

Review Article

Non-Negative Blind Source Separation and the Identifiability of Cell-Type Deconvolution in Spatial Transcriptomics: A Mathematical Review

Amit Walinjar¹

1. Vitacore Research, London, United Kingdom

Cell-type deconvolution in spatial transcriptomics is, in signal-processing terms, a non-negative blind source separation (BSS) problem: an observed mixture $Y = AS + E$ must be resolved into a mixing matrix of cell-type proportions A and a set of source signatures S , from counts alone. The biological literature has answered this with more than twenty competing algorithms and a series of benchmarks that agree on only one thing—no method wins consistently—without a formal account of why. This review supplies that account from the source-separation side. We show that the identifiability theory developed for BSS, independent component analysis, and non-negative matrix factorisation—and matured over two decades of hyperspectral unmixing—transfers directly to the deconvolution problem, because it is literally the same model on the same equation. Three conditions organise the entire field: full column rank of the signature matrix (reference completeness), which governs supervised methods; the condition number $\kappa(S)$, which bounds noise amplification when two sources are nearly collinear; and a separability or minimum-volume condition, which governs the unsupervised case. We further show that the negative-binomial nature of count data makes the squared-error objectives most methods minimise a high-count approximation rather than the maximum-likelihood estimator. We locate the major published algorithms within this framework, demonstrate through a worked simulation that recovered-proportion error tracks the $\kappa(S)$ bound, and show that each condition reproduces a documented benchmark observation— including the much-noted result that a plain regression keeps pace with half the dedicated tools. The practical payoff is that all three conditions are computable on the reference before any algorithm is run, turning method selection from a benchmark lookup into a short calculation.

1. Introduction

Spatial transcriptomics (ST) reads gene expression across thousands of positions in a tissue section without discarding where each reading came from. That spatial context is the whole point, and it is what dissociation-based single-cell methods throw away. *Nature Methods* called the approach its Method of the Year in 2020 ^[1], and the technology has since pushed out of method development and into the clinic, most visibly in oncology. The reason is that the spatial arrangement of the tumour microenvironment carries therapeutic information you cannot get any other way: profiling 274,000 cells across twenty pretreatment lesions showed that the balance of immune-supportive to immune-suppressive cells separates complete from partial responders to immunotherapy ^[2], and spatial immune niches in metastatic renal carcinoma sort checkpoint-inhibitor responders where bulk sequencing is blind to them ^[3]. In both studies, the clinical call comes down to knowing which cell types are where.

The trouble is that this is never measured directly. The platforms in routine use—10x Visium and its relatives—read expression from barcoded spots, and a spot of that size holds several cells, often of different types ^[4]. Every spot is therefore a blend, and pulling the per-spot cell-type composition back out of that blend is a computational problem, not an experimental one. This is deconvolution, and for spot-based ST it is the bottleneck.

The community has answered with an abundance of algorithms rather than a theory. Reference-anchored methods dominate practice: RCTD fits a Poisson model against a single-cell reference ^[5], Cell2location wraps a Bayesian prior with spatial smoothing around the same idea ^[6], and SpatialDWLS leans on dampened weighted least squares ^[7]. Then came the reference-free and deep-learning entrants—graph neural networks ^{[8][9]}, single-cell-resolution regression ^[10], non-negative least squares ^[11], semi-supervised transfer ^[12], and independent component analysis run straight on the spot matrix ^[13]. Recent reviews can catalogue this growing list but cannot rank it ^[14].

Benchmarking was supposed to settle things and instead documented the mess. An evaluation of 18 methods over 50 datasets put CARD, Cell2location and Tangram on top on average, found no method that led everywhere, and fell back on a decision tree for picking one ^[15]. The Spotless benchmark—11 methods, 63 silver and 3 gold standards—landed two results the field has not really absorbed: Cell2location and RCTD led overall, yet a plain regression beat almost half of the purpose-built spatial

methods, and every method fell apart once cell types were rare or highly abundant [16]. A separate evaluation on cardiovascular and kidney tissue concluded only that the best choice “depends on the use case” [17]. The shared signature across all three is performance that is real but unpredictable—selection by folklore.

What is missing is the question a signal-processing engineer asks first. None of these algorithms states the conditions under which its model has a *unique* solution, and none predicts failure before the method is run. They each fit a model and report a score; not one asks whether the separation it attempts is identifiable to begin with. That question was answered long ago for the very same mathematical object—the non-negative linear mixture. Blind source separation, independent component analysis and non-negative matrix factorisation all come with identifiability theory [18][19][20], and remote sensing in particular has spent two decades unmixing hyperspectral images, where each pixel is a mixture of material spectra, into endmembers and their abundances, with provable conditions on rank, conditioning and source geometry [21][22][23]. Swap “endmember” for “cell-type signature” and “abundance” for “cell-type proportion” and it is the same problem on the same equation.

That correspondence is the subject of this review. We do not offer a new algorithm, and we do not claim priority for applying a separation method to ST data—ICA-based tools already exist [13]. The aim is consolidation: to show that deconvolution is a non-negative blind source separation problem, and that once it is written as one, the identifiability conditions the biological field has managed without follow directly from established signal-processing theory. With them, the scattered failures stop looking random, the algorithm zoo sorts itself by which assumption each method depends on, and method selection becomes a short calculation. Section 2 lays out the model, Section 3 maps the algorithms onto it, Section 4 turns the conditions into selection criteria, Section 5 follows the consequences into the clinic, and Section 6 takes stock.

2. The Formal Signal Model


2.1. Observation model

Write $Y \in \mathbb{R}^{n \times g}$ for the observed expression over n spots and g genes, $A \in \mathbb{R}_{\geq 0}^{n \times k}$ for the cell-type proportions over k types, and $S \in \mathbb{R}_{\geq 0}^{k \times g}$ for the cell-type signatures. Deconvolution is the linear instantaneous mixing model

$$Y = AS + E, \tag{1}$$

with $A \geq 0$, $S \geq 0$, the simplex constraint $A\mathbf{1}_k = \mathbf{1}_n$ so that each spot's proportions sum to one, and E collecting sequencing noise, dropout and ambient RNA. Read row by row, Equation (1) says each spot is a convex blend of the k signatures, weighted by that spot's proportions. This is the canonical blind source separation setup: A is the mixing matrix, the rows of S are the sources, and deconvolution recovers both from Y alone. The decomposition is shown schematically in Figure 1.

Figure 1. Spatial transcriptomics deconvolution as a linear mixing model $Y = AS + E$



Vitacore Research | awalkr@vitacore-research.org

Figure 1. Spatial transcriptomics deconvolution as a linear mixing model. The observed spots-by-genes matrix Y factorises into cell-type proportions A and cell-type signatures S under non-negativity and simplex constraints, with an additive count-noise term E . This is the canonical non-negative blind source separation model.

2.2. Noise model

Most methods treat the noise as Gaussian by implication, because they minimise a squared error or a weighted least squares. The count nature of the data contradicts this. Unique-molecular-identifier counts are well described by a negative-binomial (NB) distribution [24][25], and careful analysis of technical replicates places them close to Poisson with moderate overdispersion [26]. So the honest model is

$$Y_{ij} \sim NB\left((AS)_{ij}, \theta\right), \quad (2)$$

with mean $(AS)_{ij}$ and dispersion θ , rather than $E_{ij} \sim \mathcal{N}(0, \sigma^2)$. Equation (2) matters because least squares is the maximum-likelihood estimator only when the noise is Gaussian; under counts it is

misspecified, and the misspecification bites hardest at low depth, where the mean is small and the two likelihoods part company.

2.3. Optimality criterion

The estimator that respects the data maximises the NB log-likelihood:

$$(\hat{A}, \hat{S}) = \arg \max_{A, S \geq 0} \sum_{i,j} \left[Y_{ij} \log \frac{(AS)_{ij}}{(AS)_{ij} + \theta} + \theta \log \frac{\theta}{(AS)_{ij} + \theta} \right] \quad (3)$$

Equation (3) is the correct objective. Gaussian least squares, $\min \|Y - AS\|_F^2$, drops out of it only in the large-mean limit at fixed variance—so the squared-error objective is a high-count approximation, fine where counts are large and increasingly wrong where they are not.

3. Identifiability and the Algorithm Landscape

The payoff from the BSS framing is that identifiability—whether a unique solution even exists—has already been worked out for each case: sources known, partly known, or unknown. Those three cases are exactly supervised, semi-supervised and unsupervised deconvolution.

3.1. Supervised regime: reference completeness

When a single-cell reference supplies $S = S_{ref}$, deconvolution collapses, per spot, into a constrained regression. For spot i , with $y_i \in \mathbb{R}^g$ and $a_i \in \mathbb{R}^k$,

$$\hat{a}_i = \arg \min_{a_i \geq 0, a_i^\top \mathbf{1} = 1} \|y_i - S^\top a_i\|^2, \quad (4)$$

and the proportions are pinned down uniquely precisely when

$$\boxed{\text{identifiable} \Leftrightarrow \text{rank}(S) = k} \quad (5)$$

that is, when the reference signatures are linearly independent. Strip away the packaging and Equation (4) is what RCTD, Cell2location and SpatialDWLS ^{[5][6][7]} all are—supervised solutions to one regression, differing in likelihood and regularisation, not in the problem they face. Equation (5) also explains the Spotless result that a plain regression matched half the dedicated methods ^[16]. In this regime, with a well-conditioned reference, deconvolution simply *is* constrained regression; the extra machinery buys variance, not identifiability. The failure mode is just as formal: if a cell type sits in the

tissue but is missing from the reference, then $\text{rank}(S) < k_{\text{true}}$, the span condition breaks, and the orphaned signal gets folded into whichever types are present, inflating their proportions.


3.2. Conditioning: the similar-source problem

A complete reference is still not enough on its own. How violently noise gets amplified into the proportion estimates is bounded by the condition number of the signature matrix:

$$\boxed{\frac{\|\hat{a}_i - a_i\|}{\|a_i\|} \lesssim \kappa(S) \frac{\|e_i\|}{\|y_i\|}} \quad (6)$$

When two cell types express almost the same programme—CD4 and CD8 T-cell subtypes are the textbook case—the rows of S run nearly parallel, $\kappa(S)$ blows up, and a little measurement noise turns into a large error in the recovered proportions. This is the identical pathology that remote sensing hits when two materials have nearly the same spectrum, and remote sensing handles it by geometric preconditioning that provably drives the effective condition number down ^[22]. Equation (6) predicts, with no new experiment, the benchmark finding that every method struggles on transcriptionally similar cell types ^[16]; that degradation is a property of the data’s conditioning, not a flaw anyone can engineer out. Figure 2 shows the bound in a controlled simulation: as $\kappa(S)$ grows, the relative error in recovered proportions rises in step with it.

Figure 2. Proportion error grows with signature conditioning
(illustrative simulation; fixed seed)



Vitacore Research | awalkr@vitacore-research.org

Figure 2. Proportion error grows with signature conditioning. Illustrative simulation (fixed seed): synthetic signature matrices with increasing condition number $\kappa(S)$ are used to generate noisy spot mixtures, which are then deconvolved by constrained least squares. The mean relative error in recovered proportions tracks the theoretical bound of Equation (6), confirming that ill-conditioning—transcriptionally similar cell types—is the source of instability.

3.3. Unsupervised regime: separability and minimum volume

With neither A nor S in hand, deconvolution is blind source separation in the full sense, and non-negativity by itself does not buy uniqueness—there can be many factorisations $Y = A_1 S_1 = A_2 S_2$ [20] [27]. Something extra is needed. The cleanest such condition is separability, the existence of an anchor or marker gene for each cell type:

$$\forall t \in \{1, \dots, k\} \exists j : S_{tj} > 0, S_{t'j} = 0 \forall t' \neq t \quad (7)$$

and where exact anchors are too much to ask, a minimum-volume condition on the convex hull of the data restores uniqueness instead—and does so even when the number of sources has been overestimated [23]. Equation (7) is the formal reason unsupervised and NMF-style deconvolution

wobbles: without a marker-gene or minimum-volume guarantee, the method hands back *a* factorisation rather than *the* factorisation, and which one depends on where it started. ICA tools such as CellsFromSpace ^[13] swap statistical independence in for these conditions—a different identifiability assumption, not an escape from needing one.

3.4. Summary

Three conditions, then, organise the whole field, and every published algorithm lands somewhere on the grid of Figure 3, defined by which conditions it secures and which it leaves to luck. Supervised methods are identifiable when the reference has full column rank and stable when $\kappa(S)$ stays small; unsupervised methods are identifiable when the sources are separable or minimum-volume; and all of them are correctly specified only under an NB likelihood.

Figure 3. Existing spatial-transcriptomics deconvolution algorithms located within the blind source separation framework

Algorithm	Regime	BSS equivalent	Identifiability assumption	Known failure condition
RCTD	Supervised	Supervised NMF (Poisson likelihood)	Reference S complete, full column rank	Missing cell type $\rightarrow$ residual absorbed into present types
Cell2location	Supervised	Bayesian supervised deconvolution	S known; prior enforces non-neg. + smoothness	Smoothness prior biases estimates at sharp tissue boundaries
SpatialDWLS	Supervised	Dampened weighted least squares	S known; heuristic gene weighting	Weights not ML under true (NB) noise model
SpaDecon	Semi-supervised	Partial- S transfer learning	Partial signatures + spatial/histology	Constraints implicit, not stated as identifiability conditions
ICA / NMF (CellsFromSpace)	Unsupervised	Classic BSS (ICA or NMF)	Source independence OR non-neg. + sparsity	Non-unique without separability or minimum-volume

Vitacore Research | awalkr@vitacore-research.org


Figure 3. Existing algorithms located within the BSS framework. Each spatial-transcriptomics deconvolution method is a solution within one supervision regime, distinguished by its likelihood and regularisation rather than by the problem it addresses, and each carries an identifiability assumption whose violation produces a documented failure mode.

4. Algorithm Selection as a Calculation

The conditions of Section 3 are not just hindsight. Each is computable up front, on the reference, which is what turns selection from a benchmark lookup into arithmetic (Figure 4). First, the reference span: before fitting anything, one can ask whether the spots lie in the column space of a candidate signature matrix; a

residual analysis after a trial fit shows tell-tale structure when types are missing. A reference that fails this is not rescued by a cleverer algorithm—it is rescued by a better reference. Second, the conditioning: $\kappa(S)$ falls straight out of the reference signatures, no spatial data required, and a large value warns that the correlated types will come back unstable no matter which method runs. Third, the regime: whether a complete single-cell reference exists decides whether the problem is well-posed regression or an unsupervised problem in which separability must be argued. This is the real content under the empirical decision trees of earlier benchmarks [\[14\]\[15\]](#).

Figure 4. Identifiability-based algorithm selection for spatial-transcriptomics deconvolution



Vitacore Research | awalkr@vitacore-research.org

Figure 4. Identifiability-based algorithm selection. The supervision regime, the existence of marker genes, and the condition number $\kappa(S)$ are all computable from the reference before any deconvolution is run, yielding concrete guidance rather than a benchmark lookup.

The guidance that comes out is concrete rather than hand-wavy. With a complete, well-conditioned reference, the problem is well-posed and the choice of method barely matters—a regularised regression or an established Poisson method will land in much the same place, which is what the benchmark equivalences already show [\[16\]](#). With a complete but ill-conditioned reference, the priority is dragging $\kappa(S)$ down by relabelling or preconditioning, not auditioning methods. With no reference, the only

question worth asking is whether marker genes exist; if they do, separability-aware methods are justified, and if they do not, no unsupervised result should be reported as though it were unique.

5. Clinical and Regulatory Implications

The framework speaks straight to the settings that make ST worth the trouble. Tumour-microenvironment composition predicts immunotherapy response through quantities like the immune-supportive-to-suppressive ratio [2] and spatial-niche signatures [3]—and those are ratios of deconvolved proportions, so they inherit whatever instability the deconvolution carries in with it. A high $\kappa(S)$ between effector and regulatory T-cell signatures means the very ratio the clinical decision rests on is the one most sensitive to noise. The framework can tell a clinician not just what was estimated but whether the estimate was identifiable in the first place—which is the distinction that matters once an output starts steering treatment.

Longitudinal monitoring makes the point sharper. Comparing composition across timepoints assumes the deconvolution is stable under the batch and parameter drift that separates the measurements; instability once chalked up to biological change may instead be a violated condition. For software-as-a-medical-device review, the framework offers what regulators already expect of any measurement system: an analytical specification. A pipeline going forward for clinical use can state its identifiability conditions, report $\kappa(S)$ for its reference, and show that its operating regime is well-posed—the computational counterpart of characterising a signal chain before you certify the instrument hanging off the end of it.

6. Discussion

The contribution of this review is a reframing, and reframings earn their keep only if something falls out. By writing deconvolution as $Y = AS + E$ and asking the identifiability question that signal processing asks of any separation problem, the field's central puzzle—no method wins, failures look idiosyncratic—resolves into three conditions on rank, conditioning and source structure, each of which reproduces a benchmark result it was never fitted to.

The sharpest of those is the Spotless finding that a plain regression kept pace with half the dedicated methods [16]. Read as a benchmark, that is an embarrassment nobody quite explained. Read through identifiability, it is a prediction: in the supervised regime with a well-conditioned reference, deconvolution reduces exactly to constrained regression, so of course a regression competes. The collapse

of every method on rare, abundant and similar cell types ^{[15][16]} is likewise forecast by the conditioning bound rather than stumbled upon.

Set against prior work, this is unification, not a turf claim. RCTD, Cell2location and SpatialDWLS ^{[5][6]} ^[7] are not pushed aside—they are located, each revealed as a supervised solution distinguished by its likelihood and regularisation. The benchmarking decision trees ^{[14][15]} are not contradicted but derived, their heuristic branches shown to be the regime and conditioning splits written out exactly. And the hyperspectral-unmixing literature ^{[21][22][23]} is not borrowed by loose analogy; it is the same mathematics on the same model, which is why its identifiability results transfer at all.

There is a disciplinary explanation for why the field went straight to machine learning rather than stopping to ask the identifiability question first. Coming from telecommunications and sensor engineering, the first thing you are trained to ask when someone hands you a separation problem is not “which algorithm should I use?” but “is this problem even separable under these conditions?” Before building a receiver in telco, you characterise the channel. Before designing a sensor interface, you specify the noise floor. The engineering habit is to understand the conditions for a unique solution before committing resources to finding it. Genomics, moving at a different pace and from a different tradition, largely skipped that step. The pressure to publish working software and demonstrate empirical performance on held-out data is intense, and it crowds out the slower work of asking whether the problem is well-posed to begin with. This does not reflect badly on the field—the algorithms that emerged are genuinely useful—but it does explain why a framework that signal processing would have reached for first took this long to arrive. The separability question was always there waiting. It just needed carrying across the disciplinary border.

Several limitations bound the claims, and they belong to the model rather than being oversights in it. Linear instantaneous mixing ignores spillover between neighbouring spots and any non-linearity in capture; real tissue has both, and the framework degrades as they grow. The mixing is linear in expectation, with count noise carried by the NB likelihood—a generalised-linear framing, not literal additive noise. The number of cell types is taken as known in the supervised and semi-supervised regimes, and the unsupervised case inherits the old difficulty of choosing rank in NMF. Conditioning is a joint property of the reference and the biology, so the framework can diagnose an ill-conditioned problem but cannot dissolve one when two cell types really are near-identical in expression. And the scope is spot-based sequencing ST; imaging-based platforms creeping toward single-cell resolution ^[14] largely sidestep the mixing problem and fall outside this treatment. These same limitations

mark the natural directions for the identifiability programme: computing $\kappa(S)$ and span diagnostics across benchmark datasets to test whether they predict per-dataset method performance ahead of time, and building a formal identifiability test to run as a pre-analysis quality check.

7. Conclusions

The spatial-transcriptomics deconvolution literature has spent five years running algorithms against each other. It would have been faster to ask first whether the problem was well-posed. It is a non-negative blind source separation problem, and once it is written as one, three computable conditions fall out naturally: reference completeness, signature conditioning, and source separability. These organise the existing methods, explain the failures that benchmarking documented without explaining, and replace the current guesswork with criteria evaluable on the reference data before any deconvolution is run. The correct noise model for counts adds one more clause that the field has largely ignored: squared-error objectives are a high-count approximation, not the right estimator, and the gap matters most exactly where spatial transcriptomics is already weakest.

Reviewing the field through the identifiability lens points to concrete next steps: validating the identifiability conditions as quantitative predictors of per-dataset method performance across the major public benchmarks; building a pre-analysis identifiability check that research groups can run routinely before committing to a deconvolution result; extending the framework to imaging-based platforms approaching single-cell resolution, where the mixing model itself changes; and laying the formal groundwork for regulatory specification of ST pipelines under software-as-a-medical-device review. The identifiability conditions are stated in computable form precisely so that experimental and computational groups can test them against their own data. That is where the framework earns its keep.

Statements and Declarations

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Potential Competing Interests

The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data Availability

This work is theoretical and uses no primary data. Everything referenced is drawn from the cited published literature, available through the respective publishers. The simulation in Figure 2 is fully synthetic and reproducible from its stated random seed.

Prior Dissemination

An earlier version of this framework was deposited as a preprint on Zenodo (DOI: 10.5281/zenodo.20760315). The present version is a revised and re-framed account of the same work. This manuscript has not been published in, and is not under consideration at, any peer-reviewed journal.

Use of Generative AI

Generative AI tools were used to assist with language editing, formatting, and LaTeX preparation. All scientific content, analysis, and conclusions are the author's own, and the author has reviewed and takes full responsibility for the manuscript.

References

1. [△]Nature Methods (2021). "Method of the Year 2020: Spatially Resolved Transcriptomics." *Nat Methods*. **18** (1):1. doi:[10.1038/s41592-020-01042-x](https://doi.org/10.1038/s41592-020-01042-x).
2. [△][△]Abdulrahman Z, Slieker R, McGuire D, Welters M, van Poelgeest M, van der Burg S (2025). "Single-Cell Spatial Transcriptomics Unravels Cell States and Ecosystems Associated with Clinical Response to Immunotherapy." *J Immunother Cancer*. **13**(3):e011308. doi:[10.1136/jitc-2024-011308](https://doi.org/10.1136/jitc-2024-011308).
3. [△][△]Kalogirou C, Krebs M, Kunz A, Hahn O, Kübler H, Schwinger M, Hartmann A, Schmieder A, Eckstein M (2025). "Spatial Transcriptomic Profiling of Metastatic Renal Cell Carcinoma Identifies Chemokine-Driven Macrophage and CD8⁺ T-Cell Interactions Predictive of Immunotherapy Response." *J Immunother Cancer*. **13**:e012991. doi:[10.1136/jitc-2025-012991](https://doi.org/10.1136/jitc-2025-012991).

4. ^aStåhl P, Salmén F, Vickovic S, Lundmark A, Navarro J, Magnusson J, Giacomello S, Asp M, Westholm J, Huss M, Mollbrink A, Linnarsson S, Codeluppi S, Borg Å, Ponten F, Costea P, Sahlen P, Mulder J, Bergmann O, Lundberg J, Frisen J (2016). "Visualization and Analysis of Gene Expression in Tissue Sections by Spatial Transcriptomics." *Science*. **353**(6294):78–82. doi:[10.1126/science.aaf2403](https://doi.org/10.1126/science.aaf2403).
5. ^a, ^b, ^cCable D, Murray E, Zou L, Goeva A, Macosko E, Chen F, Irizarry R (2022). "Robust Decomposition of Cell Type Mixtures in Spatial Transcriptomics." *Nat Biotechnol*. **40**(4):517–526. doi:[10.1038/s41587-021-00830-w](https://doi.org/10.1038/s41587-021-00830-w).
6. ^a, ^b, ^cKleshchevnikov V, Shmatko A, Dann E, Aivazidis A, King H, Li T, Elmentaite R, Lomakin A, Kedlian V, Gayoso A, Jain M, Park J, Ramona L, Tuck E, Arutyunyan A, Vento-Tormo R, Gerstung M, James L, Stegle O, Bayraktar O (2022). "Cell2location Maps Fine-Grained Cell Types in Spatial Transcriptomics." *Nat Biotechnol*. **40**(5):661–671. doi:[10.1038/s41587-021-01139-4](https://doi.org/10.1038/s41587-021-01139-4).
7. ^a, ^b, ^cDong R, Yuan G (2021). "SpatialDWLS: Accurate Deconvolution of Spatial Transcriptomic Data." *Genome Biol*. **22**(1):145. doi:[10.1186/s13059-021-02362-7](https://doi.org/10.1186/s13059-021-02362-7).
8. ^aLong Y, Ang K, Li M, Chong K, Sethi R, Zhong C, Xu H, Ong Z, Sachaphibulkij K, Chen A, Zeng L, Fu H, Wu M, Lim L, Liu L, Chen J (2023). "Spatially Informed Clustering, Integration, and Deconvolution of Spatial Transcriptomics with GraphST." *Nat Commun*. **14**:1155. doi:[10.1038/s41467-023-36796-3](https://doi.org/10.1038/s41467-023-36796-3).
9. ^aLi Y, Luo Y (2024). "STdGCN: Spatial Transcriptomic Cell-Type Deconvolution Using Graph Convolutional Networks." *Genome Biol*. **25**:206. doi:[10.1186/s13059-024-03353-0](https://doi.org/10.1186/s13059-024-03353-0).
10. ^aZhou Z, Zhong Y, Zhang Z, Ren X (2023). "Spatial Transcriptomics Deconvolution at Single-Cell Resolution Using Redeconve." *Nat Commun*. **14**:7475. doi:[10.1038/s41467-023-43600-9](https://doi.org/10.1038/s41467-023-43600-9).
11. ^aChen Y, Ruan F, Wang J (2024). "NLSDeconv: An Efficient Cell-Type Deconvolution Method for Spatial Transcriptomics Data." *Bioinformatics*. **41**(1):btac747. doi:[10.1093/bioinformatics/btac747](https://doi.org/10.1093/bioinformatics/btac747).
12. ^aColeman K, Hu J, Schroeder A, Lee E, Li M (2023). "SpaDecon: Cell-Type Deconvolution in Spatial Transcriptomics with Semi-Supervised Learning." *Commun Biol*. **6**(1):378. doi:[10.1038/s42003-023-04761-x](https://doi.org/10.1038/s42003-023-04761-x).
13. ^a, ^b, ^cMarie C, others (2023). "CellsFromSpace: A Fast, Accurate and Reference-Free Tool to Deconvolve and Annotate Spatially Distributed Omics Data." 2023 Preprint. doi:[10.1101/2023.08.30.555558](https://doi.org/10.1101/2023.08.30.555558).
14. ^a, ^b, ^c, ^dGaspard-Boulin L, Cavalli F, Theis F, Morrissy A (2025). "Cell-Type Deconvolution Methods for Spatial Transcriptomics." *Nat Rev Genet*. doi:[10.1038/s41576-025-00845-y](https://doi.org/10.1038/s41576-025-00845-y).
15. ^a, ^b, ^c, ^dLi H, Zhou J, Li Z, Chen S, Liao X, Zhang B, Zhang R, Wang Y, Sun S, Gao X (2023). "A Comprehensive Benchmarking with Practical Guidelines for Cellular Deconvolution of Spatial Transcriptomics." *Nat Commun*. **14**(1):1548. doi:[10.1038/s41467-023-37168-7](https://doi.org/10.1038/s41467-023-37168-7).

16. ^{a, b, c, d, e, f}Sang-aram C, Browaeys R, Seurinck R, Saeys Y (2024). "Spotless, a Reproducible Pipeline for Benchmarking Cell Type Deconvolution in Spatial Transcriptomics." *eLife*. **12**:RP88431. doi:[10.7554/eLife.88431](https://doi.org/10.7554/eLife.88431).
17. ^ΔSlabowska A, Pyke C, Owczarek O, others (2024). "A Systematic Evaluation of State-of-the-Art Deconvolution Methods in Spatial Transcriptomics: Insights from Cardiovascular Disease and Chronic Kidney Disease." *Front Bioinform*. **4**:1352594. doi:[10.3389/fbinf.2024.1352594](https://doi.org/10.3389/fbinf.2024.1352594).
18. ^ΔComon P (1994). "Independent Component Analysis, a New Concept?." *Signal Process*. **36**(3):287–314. doi:[10.1016/0165-1684\(94\)90029-9](https://doi.org/10.1016/0165-1684(94)90029-9).
19. ^ΔLee D, Seung H (1999). "Learning the Parts of Objects by Non-Negative Matrix Factorization." *Nature*. **401**(6755):788–791. doi:[10.1038/44565](https://doi.org/10.1038/44565).
20. ^{a, b}Donoho D, Stodden V (2003). "When Does Non-Negative Matrix Factorization Give a Correct Decomposition into Parts?."
21. ^{a, b}Feng X, Li H, Wang R, Du Q, Jia X, Plaza A (2022). "Hyperspectral Unmixing Based on Nonnegative Matrix Factorization: A Comprehensive Review." *IEEE J Sel Topics Appl Earth Obs Remote Sens*. **15**:4414–4436.
22. ^{a, b, c}Lin C, Bioucas-Dias J (2020). "Nonnegative Blind Source Separation for Ill-Conditioned Mixtures via John Ellipsoid." *IEEE Trans Neural Netw Learn Syst*. **32**(5):2209–2223. doi:[10.1109/TNNLS.2020.3002618](https://doi.org/10.1109/TNNLS.2020.3002618).
23. ^{a, b, c}Leplat V, Gillis N, Ang A (2020). "Blind Audio Source Separation with Minimum-Volume Beta-Divergence NMF." *IEEE Trans Signal Process*. **68**:3400–3410.
24. ^ΔChen W, Li Y, Easton J, Finkelstein D, Wu G, Chen X (2018). "UMI-Count Modeling and Differential Expression Analysis for Single-Cell RNA Sequencing." *Genome Biol*. **19**:70. doi:[10.1186/s13059-018-1438-9](https://doi.org/10.1186/s13059-018-1438-9).
25. ^ΔHafemeister C, Satija R (2019). "Normalization and Variance Stabilization of Single-Cell RNA-Seq Data Using Regularized Negative Binomial Regression." *Genome Biol*. **20**:296. doi:[10.1186/s13059-019-1874-1](https://doi.org/10.1186/s13059-019-1874-1).
26. ^ΔLause J, Berens P, Kobak D (2021). "Analytic Pearson Residuals for Normalization of Single-Cell RNA-Seq UMI Data." *Genome Biol*. **22**:258. doi:[10.1186/s13059-021-02451-7](https://doi.org/10.1186/s13059-021-02451-7).
27. ^ΔDeville Y (2019). "From Separability/Identifiability Properties of Bilinear and Linear-Quadratic Mixture Matrix Factorization to Factorization Algorithms." *Digit Signal Process*. **87**:21–33.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.