieved a JSD of 0.0561, ranking 3rd out of 13 methods, behind only RCTD (0.0334) and Cell2Location (0.0352), and substantially outperforming the NNLS baseline (0.1056). On the AUPR metric, FlashDeconv achieved 0.66, ranking 7th. The predicted hepatocyte proportion (68.4%) aligned closely with the expected dominant presence of hepatocytes in liver tissue.

Beyond accuracy metrics, we evaluated FlashDeconv's robustness to reference protocol variations using the Spotless liver reference sensitivity test. This test measures prediction consistency when using three different reference protocols: exVivo scRNA-seq, inVivo scRNA-seq, and snRNA-seq. FlashDeconv achieved a stability JSD of 0.0138, ranking 1st out of 13 methods—25% better than RCTD (0.0185), the second-best performer. This robustness indicates that FlashDeconv's predictions are less sensitive to the specific reference protocol choice, a practically valuable property given the heterogeneity of available scRNA-seq references.

The melanoma dataset includes three Visium slides (7,557 spots) with Molecular Cartography-derived ground truth proportions for 7 cell types. We evaluated FlashDeconv using the default sketch dimension ($d = 512$) and the full reference (15 cell types), identifying a fundamental trade-off between distribution-based similarity and absolute abundance estimation. When optimized for the Jensen-Shannon Divergence (JSD) metric—which penalizes small relative errors in rare cell types—FlashDeconv using Pearson residual preprocessing achieved a JSD of 0.015, ranking 3rd among 13 methods, behind only Cell2Location and SPOTlight. Under this configuration, FlashDeconv accurately captured rare T-cell populations (3.3% predicted vs. 0.3% with log-CPM; ground truth 4.7%) but underestimated the dominant melanocytic cells (75.5% vs. 84.8% ground truth). Alternatively, using log-CPM preprocessing prioritized the dominant signal, yielding a more accurate melanocytic proportion (81.2%) but a higher JSD (0.033, rank 7) due to increased sparsity in rare cell estimates. This sensitivity analysis demonstrates that FlashDeconv's flexible preprocessing allows users to prioritize either broad-spectrum distribution fidelity or dominant-cell accuracy, with both modes maintaining robustness against the extreme collinearity of malignant cell states ($\kappa = 63.4$; Supplementary Note 2).

These results demonstrate that FlashDeconv achieves competitive accuracy on real spatial data, with performance varying across tissue types. The liver dataset represents a favorable scenario: well-separated cell type signatures and moderate spatial structure align with FlashDeconv's design, yielding strong performance (rank 3). Notably, rare cell type markers in liver (Cholangiocytes 1.16%, Mesothelial cells 0.53%, Portal/Central Vein Endothelial cells ∼2%) exhibit significantly higher leverage scores than markers of abundant types ($p = 0.0062$), a separation absent in variance-based ranking ($p = 0.97$; Fig. 2)—extending the variance-leverage decoupling established on brain tissue (Section 2) to a second organ. On the melanoma dataset, where malignant cell states exhibit extreme transcriptomic similarity ($r > 0.98$ for all 42 pairwise correlations; Supplementary Fig. S4), FlashDeconv with Pearson preprocessing achieves accuracy comparable to top probabilistic methods (rank 3). This demonstrates that appropriate preprocessing can partially overcome the collinearity challenge inherent in linear deconvolution, though at the cost of reduced accuracy for dominant

populations—a trade-off that users can navigate based on their analytical priorities (Supplementary Note 2). This variability is consistent with the Spotless study's conclusion that no single method excels universally across all tissue contexts [12]; complete performance metrics for all 13 methods across synthetic benchmarks, gold standards, and case studies are provided in Supplementary Data .

Linear scalability enables atlas-scale deconvolution

The defining advantage of FlashDeconv is its scalability. We benchmarked runtime and memory usage across datasets ranging from 10^3 to 10^6 spots (Fig. 3c–e). For a dataset with 100,000 spots, FlashDeconv completed the analysis in under 4 seconds. In contrast, deep learning methods face severe computational constraints at much smaller scales: benchmark studies report that Stereoscope requires over 8 hours on datasets with only 10,000 spots [12], and Cell2Location exhibits runtimes nearly 100-fold longer than regression-based alternatives [35]. Even fast regression-based methods face scalability constraints: per RCTD documentation, 11,000 spots require approximately 20 minutes on 4 cores; extrapolating linearly, 1 million spots would require over 30 hours—compared to FlashDeconv's 3 minutes for the same scale.

Crucially, FlashDeconv exhibits $O(N)$ linear scaling for both time and memory. On commodity hardware (32GB unified memory, no GPU), we successfully deconvolved a simulated 1-million-spot dataset in approximately 3 minutes. In contrast, probabilistic methods encounter severe scalability barriers at much smaller scales: CARD developers report that Cell2Location could not be applied to Slide-seqV2 data (∼20,000 spots) due to computational burden [13], and benchmark studies observe memory errors when processing datasets of similar size on standard GPU hardware [36]. This efficiency enables interactive, iterative analysis of atlas-scale data without specialized hardware. Beyond raw throughput, linear complexity also opens systematic multi-resolution analysis—deconvolving the same tissue at multiple bin sizes to probe how spatial resolution affects biological inference—an investigation we pursue in Section 2 after first demonstrating practical utility on human cancer data.

This scalability advantage is rooted in two fundamental algorithmic design choices. First, FlashDeconv operates in a compressed sketch space rather than the full gene space. Standard regression-based methods scale with the number of genes G: NNLS's active-set algorithm requires $O(N \cdot G \cdot K^2)$ due to iterative QR decomposition over up to K variables, and RCTD's iteratively reweighted least squares involves Hessian computation over all cell type pairs, also yielding $O(N \cdot G \cdot K^2)$. In contrast, FlashDeconv solves the regression problem in a d-dimensional sketch space where $d \ll G$, reducing the complexity to $O(N \cdot d \cdot K)$ (Methods). With typical parameters $d = 512$ and $G \approx 20{,}000$, this dimensional reduction provides speedup compared to full-transcriptome methods. We note that methods using HVG pre-selection (typically 2,000 genes) would see a smaller speedup factor of approximately 4× from sketching alone. However, FlashDeconv's scalability advantage stems primarily from two additional factors: (1) precomputation of $H = X_{\text{sketch}} Y_{\text{sketch}}^T$ makes BCD iteration complexity $O(N \cdot K^2)$ independent of gene dimension (Supplementary Note 1), and (2) sparse graph Laplacian regularization scales as $O(N \cdot k)$ rather than $O(N^2)$ for dense kernel methods. Furthermore, leverage-score weighting provides substantial

accuracy advantages over HVG selection, particularly for rare cell types (Supplementary Fig. S7). These combined advantages explain our empirical benchmarks, where FlashDeconv consistently outperformed gene-space methods by 1-2 orders of magnitude.

Second, FlashDeconv employs sparse graph-based spatial regularization rather than dense covariance modeling. Methods like CARD that model spatial dependencies through dense $N \times N$ kernel matrices incur $O(N^2)$ memory costs and $O(N^2 \cdot K)$ per-iteration complexity for matrix operations, rendering them prohibitive for datasets with $N > 10{,}000$ spots. FlashDeconv instead constructs a sparse k-nearest-neighbor graph, where each spot connects to only $k = 6$ neighbors. The resulting graph Laplacian regularization scales as $O(N \cdot k)$ in both time and space, maintaining linear complexity even as N grows to millions. A detailed complexity comparison across methods is provided in Supplementary Table S1.

Atlas-scale deconvolution of human colorectal cancer at subcellular resolution

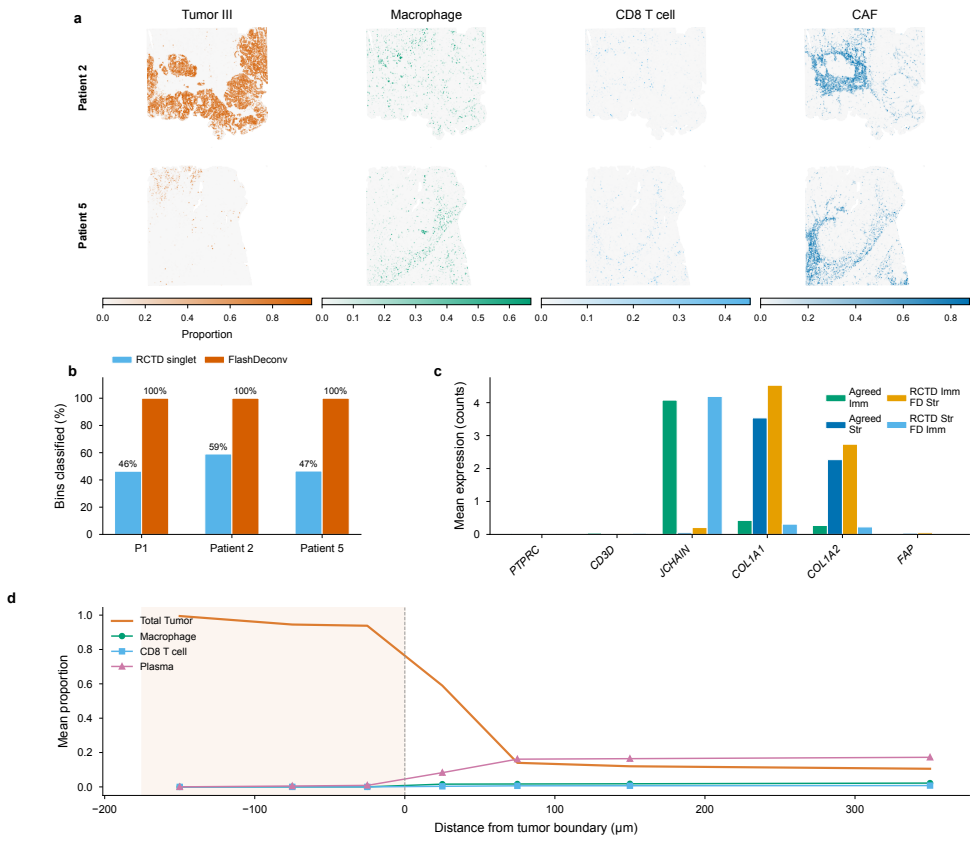

Fig. 4 Atlas-scale deconvolution of human colorectal cancer at subcellular resolution. **(a)** Spatial distribution of four key cell types (Tumor III, Macrophage, CD8 T cell, CAF) in two representative CRC patients (P2, P5) at 8 μm Visium HD resolution, showing heterogeneous tumor microenvironment architecture. **(b)** Coverage comparison between FlashDeconv and RCTD across three patients. FlashDeconv outputs continuous proportions for 100% of bins. RCTD, run in doublet mode following Oliveira et al. [37] (assigning one to two cell types per bin), classifies only 46–59% as singlets, with 7–19% as doublets and 5–7% rejected outright. At 8 μm subcellular resolution, bins are smaller than most cells and mRNA diffusion causes nearly every bin to contain mixed signals, a regime where discrete classification—even allowing two cell types—cannot capture the continuous mixing. **(c)** Independent marker gene validation on 71,769 disputed bins where FlashDeconv and RCTD disagree on immune vs. stromal classification. Canonical lineage markers (PTPRC, CD3D, JCHAIN for immune; COL1A1, COL1A2, FAP for stromal) confirm FlashDeconv's assignment in 19/22 gene-direction verdicts (86.4%), with unanimous agreement (11/11) when RCTD classifies as stromal but FlashDeconv assigns as immune. **(d)** Immune cell infiltration gradient as a function of signed distance from the tumor-stroma boundary, averaged across patients (shaded region = tumor interior). Total tumor proportion drops from ∼94% at −25 μm to ∼14% at +75 μm. Among immune populations, Plasma cells exhibit the steepest gradient (∼16% in peri-tumoral stroma), followed by macrophages (∼2%) and CD8 T cells (∼0.8%). Cancer-associated fibroblasts define a distinct stromal layer peaking at ∼14% near the boundary. Data: Visium HD CRC cohort from Oliveira et al. [37], 3 patients, 1,595,565 total bins at 8 μm resolution, 38 cell types. Processing time: 153 seconds for the entire cohort (∼10,400 bins/s).

To demonstrate FlashDeconv's utility for atlas-scale human cancer analysis, we applied it to Visium HD data from a colorectal cancer (CRC) cohort [37] comprising three patients with 1,595,565 total bins at 8 μm subcellular resolution and 38 cell types. FlashDeconv processed the entire cohort in 153 seconds at a throughput of approximately 10,400 bins per second, generating spatially resolved cell type proportion maps (Fig. 4a).

A key advantage of FlashDeconv's regression framework emerges at subcellular resolution, where discrete classification methods break down. Oliveira et al. [37] ran RCTD in doublet mode—the recommended setting for high-resolution platforms, which assigns each bin to one or two cell types. Even with this more permissive classification, RCTD rejects 5–7% of bins outright and classifies only 46–59% as singlets across the three patients, with the remainder labeled as doublet-certain or doublet-uncertain (Fig. 4b). At 8 μm resolution, individual bins are smaller than most cells, and mRNA diffusion causes nearly all bins to contain mixed signals from multiple cell types—a regime where even two discrete labels cannot capture the underlying continuous mixing. FlashDeconv's regression framework outputs continuous proportion estimates across all 38 cell types for every bin, achieving 100% coverage without discarding any measurements.

To independently validate accuracy on the 71,769 bins where FlashDeconv and RCTD disagree on immune versus stromal classification, we examined expression of canonical lineage markers (PTPRC, CD3D, JCHAIN, COL1A1, COL1A2, FAP) that serve as ground truth independent of either method's output (Fig. 4c). FlashDeconv's assignment was confirmed as correct in 19 of 22 gene-direction verdicts (86.4%). Notably, in the 11 cases where RCTD classified bins as stromal but FlashDeconv assigned them as immune, marker gene expression unanimously supported FlashDeconv—these bins expressed immune markers (PTPRC, CD3D, JCHAIN) at levels comparable to agreed-immune bins.

Analysis of immune cell infiltration as a function of signed distance from the tumor-stroma boundary revealed a sharp transition (Fig. 4d). Immune cell types were essentially absent inside the tumor and rose within the first 50 μm of the stromal side. Plasma cells exhibited the steepest gradient, accumulating to $\sim$16% in the peritumoral stroma, consistent with their known enrichment at the tumor-stroma interface in microsatellite-stable CRC [37]. Cancer-associated fibroblasts (CAFs) peaked at 50–100 μm outside the boundary ($\sim$14%), defining a fibroblast-rich peri-tumoral zone. CD8 T cells showed the weakest infiltration ($\sim$0.8% in distant stroma), consistent with the limited T-cell infiltration characteristic of microsatellite-stable colorectal tumors [37]. These continuous spatial gradients, detectable only at subcellular resolution, recapitulate the immune infiltration patterns described in the Immunoscore literature [38], demonstrating that FlashDeconv can resolve the spatial organization of immune populations at scales relevant to established tissue classification frameworks.

Beyond global infiltration patterns, FlashDeconv's complete tissue coverage revealed discrete neutrophil inflammatory microdomains at the tumor–stroma interface in all three patients (Fig. 5a). These aggregates—bins with $\geq$10% Neutrophil proportion forming spatially coherent clusters—exhibited the strongest spatial self-enrichment of any cell type in the 38-type dataset (16–56$\times$). Analysis of the

surrounding microenvironment uncovered two distinct spatial contexts (Fig. 5b): stromal-resident aggregates were surrounded by a vascularized innate immune niche enriched for macrophages, mRegDC, and mast cells, while tumor-proximal aggregates showed broad immune depletion—a dichotomy consistent across all three patients. Six neutrophil-specific marker genes independently confirmed the identity of these microdomains (11–63-fold enrichment; Fig. 5c), with the signal amplifying rather than diminishing when restricting to high-UMI bins ($\geq$200), ruling out low-quality deconvolution artifacts (Fig. 5d). The niche structure was most pronounced at 8 μm resolution, decaying at coarser scales (Fig. 5e; Supplementary Note 9). These microdomains are largely invisible to classification-based methods: across three patients, RCTD correctly identified only 2.3% of the 16,827 hotspot bins as Neutrophil singlets, with the majority (61%) receiving no classification and the remainder assigned to incorrect cell types—a direct consequence of forcing discrete labels onto bins containing continuous cell type mixtures.

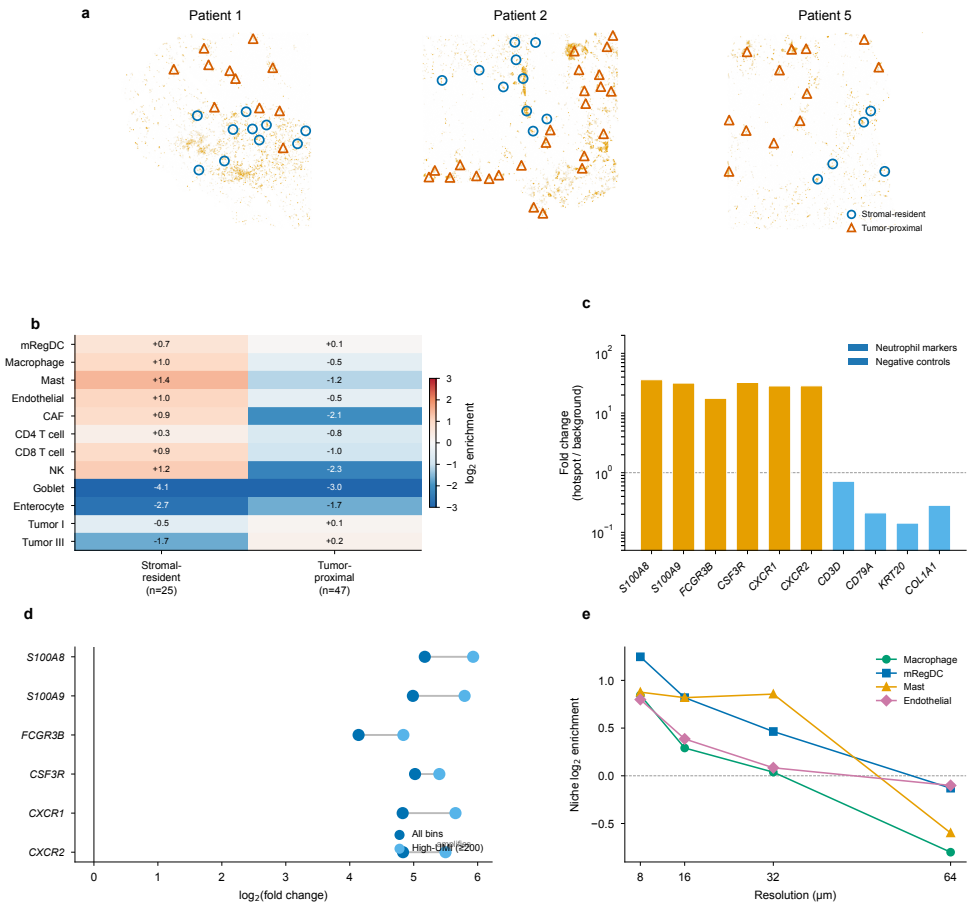

Fig. 5 Neutrophil inflammatory microdomains at the CRC tumor–stroma interface reveal two distinct spatial contexts. (a) Spatial distribution of neutrophil hotspot bins (≥10% proportion) across three patients. Aggregate centroids classified as stromal-resident (blue circles) or tumor-proximal (orange triangles) based on niche tumor content. **(b)** Niche composition of the two aggregate types (median $\log_2$ enrichment pooled across patients). Stromal-resident aggregates are surrounded by a vascularized innate immune niche (Macrophage, mRegDC, Mast, Endothelial enriched), while tumor-proximal aggregates show broad immune depletion. **(c)** Marker gene validation. Neutrophil-specific genes show 11–63-fold enrichment in hotspot versus background bins; negative control genes show no enrichment or depletion. Y-axis: log-scale fold change. **(d)** UMI-stratified robustness. Log_2 fold change of neutrophil marker genes in all hotspot bins versus high-UMI (≥200) bins, averaged across patients. The enrichment signal amplifies in high-UMI bins, ruling out low-quality deconvolution artifact. **(e)** Multi-resolution niche stability. Enrichment of key co-localized cell types decays from 8 to 64 μm, demonstrating that the microenvironment structure requires sub-cellular resolution to resolve. Data: Visium HD CRC cohort [37], 1,595,565 total bins, 38 cell types.

To orthogonally validate these deconvolution estimates, we leveraged Xenium in situ sequencing data from an adjacent serial section of the same tissue (Patient 1 [37]), which provides single-cell-resolution ground truth independent of any deconvolution

model. We annotated 289,352 Xenium cells by correlation with the scRNA-seq reference, then computationally aggregated them into spatial bins of 8–128 μm to create pseudo-bulk mixtures with known composition. FlashDeconv recovered lineage-level proportions with Pearson $r > 0.88$ across all resolutions, even at 8 μm where bins contained a median of one cell (Supplementary Note 7, Supplementary Fig. S22). Comparison of global cell type proportions across platforms revealed that at the coarser lineage level (6 categories), both FlashDeconv and RCTD agreed well with the Xenium reference ($r = 0.96$ and $r = 0.94$, respectively), but at the individual cell-type level (38 types), FlashDeconv maintained strong correspondence ($r = 0.78$) while RCTD singlet-derived proportions showed no systematic agreement ($r = -0.02$)—consistent with the information loss inherent in collapsing continuous mixtures to discrete labels.

FlashDeconv's accuracy extends beyond CRC: application to human ovarian cancer recovers the cell type composition differences between response groups reported by Denisenko et al. [39] (Supplementary Fig. S21).

Scale-space analysis of Visium HD reveals resolution-dependent information loss

While the preceding CRC analysis demonstrates FlashDeconv's practical utility for atlas-scale tumor microenvironment mapping at subcellular resolution, the 8 μm Visium HD platform raises a fundamental question beyond any single biological application: at what resolution does cellular information begin to collapse? We leveraged FlashDeconv's computational efficiency to perform a systematic scale-space analysis that would be infeasible with existing methods: at 350,000 spots (8 μm resolution), probabilistic methods face prohibitive barriers—Cell2Location cannot handle datasets beyond ∼20,000 spots [13], and Stereoscope's >8-hour runtime on 10,000 spots [12] would extrapolate to over 11 days.

Using Visium HD data from mouse small intestine (10x Genomics), we deconvolved cell type proportions across five resolutions (8, 16, 32, 64, and 128 μm) using a matched scRNA-seq reference [40]. FlashDeconv processed 351,817 bins at 8 μm resolution in 12.0 seconds—a throughput of approximately 29,000 bins per second on commodity hardware (Fig. 6a). The complete five-resolution analysis (366,975 total bins) required only 14 seconds, enabling systematic exploration of the resolution-information trade-off.

The resolution horizon.

Our analysis revealed a pronounced transition between 8 μm and 16 μm resolution (Fig. 6b). At 8 μm, 61.5% of bins were dominated by a single cell type (proportion >80%), consistent with near-single-cell purity. This fraction collapsed to 13.3% at 16 μm—a 78% reduction—and continued declining to 3.9% at 128 μm. The normalized Shannon entropy of cell type mixing increased correspondingly, from 0.23 at 8 μm to 0.48 at 16 μm. This transition represents the boundary beyond which cellular identity becomes fundamentally obscured by spatial averaging. The 8–16 μm transition observed in mouse small intestine was independently reproduced in mouse colon using Xenium ground truth without deconvolution (Supplementary Note 4), suggesting this scale is a robust feature of intestinal epithelium. Notably, this transition coincides

with the characteristic diameter of intestinal epithelial cells (∼10–15 μm), consistent with the geometric intuition that the resolution horizon occurs when measurement bins transition from capturing predominantly single cells to averaging multiple cells. For tissues with substantially different cell sizes or spatial organization patterns, the resolution horizon would be expected to shift accordingly—highlighting that this is a tissue-dependent biophysical parameter rather than a universal constant.

Cell types exhibit distinct resolution sensitivities.

Different cell types showed markedly different vulnerabilities to spatial aggregation (Fig. 6c). Intestinal stem cells—localized to crypt bases with a characteristic niche size of approximately 10–15 μm—exhibited the steepest information loss: spatial autocorrelation (Moran's I) decreased by 70% from 8 μm to 128 μm. In contrast, Paneth cells, which form tight clusters at crypt bases, showed only 17% loss of spatial coherence across the same resolution range. These differences reflect the characteristic spatial scales of each cell type's tissue organization.

Validation through crypt-villus boundary analysis.

To confirm that the observed resolution effects reflect genuine biological structure, we analyzed the sharpness of crypt-villus boundaries—the transition zone between Paneth cell-rich crypts and enterocyte-covered villi (Fig. 6d). We traced 100 linear paths from crypt cores (Paneth >50%) to villus regions (enterocyte >70%) and computed the maximum gradient of enterocyte proportion along each path. At 8 μm, the mean gradient sharpness was 0.036 ± 0.029. At 16 μm, this decreased to 0.008 ± 0.003—a 77% reduction in boundary definition. At 32 μm and coarser, crypt cores meeting our criteria could no longer be identified. This quantitative boundary analysis confirms that the information loss corresponds to anatomical blurring rather than algorithmic artifacts.

Coarse binning induces spurious colocalization.

Beyond signal loss, spatial aggregation can fundamentally distort cell-cell relationship inference (Fig. 6e). Paneth cells and Goblet cells are both secretory epithelial types, yet occupy distinct spatial niches: Paneth cells reside exclusively at crypt bases, while Goblet cells distribute along the crypt-villus axis. At 8 μm resolution, this spatial segregation manifests as a weak negative correlation ($r = -0.12$, $p < 10^{-100}$, $N = 351{,}817$), reflecting their mutual exclusion at the cellular scale. However, as resolution coarsens, both cell types become mixed within the same measurement bins, inducing a spurious positive correlation that peaks at 64 μm ($r = +0.80$). This correlation sign flip—from negative to positive—represents a qualitative reversal in the apparent cell-cell relationship, a tissue-scale instance of the modifiable areal unit problem (MAUP) from spatial statistics [41], recently highlighted in the context of spatial transcriptomics [42]. The mechanism is geometric: at fine resolution, adjacent but distinct cell populations are resolved separately, but as bin size increases beyond cells' physical dimensions, they become averaged into the same measurement unit and appear artificially co-localized. Researchers analyzing data at conventional Visium resolution (55 μm) would observe strong colocalization and might erroneously conclude that Paneth and Goblet cells

share a common microenvironment or respond to similar spatial cues. To confirm that this phenomenon reflects spatial geometry rather than deconvolution methodology, we performed ground truth validation using Xenium in situ sequencing data from mouse colon, where exact cell positions are known without deconvolution. The ground truth itself exhibits the same pattern: spatially adjacent but distinct cell clusters show negative correlation at 8 μm ($r = -0.06$) but strong positive correlation at 128 μm ($r = +0.68$), confirming that the observed sign reversal reflects tissue geometry rather than data sparsity or deconvolution methodology (Supplementary Note 4; Supplementary Fig. S15). FlashDeconv's ability to process high-resolution data efficiently enables detection of such resolution-dependent artifacts, revealing the true spatial organization that coarse binning obscures.

Leverage scores explain preservation of rare cell signals.

To verify the mechanistic connection between FlashDeconv's mathematical design and the biological findings above, we confirmed that marker genes for the key cell types exhibit high leverage scores in the intestine reference (Supplementary Fig. S14). The canonical stem cell marker *Lgr5* ranks in the top 1% by leverage (rank #283 of 27,998 genes), despite ranking only in the top 10% by variance. Overall, stem cell markers average the 95th percentile by leverage, with 5 of 6 markers in the top 5%. This confirms that the 8 μm stem cell niche detection demonstrated above is a direct consequence of leverage-weighted sketching: rare cell type signals are preserved during dimensionality reduction precisely because their markers define unique directions in gene expression space, regardless of their overall expression magnitude.

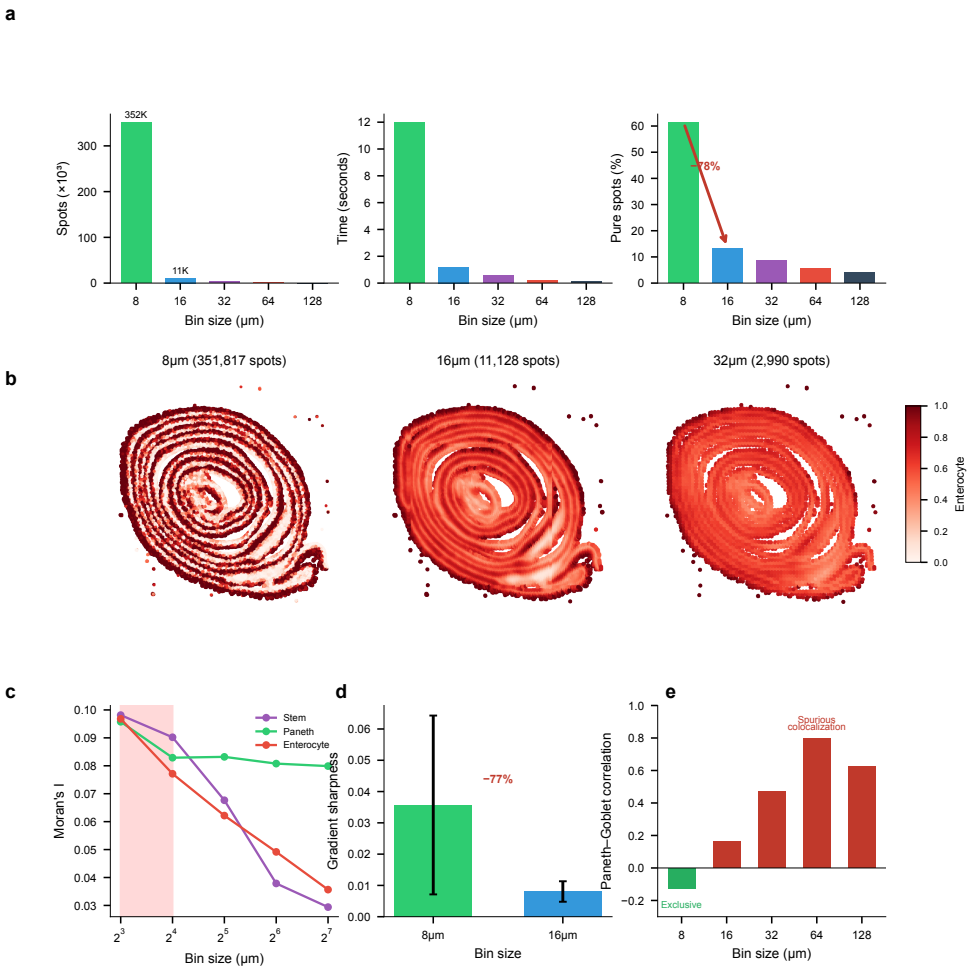

Fig. 6 Scale-space analysis of Visium HD reveals resolution-dependent information loss. **(a)** FlashDeconv performance across resolutions. Left: Number of measurement units at each bin size. Middle: Processing time demonstrating scalability to 350,000+ spots. Right: Signal purity (fraction of spots with >80% single cell type) collapses from 61.5% at 8 μm to 13.3% at 16 μm. **(b)** Spatial maps of enterocyte proportions at 8, 16, and 32 μm resolution on mouse small intestine. Fine anatomical detail visible at 8 μm becomes progressively blurred at coarser resolutions. **(c)** Resolution sensitivity varies by cell type. Stem cells (red) show the steepest decline in spatial coherence (Moran's I), while Paneth cells (blue) retain spatial structure. Shaded region indicates the 8–16 μm transition zone. **(d)** Crypt-villus boundary validation. Gradient sharpness decreases by 77% from 8 μm to 16 μm, quantifying anatomical blurring. **(e)** Spatial binning induces spurious colocalization. Paneth and Goblet cells show weak mutual exclusion at 8 μm ($r = -0.12$, $p < 10^{-100}$) but appear strongly colocalized at 64 μm ($r = +0.80$, $p < 10^{-100}$)—a correlation sign flip that could lead to incorrect biological conclusions about cell-cell interactions. Data: Visium HD Mouse Small Intestine (10x Genomics), scRNA-seq reference from Haber et al. 2017.

FlashDeconv resolves the Tuft-Stem chemosensory niche.

Beyond validating marker preservation, FlashDeconv enabled niche-resolution quantification of cell-type co-localization directly from capture-based spatial transcriptomics—recovering a spatial architecture consistent with known Tuft cell biology [43, 44] through deconvolution of sequencing-based platforms without requiring imaging-based cell segmentation. Tuft (brush) cells—rare chemosensory epithelial cells comprising only ∼0.4–2% of intestinal epithelium [45]—exhibited the highest "HVG blindness" among all cell types in our analysis (Fig. 7a): their marker genes (including *Pik3r5*, *Ptgs1*) rank 21 percentile points lower under variance-based selection than under leverage-based ranking, placing them at highest risk of being discarded by standard feature selection pipelines. At 8 μm resolution, FlashDeconv identified 2,244 focal Tuft cell niches with proportions reaching 61%—near-pure Tuft cell spots (Fig. 7b). These niches exhibited a striking spatial pattern: a 15.3-fold enrichment for intestinal stem cells, and 14.0-fold enrichment for enteroendocrine cells (Fig. 7c,e). Conversely, differentiated cell types were strongly depleted: enterocytes (0.11×), goblet cells (0.10×), and Paneth cells (0.22×)—all significantly non-random (permutation test, $p < 10^{-4}$, $n = 10{,}000$). This Tuft-Stem co-localization is anatomically consistent with Tuft cells' known localization near the intestinal stem cell zone at the crypt base [43, 44] and their recently discovered capacity to function as reserve stem cells following epithelial injury [46]. Notably, exploratory ligand-receptor analysis revealed that *Il17ra*—encoding a subunit of the IL-25 receptor through which Tuft cells signal to neighboring cells [47]—is expressed 7-fold higher in Tuft-Stem niches than in background tissue (Supplementary Fig. S20), suggesting that this spatial proximity may facilitate paracrine communication. Null model validation confirms this signal is biological rather than artifactual: independent marker gene sets show consistent spatial patterns ($r = 0.30$, $p < 10^{-200}$), co-localization is specific to stem cells (not random cell types), Tuft cell distribution exhibits significant spatial autocorrelation (Moran's I = 0.44 vs. random baseline ∼0), and the Tuft-Stem proximity is recapitulated directly in raw gene expression without deconvolution (*Pou2f3*–*Lgr5* neighbor enrichment, $p < 0.001$; Supplementary Note 6, Supplementary Fig. S19). The Tuft cell signal is exquisitely resolution-dependent: maximum Tuft proportion decreased from 61% at 8 μm to 4% at 128 μm (Fig. 7d), rendering these niches undetectable at conventional Visium resolution. This finding illustrates that resolving such niche architecture from capture-based spatial transcriptomics requires overcoming two independent barriers: rare cell markers lost to variance-based feature selection, and spatial signal diluted at resolutions exceeding the niche's characteristic scale (∼10–15 μm).

Cross-method comparison confirms niche architecture.

To assess whether this spatial architecture reflects genuine tissue biology or is specific to FlashDeconv's design, we applied RCTD—a likelihood-based method that treats each spot independently—to the same tissue at 16 μm resolution, the finest scale at which RCTD's per-spot inference remains practical on this tissue (10,952 spots; 8 μm would require >350,000 independent regressions). Both methods detected stem cell enrichment in Tuft cell hotspots, confirming that the niche reflects tissue architecture rather than a regularization artifact. However, FlashDeconv produced

markedly more spatially coherent assignments, consolidating Tuft cell signal into 9 clusters versus 168 scattered fragments for RCTD, with 7.7-fold stem cell enrichment at hotspots compared to 3.0-fold (Supplementary Fig. S24). This difference partly reflects FlashDeconv's graph Laplacian regularization, which borrows information from spatial neighbors to enforce local continuity. To verify that this coherence corresponds to genuine biology rather than over-smoothing, we examined *Il17ra*—encoding a subunit of the IL-25 receptor critical for Tuft-Stem paracrine signaling [47]—as a functional readout independent of both deconvolution methods. FlashDeconv-defined niche zones showed 5.2-fold *Il17ra* enrichment over background versus 1.5-fold for RCTD-defined zones, indicating that the spatially coherent assignments more accurately delineate functional signaling microenvironments (Supplementary Note 8).

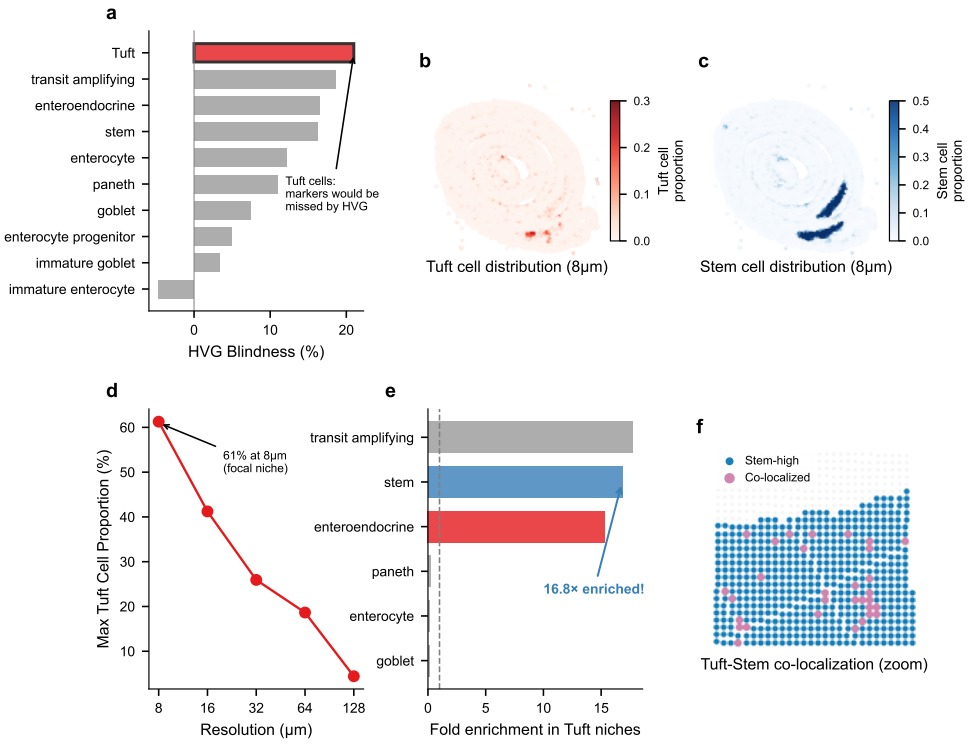

Fig. 7 FlashDeconv resolves the Tuft-Stem chemosensory niche. (a) HVG blindness ranking across intestinal cell types. HVG blindness is defined as the difference in mean percentile rank of a cell type's marker genes under variance-based versus leverage-based selection; positive values indicate systematic underweighting by HVG. Tuft (brush) cells exhibit the highest HVG blindness (21 percentile points). **(b)** Spatial distribution of Tuft cells at 8 μm resolution reveals focal niches (red spots) with proportions up to 61%. **(c)** Stem cell distribution at 8 μm shows concentration at crypt bases. **(d)** Resolution sensitivity of Tuft cell detection. Maximum proportion decreases from 61% (8 μm) to 4% (128 μm), rendering focal niches undetectable at conventional resolution. **(e)** Co-localization analysis reveals Tuft cell hotspots are enriched 15.3-fold for stem cells and 14.0-fold for enteroendocrine cells ($p < 10^{-4}$, permutation test), but depleted for differentiated cell types (enterocytes 0.11×, goblet cells 0.10×). **(f)** Spatial zoom showing Tuft-Stem co-localization at crypt bases (blue: Stem-high, pink: co-localized). Tuft cells rarely appear without adjacent stem cells, consistent with their intimate niche association. Data: Visium HD Mouse Small Intestine (10x Genomics).

Discussion

In this study, we present FlashDeconv, a deconvolution framework designed for atlas-scale spatial transcriptomics. By combining leverage-score importance sampling with sparse graph regularization, FlashDeconv achieves accuracy comparable to probabilistic methods while reducing computational cost by orders of magnitude. The concept of statistical leverage originates from regression diagnostics, where it quantifies the influence of individual observations on fitted models, and was later shown to enable

efficient matrix approximation via importance sampling in randomized numerical linear algebra [48, 49]. In single-cell biology, leverage-based sampling has been applied to subsample cells while preserving rare populations [50, 51]. In the gene space, leverage scores have been used for feature identification via CUR decomposition [52] and for variable screening in spatial transcriptomics [53]. FlashDeconv extends this principle by integrating leverage-score importance weights directly into a CountSketch randomized projection for spatial deconvolution, addressing the complementary challenge that variance-based feature selection systematically discards markers of rare cell types. Despite the inherent randomness in sketching, FlashDeconv produces highly reproducible results: across 10 runs with different random seeds, pairwise Pearson correlations exceeded 0.98, confirming that users can obtain consistent results without fixing random seeds.

The core insight underlying our approach is the distinction between geometric structure and statistical variance. Variance-based methods—including PCA and highly variable gene selection—conflate biological signal with population frequency, systematically underweighting rare populations regardless of their biological importance. Leverage scores instead capture each gene's contribution to the discriminative structure among cell types, independent of expression magnitude or population size. Our experiments confirm this decoupling: leverage ranking remains stable under 67-fold changes in cell type abundance, and the structurally informative (GOLD) gene set shows coherent functional enrichment and clear spatial organization, while variance-dominated (NOISE) genes lack both—indicating that high variance alone does not ensure biological relevance.

Our choice of Log-CPM normalization reflects an engineering trade-off rather than a claim of statistical superiority. While Pearson residuals are statistically principled for negative binomial count data, Log-CPM provides specific advantages for L_2-based sketching: its bounded norm prevents high-expression genes from dominating the compressed space. This pragmatic choice, combined with leverage-weighted sketching, allows FlashDeconv to preserve essential biological signals during dimensionality reduction.

Our case studies reveal three complementary failure modes in spatial deconvolution that FlashDeconv addresses through distinct mechanisms. The melanoma dataset represents a feature-space challenge: highly correlated malignant states ($r > 0.98$) along a continuous phenotypic spectrum create an ill-conditioned regression problem. Here, sparse NNLS naturally selects the dominant signals present in the tissue without requiring manual reference curation—a property particularly valuable for atlas-scale mapping where exhaustive reference curation is impractical. The Visium HD mouse intestine dataset represents a signal-space challenge: at 8 μm resolution with sparse counts, rare cell markers risk being discarded during dimensionality reduction. Leverage-weighted sketching preserves these geometrically distinctive signals, enabling detection of the Tuft-Stem chemosensory niche—spatial architecture that requires both rare-marker-aware feature selection and sub-16 μm resolution to resolve. Cross-method comparison with RCTD confirms that the niche is detectable by either approach, but highlights the role of spatial regularization in resolving niche architecture: spot-independent methods that treat each location in isolation disperse

the signal across the tissue, producing niche zones with weaker enrichment for the paracrine signaling gene *Il17ra* than spatially regularized estimates (Supplementary Note 8). This intimate spatial proximity between Tuft cells and intestinal stem cells aligns with recent findings that Tuft cells can dedifferentiate into functional stem cells during regeneration [46], suggesting that the niche architecture we observe may facilitate both paracrine signaling and direct cellular interconversion. Quantitatively, our measurements—15.3-fold stem cell enrichment within Tuft cell foci at ∼10–15 μm resolution—are consistent with lineage tracing and imaging studies that localize Tuft cells to the Lgr5-positive crypt base compartment [43, 44] and with the spatial range required for IL-25 paracrine signaling [47]. FlashDeconv's contribution here is not the discovery of this proximity per se, but its quantitative characterization from capture-based spatial transcriptomics—recovering a spatial architecture previously accessible only through imaging or lineage tracing, and demonstrating that it requires sub-16 μm resolution to resolve.

The CRC Visium HD analysis reveals a third challenge specific to subcellular platforms: classification-based methods that assign each bin a discrete cell type label fail when measurement bins contain mixed mRNA from multiple cells. Even RCTD's doublet mode—which allows up to two cell types per bin—classifies only 46–59% of bins as singlets and rejects 5–7% outright, because discrete labels cannot capture the continuous mixing inherent at 8 μm resolution. FlashDeconv's regression framework achieves 100% bin coverage with continuous proportions, and independent marker gene validation confirms accuracy in >86% of disputed assignments. The biological significance of this coverage gap is illustrated by the neutrophil microdomain analysis (Fig. 5): RCTD correctly classifies only 2.3% of the 16,827 hotspot bins as Neutrophil singlets, rendering the organized niche structure—neutrophils co-localized with immunoregulatory mRegDC [54], macrophages, mast cells, and endothelial cells at vascularized stromal sites—entirely undetectable. The functional consequences of this spatial architecture for antigen presentation and immune regulation warrant further investigation (Supplementary Note 9). This regression-versus-classification distinction becomes increasingly critical as spatial transcriptomics moves toward subcellular resolution, where nearly every measurement unit captures signal from multiple cell types. No single mechanism addresses all three challenges; FlashDeconv's combination of structure-preserving compression, signal-driven inference, and continuous regression provides robustness across these qualitatively different failure modes.

Beyond accuracy, FlashDeconv's computational efficiency enables systematic scale-space exploration. The phenomenon of resolution-dependent information loss upon spatial aggregation—including correlation sign reversal—is well established in spatial statistics as the modifiable areal unit problem (MAUP) [41], and has recently been identified as relevant to spatial transcriptomics [42]. Our multi-resolution analysis provides systematic, tissue-specific quantification of this effect in high-resolution spatial transcriptomics, identifying what we call the *resolution horizon*: the spatial scale at which biological segregation collapses into geometric mixing. Signal purity drops precipitously between 8 and 16 μm, while cell-cell correlations undergo sign inversion (from $r = -0.12$ to $r = +0.80$ by 64 μm), as validated in Xenium ground truth. Critically, the resolution horizon is tissue-specific and cell-type-dependent: stem cells

require 8 μm resolution while Paneth cells tolerate 32 μm—a granularity not captured by prior theoretical treatments of MAUP. FlashDeconv's linear scalability transforms resolution selection from an arbitrary experimental choice into a data-driven determination.

Our sparse graph Laplacian regularization provides coverage-dependent benefits particularly relevant for emerging high-resolution platforms (Supplementary Note). At low sequencing depth typical of Visium HD 2 μm bins, spatial regularization improves accuracy by borrowing information from neighboring spots, reducing sampling noise without introducing excessive bias. As spatial transcriptomics scales to subcellular resolution with increasingly sparse measurements, this adaptive regularization becomes increasingly important.

We acknowledge scenarios where alternative approaches may be preferable. For small datasets (<50 spots), sketching and spatial smoothing provide limited benefit. For resolving fine-grained cell states along continuous phenotypic spectra, probabilistic models with explicit uncertainty quantification remain valuable. When platform-specific batch effects dominate biological signal, methods with explicit noise modeling may be more robust, although leverage scores provide inherent robustness to batch effects through reference anchoring (Supplementary Note). A related limitation is reference incompleteness: when the spatial tissue contains cell types absent from the scRNA-seq reference, deconvolution estimates are necessarily biased. However, FlashDeconv's unnormalized regression framework provides an implicit out-of-distribution detection mechanism through reconstruction residuals, which can flag regions of potential reference incompleteness (Supplementary Fig. S18). We view FlashDeconv not as a universal replacement, but as a purpose-built tool for million-spot atlases where linear scalability is essential.

More broadly, this work illustrates that measuring geometric structure rather than statistical variance can decouple biological importance from numerical prevalence—a principle that may prove valuable in other computational biology settings where rare populations carry disproportionate functional significance. As spatial biology scales to subcellular resolution, feature selection strategies that preserve critical signals regardless of abundance will be essential to ensure computational efficiency does not compromise biological discovery.

Methods

Algorithm Overview

Algorithm 1 presents the complete FlashDeconv pipeline, which consists of five main stages: gene selection, data preprocessing, structure-preserving sketching, spatial graph construction, and optimization via block coordinate descent.

Two-Stage Feature Selection

FlashDeconv employs a two-stage feature selection strategy that combines variance-based noise filtering with reference-guided marker inclusion, followed by leverage-based

Algorithm 1 FlashDeconv: Structure-Preserving Spatial Deconvolution

Require: Spatial count matrix $Y \in \mathbb{R}^{N \times G}$, reference signatures $X \in \mathbb{R}^{K \times G}$, spatial coordinates $C \in \mathbb{R}^{N \times 2}$

Require: Parameters: d (sketch dimension), k (neighbors), λ (spatial regularization), ρ (L1 regularization)

Ensure: Cell type proportions $\beta \in \mathbb{R}^{N \times K}$

// Stage 1: Gene Selection
$\mathcal{G}_{\text{hvg}} \leftarrow \text{SelectHVG}(Y,\, n_{\text{hvg}} = 2000)$
$\mathcal{G}_{\text{marker}} \leftarrow \text{SelectMarkers}(X,\, n_{\text{marker}} = 50 \text{ per type})$
$\mathcal{G} \leftarrow \mathcal{G}_{\text{hvg}} \cup \mathcal{G}_{\text{marker}}$
$Y \leftarrow Y[:, \mathcal{G}],\; X \leftarrow X[:, \mathcal{G}]$
// Stage 2: Data Preprocessing
$\tilde{Y} \leftarrow \log(1 + 10^4 \cdot Y/\text{rowSum}(Y))$ ▷ Log-CPM transformation
$\tilde{X} \leftarrow \log(1 + 10^4 \cdot X/\text{rowSum}(X))$
// Stage 3: Structure-Preserving Sketching
$U, \Sigma, V \leftarrow \text{SVD}(\tilde{X}, r)$ ▷ Truncated SVD
$\ell_g \leftarrow \sum_{j=1}^{r} V_{gj}^2$ for all genes g ▷ Leverage scores
$p_g \leftarrow \ell_g / \sum_{g'} \ell_{g'}$ ▷ Normalize to probability
$w_g \leftarrow \sqrt{p_g \cdot G}$ ▷ Importance weights
$\Omega \leftarrow \text{BuildCountSketch}(G,\, d,\, \{w_g\})$ ▷ Weighted sketching matrix
$Y_{\text{sketch}} \leftarrow \tilde{Y}\Omega,\; X_{\text{sketch}} \leftarrow \tilde{X}\Omega$
// Stage 4: Spatial Graph Construction
$A \leftarrow \text{BuildKNNGraph}(C,\, k)$ ▷ k-NN adjacency matrix
$D \leftarrow \text{diag}(A\mathbf{1})$ ▷ Degree matrix
$L \leftarrow D - A$ ▷ Graph Laplacian
if $\lambda =$ "auto" **then**
  $\lambda \leftarrow 0.005 \cdot \text{mean}(\text{diag}(X_{\text{sketch}}^T X_{\text{sketch}}))/\text{mean}(\text{diag}(D))$
end if
// Stage 5: Optimization
$\beta \leftarrow \text{BlockCoordinateDescent}(Y_{\text{sketch}},\, X_{\text{sketch}},\, L,\, \lambda,\, \rho)$ ▷ Algorithm 2
// Normalize to proportions
$\beta \leftarrow \beta/\text{rowSum}(\beta)$ ▷ Normalize each spot to sum to 1
return β

importance weighting during sketching. This hybrid approach addresses the inherent tension between noise reduction and signal preservation: purely variance-based selection (e.g., PCA on all genes) would carry forward technical noise and housekeeping genes, while purely leverage-based selection on raw noisy data could amplify measurement artifacts.

Stage 1: Variance-based noise filtering (HVG selection). We first apply Highly Variable Gene (HVG) selection as a coarse filter to remove uninformative genes dominated by technical noise. Genes are binned by mean expression, and normalized dispersion is computed within each bin using the Seurat v3 method [32]. We select the top 2,000 genes with normalized dispersion exceeding 0.5 and mean expression

between 0.0125 and 3.0 (log-scale). This stage removes housekeeping genes with near-constant expression and lowly-expressed genes dominated by dropout noise, reducing the input space from $G \approx 20{,}000$ to $G' \approx 2{,}000$ genes. We validated this two-stage strategy by comparing it against whole-transcriptome leverage calculation (without HVG filtering): the two approaches showed 92.4% overlap in top-ranked leverage genes, with all discordant genes belonging exclusively to the high-variance NOISE quadrant (Supplementary Fig. S11). This confirms that HVG pre-filtering effectively removes technical artifacts without discarding biologically important high-leverage genes.

Stage 2: Reference-guided marker inclusion. Independently from HVG selection, we identify cell-type marker genes from the *whole transcriptome* of the reference matrix X. For marker identification, we compute for each gene the expression differential between its maximum expression across cell types and its second-highest expression. For each cell type, the top 50 genes where that type has the highest expression and largest differential are selected as markers. The final gene set is the union of HVGs and markers ($\mathcal{G} = \mathcal{G}_{\text{hvg}} \cup \mathcal{G}_{\text{marker}}$), typically $|\mathcal{G}| \approx 2{,}500$ genes. Leverage scores are then computed on this combined gene set for structure-preserving sketching. Critically, the subsequent sketching step (Section 2) applies leverage-score-based importance sampling to this filtered gene set, ensuring that rare cell type markers are preserved in the final $d = 512$ dimensional sketch.

This two-stage design avoids the "variance-only" pitfall where rare cell signals are lost during dimension reduction, while simultaneously preventing leverage scores from being computed on noisy, uninformative genes. The combination of variance filtering (for noise reduction) followed by leverage-based sketching (for structure preservation) enables FlashDeconv to compress $G \approx 20{,}000$ genes to $d = 512$ dimensions without sacrificing rare cell type detection accuracy (Algorithm 1, lines 2–5, 11–13). Empirical validation on liver scRNA-seq data confirms that leverage-based importance sampling successfully enriches rare cell type markers by 2.5-fold compared to variance-based methods (Fig. 2).

Data Preprocessing

We evaluated multiple variance-stabilizing transformations—including Pearson residuals commonly used for negative binomial data—and selected Log-CPM as the default. Our signal-to-noise analysis demonstrates that Log-CPM provides superior properties for sketching-based methods by avoiding both the high-expression saturation and low-expression noise amplification inherent to Pearson residuals (Supplementary Note 2). Let $Y \in \mathbb{R}^{N \times G}$ be the spatial count matrix. We normalize by library size and log-transform:

$$\tilde{Y}_{ig} = \log\left(1 + 10^4 \cdot \frac{Y_{ig}}{\sum_g Y_{ig}}\right) \tag{1}$$

The reference matrix $X \in \mathbb{R}^{K \times G}$ is transformed identically to ensure the regression problem $Y \approx \beta X$ operates in a consistent feature space.

Structure-Preserving Randomized Sketching

Randomized sketching compresses high-dimensional data via random projections—instead of computing expensive eigendecompositions (as in PCA), we randomly assign each gene to one of d lower-dimensional feature groups, achieving similar dimensionality reduction at a fraction of the computational cost. To reduce the dimensionality of the regression problem from G genes to d features ($d \ll G$), we construct a sparse sketching matrix $\Omega \in \mathbb{R}^{G \times d}$ using a leverage-score-weighted CountSketch transform [55].

First, we compute the statistical leverage scores ℓ_g from the singular value decomposition of the transformed reference $\tilde{X} \in \mathbb{R}^{K \times G}$:

$$\ell_g = \sum_{j=1}^{r} (V_{gj})^2 \tag{2}$$

where $V \in \mathbb{R}^{G \times r}$ contains the right singular vectors of $\tilde{X}$ (corresponding to genes) and r is the numerical rank. Genes with high leverage scores correspond to discriminative markers that distinguish between cell types; for example, in liver scRNA-seq data, *Rspo3* achieves the highest leverage score among all genes as a marker of Central Vein Endothelial cells (2% of cells), while abundant Hepatocyte markers exhibit lower leverage due to transcriptional overlap with other lineages.

We then construct Ω as follows. For each gene g, we assign it to exactly one sketch dimension $h(g) \in \{1, \ldots, d\}$ via a random hash function, and set all other entries to zero:

$$\Omega_{g,j} = \begin{cases} s_g \cdot w_g & \text{if } j = h(g) \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where $s_g \in \{-1, +1\}$ is a random sign (Rademacher variable) and $w_g = \sqrt{p_g \cdot G}$ is a leverage-based importance weight, with $p_g = \ell_g / \sum_{g'} \ell_{g'}$ denoting the normalized leverage probability. This weighting ensures that high-leverage genes (often marker genes with low total abundance but high discriminative power) contribute proportionally more to the sketch, preserving their signal in the low-dimensional projection. This amplification is critical because CountSketch hashes each gene to exactly one of d dimensions: when a rare cell marker and a high-abundance housekeeping gene collide (are assigned to the same dimension), their contributions are summed, and without importance weighting the abundant gene's magnitude can overwhelm the marker signal. Leverage-proportional scaling provides asymmetric amplification—empirically, structurally informative genes (low variance, high leverage) exhibit 3–12× higher leverage-to-magnitude ratios than variance-dominated genes across six tissues (mean 6.2×; Fig. 2b)—ensuring that rare cell type signals survive hash collisions while technical noise is relatively attenuated. This selective weighting mitigates hash collision variance [27] by maintaining adequate signal-to-interference ratio in the compressed space. Finally, columns of Ω are normalized to have unit ℓ_2 norm, scaled by $\sqrt{G/d}$ to approximately preserve Frobenius norms.

The data are then projected into the sketch space:

$$Y_{\text{sketch}} = \tilde{Y}\Omega, \quad X_{\text{sketch}} = \tilde{X}\Omega \tag{4}$$

This allows us to solve the deconvolution problem in the d-dimensional sketch space (default $d = 512$; validated in Supplementary Fig. S1) with theoretically bounded approximation error [26, 56, 57]. Since Y primarily lies in the column space of X (i.e., $Y \approx X\beta + \epsilon$), leverage scores computed from X effectively capture the geometric structure of both matrices, ensuring that Ω preserves not only X but also the projection of Y onto the biological subspace spanned by reference cell types.

Spatial Regularization and Optimization

Spatial Graph Construction. Given spatial coordinates $C \in \mathbb{R}^{N\times 2}$, we construct a k-nearest neighbor graph to encode spatial proximity. For each spot i, let $\mathcal{N}_k(i)$ denote its k nearest neighbors under Euclidean distance. The binary adjacency matrix $A \in \{0,1\}^{N\times N}$ is defined as:

$$A_{ij} = \begin{cases} 1 & \text{if } j \in \mathcal{N}_k(i) \text{ or } i \in \mathcal{N}_k(j) \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

The symmetrization ensures an undirected graph. We use $k = 6$ by default, which matches the hexagonal packing geometry of standard Visium arrays; for square-grid platforms such as Visium HD, $k = 6$ captures the four cardinal plus two closest diagonal neighbors. Sensitivity analysis confirms that performance is robust across $k \in \{4, 6, 8, 12, 20\}$ (Supplementary Fig. S2a, RMSE variation $< 0.1\%$), making this default applicable across diverse spatial geometries. The graph Laplacian is then $L = D - A$, where D is the diagonal degree matrix with $D_{ii} = \sum_j A_{ij}$.

Optimization Problem. We formulate the deconvolution task as a graph-regularized non-negative least squares problem:

$$\min_{\beta\geq 0} \frac{1}{2}\|Y_{\text{sketch}} - \beta X_{\text{sketch}}\|_F^2 + \lambda \text{Tr}(\beta^T L\beta) + \rho\|\beta\|_1 \tag{6}$$

where $\|\beta\|_1 = \sum_{i,k} |\beta_{ik}|$ denotes the element-wise ℓ_1 norm promoting sparsity.

Notably, we solve for unnormalized cell type abundances $\beta \geq 0$ rather than directly optimizing proportions on the probability simplex ($\sum_k \beta_{ik} = 1$). This formulation is better understood as *cellular density estimation*: β_{ik} represents the absolute abundance of cell type k at spot i, not its relative proportion. This approach offers three advantages: (1) it simplifies optimization to standard non-negative least squares; (2) the spatial Laplacian term encourages similar *total* cell densities across neighboring spots, a physically meaningful constraint; and (3) the unnormalized sum $\sum_k \beta_{ik}$ implicitly captures spot-level capture efficiency variation, which would otherwise be forced into proportion estimates under simplex constraints (Supplementary

Fig. S13). Cell type proportions are obtained via post-hoc row normalization, equivalent to maximum likelihood estimation under a multinomial model with unknown scale (Algorithm 1, line 20).

Scale-Invariant Regularization. To ensure consistent regularization across datasets with varying signal magnitudes, both λ and ρ are internally scaled relative to the Gram matrix $G = X_{\text{sketch}} X_{\text{sketch}}^T$. Let $\bar{G} = \text{mean}(\text{diag}(G))$ denote the mean Gram diagonal, which characterizes the data fidelity gradient scale. For spatial regularization:

$$\lambda = \alpha \cdot \frac{\bar{G}}{\text{mean}(\text{degree}(A))} \tag{7}$$

where $\alpha \in [0, 1)$ is a dimensionless tuning coefficient (default 0.005) that controls the relative contribution of the spatial term to the Hessian diagonal. For sparsity regularization, the partial residual r_{ik} in the BCD update scales as $O(\bar{G})$, so the soft-thresholding parameter must be commensurate:

$$\tilde{\rho} = \rho \cdot \bar{G} \tag{8}$$

This makes the tuning coefficients α and ρ dimensionless fractions independent of sequencing depth or spot density; the resulting λ and $\tilde{\rho}$ inherit the scale of the data automatically. Sensitivity analysis confirms robust performance across varying regularization strengths (Supplementary Fig. S9). The problem is solved using a fast Block Coordinate Descent (BCD) algorithm implemented in Python with Numba acceleration.

Block Coordinate Descent Solver

We solve the optimization problem via coordinate descent, updating one element β_{ik} at a time while holding all others fixed. Let $G = X_{\text{sketch}} X_{\text{sketch}}^T \in \mathbb{R}^{K \times K}$ denote the precomputed Gram matrix. The partial derivative of the objective with respect to β_{ik} (ignoring the ℓ_1 term and non-negativity constraint) is:

$$\frac{\partial \mathcal{L}}{\partial \beta_{ik}} = -\left(X_{\text{sketch}} Y_{\text{sketch}}^T\right)_{ki} + \sum_{j=1}^{K} G_{kj} \beta_{ij} + \lambda \cdot 2(L\beta)_{ik} \tag{9}$$

The spatial regularization term expands as $(L\beta)_{ik} = d_i \beta_{ik} - \sum_{n \in \mathcal{N}(i)} \beta_{nk}$, where $d_i = |\mathcal{N}(i)|$ is the degree of spot i. Setting the gradient to zero and solving for β_{ik} yields the closed-form update:

$$\beta_{ik} \leftarrow \frac{\left(X_{\text{sketch}} Y_{\text{sketch}}^T\right)_{ki} - \sum_{j \neq k} G_{kj} \beta_{ij} + \lambda \sum_{n \in \mathcal{N}(i)} \beta_{nk}}{G_{kk} + \lambda d_i} \tag{10}$$

To incorporate ℓ_1 regularization and non-negativity, we apply the proximal operator for the non-negative ℓ_1 penalty: $\beta_{ik} \leftarrow \max(0, \mathcal{S}_{\tilde{\rho}}(r_{ik}) / (G_{kk} + \lambda d_i))$, where $\mathcal{S}_{\tilde{\rho}}(x) = \text{sign}(x) \cdot \max(|x| - \tilde{\rho}, 0)$ is the soft-thresholding operator with the scale-invariant threshold $\tilde{\rho} = \rho \cdot \bar{G}$ (Section 2). Algorithm 2 details the complete procedure [58]. Importantly,

Algorithm 2 Block Coordinate Descent for Spatial Deconvolution

Require: Sketched data $Y_{\text{sketch}} \in \mathbb{R}^{N \times d}$, sketched reference $X_{\text{sketch}} \in \mathbb{R}^{K \times d}$
Require: Graph Laplacian $L \in \mathbb{R}^{N \times N}$, regularization parameters λ, ρ, tolerance $\tau = 10^{-4}$
Ensure: Cell type abundances $\beta \in \mathbb{R}^{N \times K}$

```
// Precomputation
G ← X_sketch X_sketch^T                                  ▷ Gram matrix ∈ R^{K×K}
ρ̃ ← ρ · mean(diag(G))                                    ▷ Scale-invariant sparsity threshold
Extract neighbor indices N(i) from L for all spots i
// Initialization
β ← 0_{N×K}                                              ▷ Initialize with zeros
t ← 0                                                    ▷ Iteration counter
// Iterative Optimization
repeat
    β^old ← β
    for i = 1 to N do                                    ▷ Iterate over all spots
        N_i ← N(i)                                       ▷ Neighbors of spot i
        d_i ← |N_i|                                      ▷ Degree of spot i
        for k = 1 to K do                                ▷ Iterate over all cell types
            // Compute partial residual
            r_ik ← (X_sketch Y_sketch^T)_ki              ▷ Data fit term
            r_ik ← r_ik − Σ_{j≠k} β_ij · G_jk            ▷ Subtract other cell types
            r_ik ← r_ik + λ Σ_{n∈N_i} β_nk               ▷ Spatial smoothing term
            // Soft-thresholding with non-negativity
            s ← sign(r_ik) · max(|r_ik| − ρ̃, 0)          ▷ S_ρ̃(r_ik)
            β_ik ← max(0, s / (G_kk + λ d_i))            ▷ Update with projection
        end for
    end for
    // Check convergence
    δ ← ‖β − β^old‖_F / ‖β^old‖_F                        ▷ Relative change
    t ← t + 1
until δ < τ or t ≥ t_max                                 ▷ Converged or max iterations
return β
```

our objective decomposes as $f(\beta) + g(\beta)$, where f (data fidelity and Laplacian terms) is smooth and convex, and g (ℓ_1 penalty and non-negativity constraint) is separable and convex. For this class of composite convex problems, block coordinate descent with proximal updates is guaranteed to converge to the global optimum [59]; empirical validation on real biological data is provided in Supplementary Fig. S3.

The key computational advantages of this algorithm are: (1) the Gram matrix G is precomputed once with size $K \times K$ (typically $K \approx 10$–20), and the cross-product $X_{\text{sketch}} Y_{\text{sketch}}^T$ is cached to avoid redundant computation across iterations, (2) neighbor lookups are $O(k)$ per spot where $k \ll N$, and (3) the inner loops are JIT-compiled

using Numba for near-C performance. The overall complexity per iteration is $O(N \cdot K \cdot (K + k))$, which is linear in the number of spots.

Benchmarking

We evaluated FlashDeconv on the Spotless benchmark suite [12], which provides standardized ground-truth data for deconvolution method comparison:

- Silver Standard: 56 synthetic datasets (6 tissues $\times$ 9 abundance patterns, plus replicates) generated by computationally mixing scRNA-seq profiles. Each dataset contains 1,000–5,000 pseudo-spots with known cell type proportions, enabling systematic evaluation across diverse biological contexts (brain cortex, cerebellum [cell and nucleus references], hippocampus, kidney, skin) and abundance scenarios (dominant types, rare types, uniform distribution).
- Gold Standard: Real spatial transcriptomics data with ground-truth proportions derived from co-registered imaging. STARMap (1 dataset, 108 spots, mouse visual cortex) provides subcellular-resolution validation, while seqFISH+ (7 fields of view, <10 spots each, mouse cortex and olfactory bulb) tests performance on extremely small sample sizes.

To ensure strictly comparable evaluation, performance metrics for competing methods (Cell2Location, RCTD, Stereoscope, etc.) were obtained directly from the official Spotless benchmark results [12]. FlashDeconv was evaluated on identical source datasets using the same ground-truth labels and metric computation procedures.

Performance metrics include Root Mean Square Error (RMSE), Pearson correlation coefficient, and Area Under the Precision-Recall Curve (AUPR) for rare cell type detection. For Gold Standard data, we additionally report Jensen-Shannon Divergence (JSD) following the Spotless protocol. All runtime and memory benchmarks were performed on an Apple MacBook Pro with M2 Max chip (32GB unified memory) running macOS, representing consumer-grade hardware without GPU acceleration.

Data Availability

The Spotless benchmark datasets are available at Zenodo (https://zenodo.org/records/10277187) with code at https://github.com/saeyslab/spotless-benchmark. The Mouse Brain Visium dataset was obtained from the cell2location data portal (https://cell2location.cog.sanger.ac.uk/tutorial/); raw data are available at ArrayExpress (accession E-MTAB-11114 for Visium data, E-MTAB-11115 for snRNA-seq reference) [5]. The Visium HD Mouse Small Intestine (FFPE) dataset was obtained from 10x Genomics (https://www.10xgenomics.com/datasets/visium-hd-cytassist-gene-expression-libraries-of-mouse-intestine).
The Xenium Fresh Frozen Mouse Colon dataset, used for ground truth validation of deconvolution accuracy with single-cell resolution data, was obtained from 10x Genomics (https://www.10xgenomics.com/datasets/fresh-frozen-mouse-colon-with-xenium-multimodal-cell-segmentation-1-standard).
The intestinal scRNA-seq reference was obtained from Haber et al. [40]; we used the pre-processed version from Zenodo (https://zenodo.org/records/4447233), which

contains 10,896 cells with cell type annotations. The original data is deposited at GEO (GSE92332). The human ovarian cancer Visium dataset is available at GEO (https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE211956) [39]. The Visium HD colorectal cancer cohort data and associated Chromium Flex scRNA-seq reference are from Oliveira et al. [37] and available at Zenodo (https://doi.org/10.5281/zenodo.15042463) and 10x Genomics (https://www.10xgenomics.com/datasets/visium-hd-cytassist-gene-expression-libraries-of-human-crc).

Code Availability

FlashDeconv is implemented as an open-source Python package available at https://github.com/cafferychen777/flashdeconv. Complete analysis scripts for reproducing all figures and benchmarks presented in this paper are available at https://github.com/cafferychen777/flashdeconv-reproducibility.

Acknowledgements. We acknowledge the publicly available datasets used in our benchmarking, including those from the Spotless benchmark suite, Cell2location mouse brain data, and 10x Genomics Visium HD datasets. We also thank the developers of the open-source tools utilized in this work.

Funding. This work was supported by the National Institutes of Health R01GM144351 (J.C. and X.Z.), National Science Foundation DMS1830392, DMS2113359, DMS1811747 (X.Z.), National Science Foundation DMS2113360, and Mayo Clinic Center for Individualized Medicine (J.C.).

Author Contributions. C.Y., X.Z., and J.C. jointly conceptualized the study. C.Y. developed the methodology, implemented the computational framework, conducted formal analysis, and performed data curation, visualization, and validation. X.Z. and J.C. jointly supervised the research, secured funding, provided resources, served as corresponding authors, and administered the project, contributing equally in these senior roles. C.Y., X.Z., and J.C. all participated in manuscript writing and revision.

Competing Interests. The authors declare no competing interests.

References

[1] Ståhl, P. L. *et al.* Visualization and analysis of gene expression in tissue sections by spatial transcriptomics. *Science* **353**, 78–82 (2016).

[2] Chen, A. *et al.* Spatiotemporal transcriptomic atlas of mouse organogenesis using dna nanoball-patterned arrays. *Cell* **185**, 1777–1792 (2022).

[3] Janesick, A. *et al.* High resolution mapping of the tumor microenvironment using integrated single-cell, spatial and in situ analysis. *Nature Communications* **14**, 8353 (2023).

[4] Gaspard-Boulinc, L. C., Gortana, L., Walter, T., Barillot, E. & Cavalli, F. M. G. Cell-type deconvolution methods for spatial transcriptomics. *Nature Reviews Genetics* **26**, 828–846 (2025).

[5] Kleshchevnikov, V. *et al.* Cell2location maps fine-grained cell types in spatial transcriptomics. *Nature Biotechnology* **40**, 661–671 (2022).

[6] Andersson, A. *et al.* Single-cell and spatial transcriptomics enables probabilistic inference of cell type topography. *Communications Biology* **3**, 565 (2020).

[7] Wang, X., Park, J., Susztak, K., Zhang, N. R. & Li, M. Bulk tissue cell type deconvolution with multi-subject single-cell expression reference. *Nature Communications* **10**, 380 (2019).

[8] Dong, R. & Yuan, G.-C. Spatialdwls: accurate deconvolution of spatial transcriptomic data. *Genome Biology* **22**, 145 (2021).

[9] Cable, D. M. *et al.* Robust decomposition of cell type mixtures in spatial transcriptomics. *Nature Biotechnology* **40**, 517–526 (2022).

[10] Elosua-Bayes, M., Nieto, P., Mereu, E., Gut, I. & Heyn, H. Spotlight: seeded nmf regression to deconvolute spatial transcriptomics spots with single-cell transcriptomes. *Nucleic Acids Research* **49**, e50–e50 (2021).

[11] Biancalani, T. *et al.* Deep learning and alignment of spatially resolved single-cell transcriptomes with tangram. *Nature Methods* **18**, 1352–1362 (2021).

[12] Sang-aram, C., Browaeys, R., Seurinck, R. & Saeys, Y. Spotless, a reproducible pipeline for benchmarking cell type deconvolution in spatial transcriptomics. *eLife* **12**, RP88431 (2024).

[13] Ma, Y. & Zhou, X. Spatially informed cell-type deconvolution for spatial transcriptomics. *Nature Biotechnology* **40**, 1349–1359 (2022).

[14] Luecken, M. D. & Theis, F. J. Current best practices in single-cell rna-seq analysis: a tutorial. *Molecular Systems Biology* **15**, e8746 (2019).

[15] Heumos, L. *et al.* Best practices for single-cell analysis across modalities. *Nature Reviews Genetics* **24**, 550–572 (2023).

[16] Batlle, E. & Clevers, H. Cancer stem cells revisited. *Nature Medicine* **23**, 1124–1134 (2017).

[17] Kreso, A. & Dick, J. E. Evolution of the cancer stem cell model. *Cell Stem Cell* **14**, 275–291 (2014).

[18] Rafii, S., Butler, J. M. & Ding, B.-S. Angiocrine functions of organ-specific endothelial cells. *Nature* **529**, 316–325 (2016).

[19] Amezquita, R. A. *et al.* *Orchestrating Single-Cell Analysis with Bioconductor* (Bioconductor, 2020). URL https://bioconductor.org/books/release/OSCA/. Chapter 8: Feature Selection.

[20] Jindal, A., Gupta, P., Jayadeva & Sengupta, D. Discovery of rare cells from voluminous single cell expression data. *Nature Communications* **9**, 4719 (2018).

[21] Jiang, L., Chen, H., Pinello, L. & Yuan, G.-C. GiniClust: detecting rare cell types from single-cell gene expression data with Gini index. *Genome Biology* **17**, 144 (2016).

[22] Townes, F. W., Hicks, S. C., Aryee, M. J. & Irizarry, R. A. Feature selection and dimension reduction for single-cell RNA-Seq based on a multinomial model. *Genome Biology* **20**, 295 (2019).

[23] Hafemeister, C. & Satija, R. Normalization and variance stabilization of single-cell rna-seq data using regularized negative binomial regression. *Genome Biology* **20**, 296 (2019).

[24] Johnson, W. B. & Lindenstrauss, J. Extensions of lipschitz mappings into a hilbert space. *Contemporary Mathematics* **26**, 189–206 (1984).

[25] Drineas, P., Magdon-Ismail, M., Mahoney, M. W. & Woodruff, D. P. Fast approximation of matrix coherence and statistical leverage. *Journal of Machine Learning Research* **13**, 3475–3506 (2012).

[26] Drineas, P. & Mahoney, M. W. Randnla: randomized numerical linear algebra. *Communications of the ACM* **59**, 80–90 (2016).

[27] Clarkson, K. L. & Woodruff, D. P. Low rank approximation and regression in input sparsity time. *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)* 81–90 (2013).

[28] Chen, L., Liu, H., Kocher, J.-P. A., Li, H. & Chen, J. glmgraph: an R package for variable selection and predictive modeling of structured genomic data. *Bioinformatics* **31**, 3991–3993 (2015).

[29] Liu, Y. *et al.* SDePER: a hybrid machine learning and regression method for cell-type deconvolution of spatial barcoding-based transcriptomic data. *Genome Biology* **25**, 271 (2024).

[30] Zhang, M., Parker, J., An, L., Liu, Y. & Sun, X. Flexible analysis of spatial transcriptomics data (FAST): a deconvolution approach. *BMC Bioinformatics* **26**, 35 (2025).

[31] Rue, H. & Held, L. *Gaussian Markov Random Fields: Theory and Applications* Vol. 104 of *Monographs on Statistics and Applied Probability* (Chapman & Hall/CRC, 2005).

[32] Stuart, T. *et al.* Comprehensive integration of single-cell data. *Cell* **177**, 1888–1902 (2019).

[33] Chen, E. Y. *et al.* Enrichr: interactive and collaborative HTML5 gene list enrichment analysis tool. *BMC Bioinformatics* **14**, 128 (2013).

[34] Xie, Z. *et al.* Gene set knowledge discovery with Enrichr. *Current Protocols* **1**, e90 (2021).

[35] Li, H. *et al.* A comprehensive benchmarking with practical guidelines for cellular deconvolution of spatial transcriptomics. *Nature Communications* **14**, 1548 (2023).

[36] Li, B. *et al.* Benchmarking spatial and single-cell transcriptomics integration methods for transcript distribution prediction and cell type deconvolution. *Nature Methods* **19**, 662–670 (2022).

[37] Oliveira, M. F. d. *et al.* High-definition spatial transcriptomic profiling of immune cell populations in colorectal cancer. *Nature Genetics* **57**, 1512–1523 (2025).

[38] Galon, J. *et al.* Type, density, and location of immune cells within human colorectal tumors predict clinical outcome. *Science* **313**, 1960–1964 (2006).

[39] Denisenko, E. *et al.* Spatial transcriptomics reveals discrete tumour microenvironments and autocrine loops within ovarian cancer subclones. *Nature Communications* **15**, 2860 (2024).

[40] Haber, A. L. *et al.* A single-cell survey of the small intestinal epithelium. *Nature* **551**, 333–339 (2017).

[41] Openshaw, S. The modifiable areal unit problem. *Concepts and Techniques in Modern Geography* **38**, 1–41 (1984).

[42] Zormpas, E., Queen, R., Comber, A. & Cockell, S. J. Mapping the transcriptome: Realizing the full potential of spatial data analysis. *Cell* **186**, 5677–5689 (2023).

[43] Middelhoff, M. *et al.* Dclk1-expressing tuft cells: critical modulators of the intestinal niche? *American Journal of Physiology-Gastrointestinal and Liver Physiology* **313**, G285–G299 (2017).

[44] Gerbe, F. *et al.* Intestinal epithelial tuft cells initiate type 2 mucosal immunity to helminth parasites. *Nature* **529**, 226–230 (2016).

[45] Gerbe, F., Legraverend, C. & Jay, P. The intestinal epithelium tuft cells: specification and function. *Cellular and Molecular Life Sciences* **69**, 2907–2917 (2012).

[46] Huang, L. *et al.* Tuft cells act as regenerative stem cells in the human intestine. *Nature* **634**, 929–935 (2024).

[47] von Moltke, J., Ji, M., Liang, H.-E. & Locksley, R. M. Tuft-cell-derived il-25 regulates an intestinal ilc2–epithelial response circuit. *Nature* **529**, 221–225 (2016).

[48] Drineas, P., Mahoney, M. W. & Muthukrishnan, S. Sampling algorithms for ℓ_2 regression and applications. *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* 1127–1136 (2006).

[49] Mahoney, M. W. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning* **3**, 123–224 (2011).

[50] Hie, B., Cho, H., DeMeo, B., Bryson, B. & Berger, B. Geometric sketching compactly summarizes the single-cell transcriptomic landscape. *Cell Systems* **8**, 483–493 (2019).

[51] Hao, Y. *et al.* Dictionary learning for integrative, multimodal and scalable single-cell analysis. *Nature Biotechnology* **42**, 293–304 (2024).

[52] Bodor, A., Csabai, I., Mahoney, M. W. & Solymosi, N. rcur: an R package for CUR matrix decomposition. *BMC Bioinformatics* **13**, 103 (2012).

[53] Zhong, W., Liu, Y. & Zeng, P. A model-free variable screening method based on leverage score. *Journal of the American Statistical Association* **118**, 135–146 (2023).

[54] Maier, B. *et al.* A conserved dendritic-cell regulatory program limits antitumour immunity. *Nature* **580**, 257–262 (2020).

[55] Charikar, M., Chen, K. & Farach-Colton, M. Finding frequent items in data streams. *Theoretical Computer Science* **312**, 3–15 (2004). Preliminary version in ICALP 2002.

[56] Drineas, P., Mahoney, M. W., Muthukrishnan, S. & Sarlós, T. Faster least squares approximation. *Numerische Mathematik* **117**, 219–249 (2011).

[57] Martinsson, P.-G. & Tropp, J. A. Randomized numerical linear algebra: Foundations and algorithms. *Acta Numerica* **29**, 403–572 (2020).

[58] Kim, J., He, Y. & Park, H. Algorithms for nonnegative matrix and tensor factorizations: A unified view based on block coordinate descent framework. *Journal of Global Optimization* **58**, 285–319 (2014).

[59] Tseng, P. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications* **109**, 475–494 (2001).