

Donor-allocation effects vary with reference-cell budget in MuSiC PBMC simulations

Yunhao Jiang

biojiang777@gmail.com

University of Edinburgh

Research Article

Keywords: Cell-type deconvolution; MuSiC; reference donor allocation; reference-cell budget; PBMC; pseudobulk

Posted Date: September 28th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11156252/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Donor-allocation effects vary with reference-cell budget in MuSiC PBMC simulations

Yunhao Jiang

School of Biological Sciences, University of Edinburgh, Edinburgh, UK

Correspondence: biojiang777@gmail.com

Abstract

MuSiC averages donor profiles equally even when they are estimated from different numbers of reference cells. We tested whether the effect of unequal donor allocation depends on the total reference-cell budget. Using raw-count synthetic mixtures of six peripheral blood mononuclear cell (PBMC) types, we compared allocations while holding the three reference donors and total cells per type fixed. In 14 donors annotated as normal, the mean absolute error difference for weighted MuSiC (10:1:1 minus balanced allocation) decreased from +0.621 to -0.066 percentage points as the budget increased from 60 to 300 cells per type. The prespecified change averaged -0.687 points and was negative for 13/14 donors. The trend persisted with alternative donor grouping and additional reference draws, yet 84/168 configurations still had higher error under unequal allocation at the largest budget. In 12 measured PBMC samples, poor agreement with flow cytometry persisted through input diagnostics. Thus, donor-allocation effects were budget-dependent in these simulations: the average penalty decreased, but effects in individual configurations remained. Transfer of this trend to measured bulk remains unestablished.

Keywords

Cell-type deconvolution; MuSiC; reference donor allocation; reference-cell budget; PBMC; pseudobulk

Introduction

Single-cell references support estimation of cell-type proportions from bulk expression data [1,2]. Their total size, however, does not fully describe how they were built. For a cell type represented by 60 reference cells from three donors, the donors could contribute 20:20:20 cells or 50:5:5 cells. Both references contain the same donors and the same number of cells, but provide different amounts of information about each donor. This distinction separates donor allocation from total reference size.

Previous studies have examined preprocessing and marker selection [2], reference weighting [3,4], reference-mixture matching and donor number [5], and balancing cell types or subsampling donors [6]. scValue evaluates cell selection for downstream analyses, including deconvolution [7]. These studies address related choices in reference construction. Here, we ask how redistributing a fixed number of cells among the same donors affects accuracy, and whether that effect changes with the total cell budget.

MuSiC averages donor profiles equally while using between-donor expression variation and estimated cell sizes in fitting [1]. Equal contributions to the average need not mean equally precise profile estimates. We first tested allocation at a small budget; the observed error penalty motivated an external comparison across larger budgets. We then examined dependence on donor grouping and reference sampling, and used profile and weight diagnostics to assess sampling precision as a possible explanation. MuSiC's internal ordinary nonnegative least-squares (NNLS) output provided a comparator on the same reference basis. Finally, measured PBMC RNA-seq with independent flow-cytometry proportions allowed us to assess transfer beyond synthetic mixtures.

Methods

Data and study design

The original analysis used GSE96583 batch-2 control cells from systemic lupus erythematosus (SLE) donors cultured for six hours without interferon stimulation [8]. A fixed eligibility rule selected six donors with sufficient cells in all six retained types: B cells, CD14⁺ monocytes, CD4 T cells, CD8 T cells, FCGR3A⁺ monocytes and NK cells. The selected data comprised 11,066 cells and 35,635 unique gene identifiers. Section S1 describes eligibility, exclusions, a discrepancy in deposited singlet counts and the source correction [9].

The external analysis used a fixed CELLxGENE release associated with Perez et al. [10] and GSE174188/GSE137029. Before generating predictions, we selected all 14 donors with disease annotated as normal, disease_state annotated as na, and at least 250 cells in each retained type. We used verified raw.X counts, the authors' broad cell-type labels and all 30,172 Ensembl identifiers. Targets and references together used 93,892 unique cells initially and 106,359 after additional reference sampling. Donors whose labels matched the original dataset were excluded, although unrecognized aliases remain possible. The source correction concerned cDC1/cDC2 subtypes [11], which were excluded here; the retained annotations remain subject to error.

The original MuSiC endpoint protocol was fixed before its predictions, following earlier experiments with normalized cells (Section S7). Intermediate allocations were added after examining the endpoints. The external protocol and primary contrast were fixed after the original analysis but before external predictions. Grouping, diagnostic and reference-sampling follow-ups were each specified before their own predictions. None was publicly preregistered.

Paired reference and target construction

For each held-out donor in the original cohort, we constructed references from all ten three-donor subsets of the remaining five donors. In the external cohort, a deterministic ring formed four reference triples from 12 of the other 13 donors, leaving one donor unused in each fold. Deposited labels were used to exclude each target donor and all of their cells from the corresponding references. Donors and cells were reused across folds.

The budget B was the total number of cells per type across three reference donors; each six-type reference therefore contained $6B$ cells. At $B=60$, allocations of 20:20:20, 30:15:15, 40:10:10 and 50:5:5 corresponded to ratios of 1:1:1, 2:1:1, 4:1:1 and 10:1:1. Each unequal allocation was tested with every donor in the dominant position, contributing the most cells in all six types. The initial analyses used three paired draws, with cells sampled without replacement within each donor–type group. The original cohort was tested at $B=60$. External budgets of 60, 120 and 300 used the same allocations multiplied by one, two and five.

In the external cohort, each draw generated an ordered cell list for every donor and cell type. We reused these lists across reference triples, taking longer prefixes for larger budgets within each allocation. Whole references were not necessarily nested across allocations: increasing one donor's quota reduced the others' quotas. In the original cohort, each fold and reference triple had its own draws (Sections S3, S9 and S11).

Each donor supplied 60 fixed synthetic targets of 300 cells: one with equal cell-type proportions and 59 with compositions drawn from a symmetric Dirichlet-multinomial distribution. Both sources used the same realized cell-count vectors, with cells sampled separately from each source. Sampling within a type was with replacement in the original cohort and without replacement within each external target; cells could recur across targets. Raw counts were summed, and the sampled cell-count proportions served as ground truth. All conditions reused the same 360 original or 840 external targets.

After examining the external results, we specified an alternative deterministic ring to test dependence on reference-donor grouping. It retained all external donors, targets, ordered cell lists and method parameters,

and used the existing draws to compare balanced and 10:1:1 allocations at B=60 and 300 with every dominant-donor choice. Ten of its 14 unique reference triples differed from the original grouping (Section S14).

Estimation, outcomes and verification

We used the flat music_prop workflow in MuSiC 1.0.0, with all raw-count genes, markers=NULL and default fitting parameters; the exact code version and parameters are specified in Section S4. Both outputs use the same factorized reference basis: the equal-donor mean expression profile multiplied by the mean estimated cell size. They also share the genes used for each target. Weighted MuSiC additionally uses gene weights derived from residuals and between-donor variation. Ordinary NNLS is thus an internal comparator on a matched basis.

We calculated mean absolute error (MAE) and root mean squared error (RMSE) across the six cell-type proportion errors for each prediction, in percentage points. Within each donor, errors were averaged equally over targets, draws, reference triples and dominant-donor choices where applicable; donors then received equal weight. Each unequal allocation was paired with its shared balanced baseline. We define $I(B)=MAE(10:1:1)-MAE(balanced)$, so positive values indicate higher error under unequal allocation. The primary external contrast was weighted $I(300)-I(60)$. Secondary summaries covered absolute errors, intermediate allocations, output differences and reversals in direction. We also summarized 168 arrangement-level contrasts, each defined by a held-out donor, reference triple and dominant donor. Absolute values were taken after averaging over targets and draws (Section S16). These contrasts measure changes in accuracy rather than distances between predicted proportion vectors.

For each fixed target, cell type and donor arrangement, we calculated the sample standard deviation (SD) of predicted proportions across the three reference draws and then averaged these SDs. This conditional SD measures variation in predictions, not uncertainty in the paired MAE contrast. Shared donors, targets and reference cells make predictions and donor summaries dependent. We therefore report descriptive results without population P values or confidence intervals. Sections S6, S9 and S13 describe computational checks.

The diagnostic follow-up compared balanced and 10:1:1 allocations across all external donors, budgets and three reference draws. Eight existing targets per donor were selected by identifier without using outcomes. We measured variation in estimated expression profiles across draws, between-donor variance Σ and distances between fitted gene weights on shared genes to examine sampling precision as a possible explanation for the budget trend (Section S15). The primary endpoint retained all 60 targets per donor.

To assess precision beyond the initial three draws, a further follow-up added 12 independently randomized reference-draw blocks at B=60 and 300 for balanced and 10:1:1 allocations. All external donors and targets were retained. Each block generated ordered cell lists for every donor and type, reused across reference configurations as above. Monte Carlo standard errors (MCSEs) were calculated across complete blocks, preserving dependence within each block. The prespecified target was an MCSE of 0.05 points for weighted $I(60)$, $I(300)$ and their difference. All 12 blocks were to be completed regardless of interim results (Section S17).

Measured bulk assessment

We assessed matched PBMC RNA-seq and flow-cytometry proportions from 12 healthy adults in the Monaco S13 cohort [12]; the thirteenth participant lacked usable flow measurements. recount3 gene coverage counts for GSE107011/SRP125125 were converted to read-equivalent counts using the reported average mapped length [13]. These measured samples were paired with the external single-cell references under the same MuSiC settings.

The analysis reused all 1,344 references from the 12 additional draw blocks: 14 unique donor triples, B=60 or 300, and balanced or 10:1:1 allocations with every dominant-donor choice. To accommodate annotation

differences, we combined the six predicted types into B cells, T cells, NK cells and monocytes. The corresponding flow proportions represented 83.0–93.5% of reported PBMCs and were normalized within these four groups. MAE was calculated across the four proportions, with equal weighting of participants and reference configurations. The same signed budget contrast and whole-block MCSE were specified before this follow-up’s predictions. Subsequent diagnostics examined bulk input scale, exclusion of RPL/RPS-prefix genes and Salmon re-quantification of one sample, using balanced references at B=300 (Section S19).

Results

The average allocation penalty depended on the cell budget

At B=60 in the original cohort, weighted MAE rose from 5.634 to 5.697, 5.862 and 6.361 points across the four allocation ratios (Table 1). The 10:1:1-minus-balanced contrast was +0.728 points for weighted MuSiC and +0.117 for ordinary NNLS. Weighted donor-average error increased at every allocation step for all six donors, but this pattern held for only 95/180 weighted and 31/180 ordinary arrangements (Figure S1). The weighted-minus-ordinary difference in allocation contrasts was positive for all six donor averages and negative in 38/180 arrangements.

Table 1. Absolute MAE and signed contrasts between 10:1:1 and balanced allocations

Equal-donor means in percentage points. B is the total number of reference cells per type across three donors. $I(B)=MAE(10:1:1)-MAE(1:1:1)$, calculated before rounding. Comparisons are within each source; external rows use the initial grouping and three draws. Tables S10 and S14 report the grouping and sampling follow-ups.

Source	Output	B	1:1:1	2:1:1	4:1:1	10:1:1	I(B)
Original SLE, n=6	Weighted MuSiC	60	5.634	5.697	5.862	6.361	+0.728
Original SLE, n=6	Ordinary NNLS	60	9.110	9.137	9.147	9.227	+0.117
External normal/na, n=14	Weighted MuSiC	60	5.290	5.351	5.480	5.911	+0.621
External normal/na, n=14	Weighted MuSiC	120	5.178	5.220	5.245	5.454	+0.276
External normal/na, n=14	Weighted MuSiC	300	5.280	5.281	5.227	5.214	−0.066
External normal/na, n=14	Ordinary NNLS	60	8.552	8.603	8.619	8.720	+0.167
External normal/na, n=14	Ordinary NNLS	120	8.590	8.624	8.588	8.656	+0.066
External normal/na, n=14	Ordinary NNLS	300	8.684	8.699	8.673	8.654	−0.030


In the external cohort, the mean weighted allocation contrast $I(B)$ decreased from +0.621 to +0.276 to −0.066 points at B=60, 120 and 300 (Figure 1). The primary contrast, $I(300)-I(60)$, averaged −0.687 points, ranged from −1.152 to +0.218 and was negative for 13/14 donors. For the exception, HC-014, $I(B)$ remained negative at −0.434, −0.346 and −0.216 points: unequal allocation had lower error at every budget, but its advantage diminished as the budget grew. RMSE allocation contrasts followed the same direction, at +0.713, +0.344 and −0.119 points (Section S12).

Ordinary NNLS allocation contrasts were +0.167, +0.066 and −0.030 points, with a budget change of −0.198 points. Weighted MuSiC had lower donor-average MAE in every donor and condition in the initial analyses of both cohorts (Table 1; Figure S2). A smaller allocation penalty did not imply lower absolute prediction


error. Increasing the budget did not monotonically improve accuracy: balanced weighted MAE in the external cohort was 5.290, 5.178 and 5.280 points.

Figure 1

a Same three donors; same total cells per type



Allocation contrast: 10:1:1 – balanced




Controlled donor allocation and its budget-dependent effect. a, Design schematic: the same three donors contribute equal or unequal cell numbers at a fixed total budget B per cell type. Colours identify donors; each donor takes the dominant position in turn. b,c, $I(B) = \text{MAE}(10:1:1) - \text{MAE}(\text{balanced})$, in percentage points, for weighted MuSiC and ordinary NNLS. Grey lines show 14 held-out donor means; coloured lines show their equal-weight mean. Each donor value averages 60 targets, three draws, four triples and three dominant-donor choices. d, Weighted $I(300) - I(60)$; the orange dashed line marks its mean. Black dashed lines mark zero. A negative budget contrast indicates a decrease in the signed allocation effect. Donor summaries share reference sources; no population intervals are shown.

Effects in opposite directions persisted at the largest budget

The near-zero mean at B=300 combined effects in opposite directions: 84/168 weighted arrangements still had higher MAE under 10:1:1 allocation (Figure 2). The near-zero average therefore reflected offsetting improvements and deteriorations, not uniformly negligible effects. Mean absolute contrasts decreased from 0.708 to 0.439 to 0.402 points across the three budgets, while the signed 10th–90th percentile spans were 1.211, 1.085 and 1.246 points. The average magnitude thus decreased without a consistent narrowing of the central range. Cell-type effects also differed in direction: at B=300, contrasts included –0.424 points for CD4

T cells and +0.122 for NK cells. Sections S12 and S16 report the complete distributions, allocation-step reversals and conditional SDs.

Figure 2



Weighted allocation effects within held-out donors. B is the total number of reference cells per type across three donors. Panels show B=60, 120 and 300 in the initial external grouping. Each point is $I(B) = \text{MAE}(10:1:1) - \text{MAE}(\text{balanced})$, the difference in mean absolute error in percentage points for one reference triple and dominant donor, averaged over 60 targets and three draws. Each donor contributes 12 arrangements; vertical offsets separate points and have no scientific meaning. All panels share the same horizontal scale, and dashed lines mark zero. Positive values indicate higher error under 10:1:1 allocation. Mean I and mean |I| summarize the 168 signed contrasts and their absolute magnitudes. Shared donors, cells and targets make these descriptive configuration values dependent. Ordinary NNLS distributions are shown in Figure S4.

The budget trend persisted after changing donor grouping and adding draws


To test whether the trend depended on the particular donor triples, we changed the grouping while retaining the existing cell draws. The weighted budget contrast remained negative, averaging -0.635 points, and was negative for 13/14 donors, with HC-014 again positive. At B=300, 81/168 arrangements still had higher error under unequal allocation. Both the average budget trend and configuration-level heterogeneity therefore persisted (Tables S10–S11).

We assessed sensitivity to reference sampling with 12 additional independently randomized draw blocks. The weighted budget contrast averaged -0.665 points and was negative for 13/14 donor averages. All three prespecified MCSE targets were met. At B=300, 84/168 arrangements still had higher error under unequal allocation. Combining all 15 blocks gave a similar budget contrast (Table S14).

Profile variability and weight differences declined with budget

We examined whether profiles estimated from smaller donor samples were less precise. Across the three diagnostic draws, profile variability decreased as the minor-donor quota rose from 5 to 10 to 25 cells in all 84 donor-type groups. The mean allocation contrast in between-donor variance Σ , summed over genes, also decreased, from 1.09727×10^{-4} at $B=60$ to 2.25585×10^{-5} at $B=300$. Total-variation distance between normalized fitted gene weights declined from 0.404 to 0.288. The Σ and weight-distance decreases held for all 14 donor means, but not every configuration (Figure 3; Table S12; Figure S3).

Figure 3



Profile dispersion and fitted-weight differences declined with reference size. a, Profile variance across three draws at minor-donor quotas of 5, 10 and 25 cells per type, averaged over six types within each reference donor; logarithmic vertical axis, squared relative-expression units. b, Total-variation distance between normalized fitted gene weights under 10:1:1 and balanced allocation, restricted to shared genes and averaged over eight targets, draws and reference configurations. Grey lines show 14 donor summaries; orange squares show their equal-weight mean. Panel a describes reference donors; b describes overlapping held-out configurations. Horizontal positions are categorical. These associations support a sampling-precision explanation but do not identify the cause of the error changes. Figure S3 shows between-donor variance contrasts.

Balancing by discarding cells increased weighted error

We also examined a practical alternative to redistributing a fixed budget: balancing donors by discarding cells. Thinning the original 50:5:5 references to 5:5:5 removed 75% of their cells and increased weighted MAE by 0.547 points. Ordinary NNLS changed by -0.046 points, with mixed donor directions. This comparison reduces total reference size as well as changing allocation (Section S10).

Direct transfer to measured bulk did not recover major lineages

In 12 measured PBMC samples with flow-cytometry proportions, balanced weighted MAE was 31.676 points at $B=60$ and 31.672 at $B=300$. B-cell estimates were effectively zero in all 16,128 weighted fits, and T-cell estimates in 16,127. The weighted budget contrast was $+0.0019$ points (MCSE 0.0027). Its near-zero value therefore cannot establish allocation robustness or validate the simulation trend in measured bulk (Section S19; Figure S5).

Input-scale changes and single-sample re-quantification did not recover B and T cells. Excluding RPL/RPS-prefix genes in a single-reference control reduced weighted MAE from 31.672 to 26.741 points across the 12 participants, but T-cell estimates remained zero. These follow-ups left the transfer failure unresolved (Section S19; Table S18; Figure S6).

Discussion

These simulations show that donor-allocation effects depend on the reference-cell budget. Increasing the budget reduced the average penalty of unequal allocation, with the trend persisting after changes to donor grouping and reference draws. Yet opposite effects remained within individual configurations. Reference design should therefore be reported through cells per donor and cell type, alongside the total budget; a near-zero average contrast does not make donor allocation irrelevant. Nor does the result support balancing by discarding cells: in the thinning comparison, discarding 75% of reference cells increased error despite equal donor quotas.

Sampling precision offers a possible explanation. A donor represented by fewer cells may have a less precise expression profile. Under independent, zero-mean profile errors, the expected variance across donor estimates includes the average profile-error variance (Section S15). Declining profile dispersion, Σ contrasts and weight distances are consistent with this account. They do not establish which changes caused the MAE pattern, because allocation also changes the basis, estimated cell sizes, residuals and genes used in fitting. The comparison with ordinary NNLS cannot isolate one weighting component because both outputs share the basis.

The measured-bulk assessment shows why allocation contrasts must be interpreted alongside absolute accuracy. Near-zero contrasts accompanied failure to recover major lymphocyte populations, and the input diagnostics did not resolve that failure. Excluding RPL/RPS-prefix genes partly reduced error, but differences in RNA capture, cell-size estimation and unmodelled populations remain possible explanations. The simulation budget trend therefore remains unestablished in measured bulk; these failed fits also cannot determine whether donor allocation matters when measured samples are deconvolved accurately.

By fixing donor identities and total cells per type, our comparison isolates the effect of changing donor quotas within the tested MuSiC workflow. It complements studies of reference weighting [3,4], reference–mixture matching and donor number [5], cell-type balancing and donor subsampling [6], cell selection [7], and reference contexts and computational costs [14,15] (Table S15). The resulting guidance is conditional: evaluate donor allocation at the intended budget, and examine configuration-level effects as well as the mean. These data do not establish a universally optimal allocation or a sufficient cell budget.

Limitations

Both single-cell cohorts were selected for sufficient cells in every retained type; the external cohort required at least 250 per type. The sources differed in disease and culture context, processing, annotations, gene sets, reference composition and target-cell sampling. Differences in error between cohorts cannot be attributed to any one of these factors. Participant overlap through unrecognized donor aliases also remains unresolved.

The measured-bulk assessment covered one cohort and four aggregated lineages. Normalizing the flow proportions within those lineages does not remove expression from unmodelled cells in the bulk. Broad T-cell labels also leave differences in MAIT and $\gamma\delta$ annotation unresolved. Four-lineage errors are therefore not directly comparable with the six-type simulation errors. The gene-exclusion control used one reference and the quantification pilot one participant; neither tested the allocation-by-budget contrast.

References contained equal cell counts across types, and the same donor dominated all six types. Type-specific donor allocations were not tested, and target compositions were not sampled from empirical PBMC distributions. The 10:1:1 allocation was a stress condition; its practical prevalence and more extreme allocations remain untested.

The additional draws quantify reference-sampling precision conditional on fixed donors, targets and grouping. They do not measure uncertainty from selecting other participants or mixtures. Shared reference sources further limit population-level inference from donor summaries.

These findings apply to the tested MuSiC version and flat workflow; generalization to other deconvolution algorithms requires separate evaluation. We tested four allocation ratios and three external budgets. Each external grouping covered four of 286 possible reference triples per held-out donor, compared with all ten in the original cohort. The alternative grouping reused existing draws; grouping and sampling follow-ups covered only endpoint budgets and allocations. Three unequally spaced budgets cannot establish a linear relationship or a sufficient reference-size threshold.

Profile and weight diagnostics used three draws, sometimes at large sampling fractions, limiting precision and mechanistic interpretation. Plasmacytoid dendritic cells were excluded. Equal budgets across types and non-empirical target compositions also prevent attributing cell-type differences in error to natural cell abundance.

Within these simulations, donor allocation mattered differently at different reference-cell budgets: the average penalty of unequal allocation declined as the budget increased, while effects in individual configurations remained. Reference design therefore warrants evaluation at the intended budget, with attention to both average and configuration-level errors. Whether this pattern transfers to accurately deconvolved measured bulk remains open.

Supplementary material

Additional_file_1.pdf provides supplementary methods and results (Tables S1–S18; Figures S1–S6).

Ethics approval and consent to participate

This study used public count data and coded annotations, with no new participant recruitment or sample collection. Written advice from Dr Alasdair Ivens, University of Edinburgh, indicated that additional ethics approval was not required for secondary use of the two public anonymised single-cell datasets.

Data availability

Original counts and annotations are available under GSE96583. The external CELLxGENE release is associated with GSE174188/GSE137029; its fixed identifiers are given in Section S11. Measured PBMC RNA-seq is available under GSE107011/SRP125125 and through recount3, with matched flow proportions in Monaco et al.'s Table S6 [12]. Processed inputs, saved predictions and results are archived at Zenodo (<https://doi.org/10.5281/zenodo.22902701>). Figure source data are also available on GitHub.

Code availability

MuSiC analysis and plotting code are available at <https://github.com/ReInhard-1/bioinformatics-paper> (release v1.1.0). The corresponding code archive is included in the Zenodo deposit.

Competing interests

The author declares that he has no competing interests.

Funding

This research received no funding.

AI assistance

Codex (OpenAI) assisted with literature retrieval and organisation, debugging and modification of code written by the author, and language editing of existing text. The author reviewed and verified the AI-assisted work and takes full responsibility for the code, results, references, and final manuscript.

References

1. Wang X, Park J, Susztak K, Zhang NR, Li M. Bulk tissue cell type deconvolution with multi-subject single-cell expression reference. *Nat Commun.* 2019;10(1):380. doi:10.1038/s41467-018-08023-x.
2. Avila Cobos F, Alquicira-Hernandez J, Powell JE, Mestdagh P, De Preter K. Benchmarking of cell type deconvolution pipelines for transcriptomics data. *Nat Commun.* 2020;11(1):5650. doi:10.1038/s41467-020-19015-1. Associated funding correction: doi:10.1038/s41467-020-20288-9.
3. Dong M, Thennavan A, Urrutia E, Li Y, Perou CM, Zou F, et al. SCDC: bulk gene expression deconvolution by multiple single-cell RNA sequencing references. *Brief Bioinform.* 2021;22(1):416-427. doi:10.1093/bib/bbz166.
4. Liu G, Liu X, Ma L. DecOT: Bulk Deconvolution With Optimal Transport Loss Using a Single-Cell Reference. *Front Genet.* 2022;13:825896. doi:10.3389/fgene.2022.825896.
5. Zhao Y, Vestal BE, Moore CM. Evaluating reference-mixture matching in cell-type deconvolution with single-cell RNA-seq references. *Brief Bioinform.* 2026;27(2):bbag184. doi:10.1093/bib/bbag184.
6. Huuki-Myers LA, Montgomery KD, Kwon SH, Cinquemani S, Eagles NJ, Gonzalez-Padilla D, et al. Benchmark of cellular deconvolution methods using a multi-assay dataset from postmortem human prefrontal cortex. *Genome Biol.* 2025;26(1):88. doi:10.1186/s13059-025-03552-3.
7. Huang L, Gong W, Chen D. scValue: value-based subsampling of large-scale single-cell transcriptomic data for machine and deep learning tasks. *Brief Bioinform.* 2025;26(3):bbaf279. doi:10.1093/bib/bbaf279.
8. Kang HM, Subramaniam M, Targ S, Nguyen M, Maliskova L, McCarthy E, et al. Multiplexed droplet single-cell RNA-sequencing using natural genetic variation. *Nat Biotechnol.* 2018;36(1):89-94. doi:10.1038/nbt.4042.
9. Kang HM, Subramaniam M, Targ S, Nguyen M, Maliskova L, McCarthy E, et al. Author Correction: Multiplexed droplet single-cell RNA-sequencing using natural genetic variation. *Nat Biotechnol.* 2020;38(11):1356. doi:10.1038/s41587-020-0715-9.
10. Perez RK, Gordon MG, Subramaniam M, Kim MC, Hartoularos GC, Targ S, et al. Single-cell RNA-seq reveals cell type-specific molecular and genetic associations to lupus. *Science.* 2022;376(6589):eabf1970. doi:10.1126/science.abf1970.
11. Erratum for the Research Article “Single-cell RNA-seq reveals cell type-specific molecular and genetic associations to lupus” by R. K. Perez et al. *Science.* 2024;385(6705):eadr4064. doi:10.1126/science.adr4064.
12. Monaco G, Lee B, Xu W, Mustafah S, Hwang YY, Carré C, et al. RNA-Seq Signatures Normalized by mRNA Abundance Allow Absolute Deconvolution of Human Immune Cell Types. *Cell Rep.* 2019;26(6):1627-1640.e7. doi:10.1016/j.celrep.2019.01.041.
13. Wilks C, Zheng SC, Chen FY, Charles R, Solomon B, Ling JP, et al. recount3: summaries and queries for large-scale RNA-seq expression and splicing. *Genome Biol.* 2021;22:323. doi:10.1186/s13059-021-02533-6.
14. Dietrich A, Merotto L, Pelz K, Eder B, Zackl C, Reinisch K, et al. omnideconv: a unifying framework for using and benchmarking single-cell-informed deconvolution of bulk RNA-seq data. *Genome Biol.* 2026;27:6. doi:10.1186/s13059-026-03955-w.
15. Malik AA, Alif MN, Bratton A, Sun J, Li Q, Zhang W. Benchmarking reference-based cellular deconvolution algorithms to predict cell proportions. *Comput Biol Chem.* 2026:109370. Published online 15 September 2026; journal pre-proof. doi:10.1016/j.compbiolchem.2026.109370.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Additionalfile1.pdf](#)