RIPPLE: A SELF-DIMENSIONING FACTOR MODEL FOR SINGLE-CELL TRANSCRIPTOMICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Single-cell atlases are becoming biology's reference maps, and almost every downstream analysis runs on a low-dimensional embedding of the counts rather than on the counts themselves. Such an analysis must discover the cell types and states of interest, so the representation has to be learned without labels. Two choices govern it and are usually left unprincipled: how many factors the representation carries, and whether its coordinates mean anything. We present RIPPLE, an unsupervised negative binomial factor model that settles both from one design. A sparsity-inducing half-Cauchy prior on the per-factor loading scales collapses the factors the data do not support, and the active dimension is read off the fitted posterior by a signal-to-noise rule with no threshold to tune. The decoder is linear, so every factor that survives is a gene program that can be read directly. Under mild conditions, we prove that this rule loses no true factor, and that the fitted signal carried beyond the true factors is asymptotically weightless. On a six-dataset, eight-method benchmark, RIPPLE outperforms every unsupervised baseline on every biology-preserving metric averaged over the datasets. On the 585k-cell Human Lung Cell Atlas, a single selected factor separates each of the 61 annotated cell types at a median AUROC of 0.96, independently of how rare the cell type is.

1 INTRODUCTION

Single-cell RNA sequencing has transformed biology by resolving the cell states that bulk assays average away (Tang et al., 2009; Macosko et al., 2015). Droplet protocols now profile hundreds of thousands of cells in one experiment (Zheng et al., 2017), and the resulting *cell atlases*, such as the Human Lung Cell Atlas (HLCA), are reference maps against which new samples are read (Regev et al., 2017; Sikkema et al., 2023; CZI Cell Science Program et al., 2025). The measured counts, however, make a poor coordinate system. They are extremely sparse, and their zeros mix true biological absence with technological dropouts (Kharchenko et al., 2014; Hicks et al., 2018). The nonzero counts are over-dispersed, misspecifying Poisson models and making the negative binomial (NB) the standard alternative (Grün et al., 2014; Svensson, 2020). Furthermore, cells from different batches carry non-biological technical variation, rendering distances in the raw space incomparable (Korsunsky et al., 2019). The biology of interest is also far lower-dimensional than the tens of thousands of measured genes, which act in co-regulated programs rather than independently (Stuart et al., 2003; Langfelder & Horvath, 2008; Heimberg et al., 2016). Almost every downstream analysis, from grouping cells into cell types and discovering new states to characterizing disease and annotating a freshly sequenced cohort, is therefore answered not on the raw gene counts but on a low-dimensional *embedding* of them (Luecken & Theis, 2019).

Constructing that embedding is the hard part, and it must be done without supervision: the analysis is meant to discover the cell types and states themselves. Cast as a machine-learning problem, this is unsupervised latent-variable modeling of high-dimensional count data, where the pathologies above render traditional factorization methods unfit for the task and an explicit count likelihood is needed (Townes et al., 2019). Two fundamental challenges remain in the current literature. (i) *The latent dimension is unknown*: it is fixed by hand or by heuristic before fitting, with no guarantee the choice matches the true dimension (Kiselev et al., 2019). (ii) *The embedding must be readable*: biological discovery proceeds by asking which genes distinguish one population from another, so a factor is useful only if it names a concrete gene program.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Existing methods meet these needs only in part. *Amortized deep models* carry high representation power, but their deep decoders leave the factors uninterpretable. Specifically, hidden layers mix the embedding coordinates before they reach genes, so no factor carries a fixed gene program. Furthermore, the latent dimension must be determined a priori with no principled guidance, and changing it requires a full refit. Among them, scVI (Lopez et al., 2018) trains an NB variational autoencoder, and its semi-supervised extension scANVI (Xu et al., 2021) buys accuracy with cell type labels that are unavailable in most applications. scVIC (Xiong et al., 2024) places a Gaussian-mixture prior on the latent space, while scMAE (Fang et al., 2024) drops the count likelihood for a masked-autoencoder objective. *Linear factor models* do not suffer from the interpretability issue, since a linear decoder makes every factor a loading vector over genes, allowing it to be read directly as a gene program. Nonetheless, they also require the latent dimension to be fixed a priori. Among them, ZINB-WaVE (Risso et al., 2018) models counts with a zero-inflated NB likelihood, while scGBM (Nicol & Miller, 2025) uses a Poisson likelihood. Outside both families, Harmony (Korsunsky et al., 2019) corrects a PCA embedding with batch-specific shifts, so the corrected axes are no longer a fixed linear map from genes and thus not directly interpretable. None of these methods is at once self-dimensioning, with the latent dimension driven by the data rather than by a heuristic, and able to return an interpretable embedding accurate enough for downstream biological discovery.

We introduce **RIPPLE**, an NB factor model that meets both challenges with no dimension to tune. A half-Cauchy prior on the per-factor loading scales shrinks the factors the data do not support. The active factor set is read off once after convergence, by comparing each factor's spread across cells to the posterior noise the model already fits for it. Since that noise floor comes from the fit, no threshold has to be tuned. The decoder is *linear* in the factors, so each row of the loading matrix is a gene program a biologist can read, and a factor that survives selection is a program the data supported. A neighborhood-graph regularizer completes the model, pulling each cell toward its nearest neighbors in the embedding.

Shrinkage of this kind is well understood for Gaussian factor models (Bishop, 1998; Minka, 2000; Carvalho et al., 2010; Piironen & Vehtari, 2017). For covariate-free count models, Chen & Li (2022) prove that the number of factors can be selected consistently by comparing a penalized joint likelihood across candidate ranks. Their guarantee concerns the natural-parameter matrix rather than the factors, which they note are not identifiable under rotation. We instead allow covariates and add factor-level guarantees: every true factor lies in the selected span and has a distinct selected counterpart. Under mild conditions, we show that RIPPLE's selection rule loses no true factor, and the fitted signal it carries beyond the true factors is asymptotically weightless (Section 3 and Appendix E). Empirically (Section 4), on six datasets and eight methods, RIPPLE attains the highest average score of any unsupervised method on every biology-preserving clustering metric. Its embedding also surpasses that of scANVI, which is trained on the cell-type labels RIPPLE never sees. A neighborhood-regularized unsupervised linear factor model that picks its own dimension is therefore sufficient to match or beat the amortized deep models in common use. On the 585k-cell HLCA, it selects a compact factor set whose members separate cell types and group them by lineage, and whose decoder reproduces donor-pseudobulk differential expression by rank. Every one of the 61 annotated cell types is separated by a single factor, so the rarest populations survive the reduction rather than being averaged into the majority. RIPPLE is open-sourced at `https://anonymous.4open.science/r/ripple-61C3`.

1.1 Notation

For a matrix X, we write $X_{i\cdot}$ for its ith row, $X_{\cdot j}$ for its jth column, and X_{ij} for the (i, j) entry. For a vector θ, θ_j denotes the jth entry. A hat denotes a fitted quantity and a superscript zero the truth, so $\widehat{W}$ and $\widehat{\alpha}$ estimate W^0 and α^0. $[n]$ denotes the set $\{1, \dots, n\}$, and $\mathbb{Z}_{\geq 0}$ the nonnegative integers.

2 The RIPPLE model

RIPPLE models the scRNA-seq count matrix $Y \in \mathbb{Z}_{\geq 0}^{n \times p}$ over n cells and p genes as a mean-dispersion parametrized NB factor model. For cell $i \in [n]$ and gene $j \in [p]$, the log-mean gene expression $\log(\mu_{ij})$ is modeled as

$$\log \mu_{ij} = X_{i\cdot}^\top \beta_{\cdot j} + \gamma_{i\cdot}^\top V_{j\cdot} + W_{i\cdot}^\top \alpha_{\cdot j} + O_i, \qquad Y_{ij} \mid \mu_{ij}, \theta_j \sim \mathrm{NB}(\mu_{ij}, \theta_j). \tag{1}$$

Most importantly, the *scores* $W_{i\cdot} \in \mathbb{R}^L$, i.e., the rows of the *embedding* $W \in \mathbb{R}^{n\times L}$, are cell i's coordinates in the latent space every downstream analysis runs on, and the *factors* $\alpha_{\ell\cdot} \in \mathbb{R}^p$, i.e., the rows of the *loading matrix* $\alpha \in \mathbb{R}^{L\times p}$, are gene-space vectors, which can be read directly as gene programs. The latent dimension L is an upper bound, and the sparsity-inducing prior in Eq. (2) makes the model self-dimensioning (Appendix D.1). The other terms can absorb non-biological variation, helping $W\alpha$ represent the biology of interest. Cell covariates $X \in \mathbb{R}^{n\times H}$ hold nuisance cell information, such as experimental batch, while gene covariates $V \in \mathbb{R}^{p\times G}$ hold nuisance gene information, such as GC content. The cell coefficients $\beta \in \mathbb{R}^{H\times p}$ are gene-specific, since a batch does not shift every gene by the same amount. Similarly, the gene coefficients $\gamma \in \mathbb{R}^{n\times G}$ are cell-specific, since a gene property need not act equally in every cell. The offset $O_i = \log(d_i/\bar{d})$ rescales each cell by its total count, or library size d_i. The dispersion θ_j is gene-specific, since genes have different levels of variability.

Sparsity-inducing prior. Unlike existing factor models, which require the unknown latent dimension L to be set a priori, we let the model infer it from the data. Specifically, RIPPLE over-provisions L factors and lets a shrinkage prior collapse the ones the data do not support:

$$\lambda_\ell \sim \mathrm{HalfCauchy}(1), \qquad \alpha_{\ell j} \mid \lambda_\ell \sim \mathcal{N}(0, \lambda_\ell^2), \qquad W_{i\ell} \sim \mathcal{N}(0,1), \tag{2}$$

with weakly-informative priors on the remaining parameters (Appendix A.1). The $\mathrm{HalfCauchy}(1)$ prior concentrates its mass near zero, shrinking spurious factors toward zero. Its heavy polynomial tails also allow factors to escape the shrinkage when the data support them. Specifically, once the likelihood pulls λ_ℓ away from zero, the regularization attenuates in the tails, leaving the active factor essentially unpenalized with asymptotically unbiased shrinkage at $O(1/p)$ (Appendix A.4). In contrast, a light-tailed prior on the scale pulls harder the larger λ_ℓ grows, so the shrinkage grows with the strength of the factor instead of fading away. The shrinkage is also reversible, since the objective is smooth in $\log\lambda_\ell$, a factor shrunk small early can re-inflate once the data support it, unlike a discrete spike-and-slab. In addition to inducing sparsity, the per-factor scales λ_ℓ in Eq. (2) break the rotational invariance of the model. Since the likelihood depends on W and α only through their product $W\alpha$, it is invariant under orthogonal rotations. However, the distinct λ_ℓ's render the prior anisotropic, breaking this invariance and picking out a particular basis for the factor decomposition.

Variational inference. We fit the model by mean-field variational Bayes (VB), giving each parameter a Gaussian variational distribution with variational mean $m_\bullet$ and variance $s^2_\bullet$, while λ and γ enter as point estimates. Writing ϕ for the free variational parameters, we optimize the penalized evidence lower bound (ELBO),

$$\mathcal{L}(\phi) = \sum_{i,j} \mathbb{E}_q\big[\log \mathrm{NB}(Y_{ij} \mid \mu_{ij}, \theta_j)\big] - \underbrace{\big[\mathrm{KL}\big(q(\alpha) \,\|\, p(\alpha \mid \lambda)\big) - \log p(\lambda)\big]}_{\text{sparsity-inducing}} - \mathrm{KL}_{\mathrm{rest}} - \mathcal{R}, \tag{3}$$

where the braced term is the sparsity machinery, $\mathrm{KL}_{\mathrm{rest}}$ collects the remaining Kullback–Leibler terms, and $\mathcal{R}$ is the neighborhood penalty defined next (full details in Appendix A.4 and A.5). We isolate λ from the variational distribution because the asymptotically unbiased shrinkage holds at a stationary λ_ℓ, which a point estimate reaches exactly. Furthermore, no downstream analysis requires the posterior of λ, so nothing is lost.

To avoid an amortization gap (Cremer et al., 2018), we model the embedding $W \in \mathbb{R}^{n\times L}$ at cell level instead of having cells share an encoder, which is common in recent literature for scalability. In fact, RIPPLE remains scalable to the full HLCA core atlas (Section 4.2), and the $O(nL)$ variational parameters of cell-level inference do not dominate the fit (Appendix B.5). We achieve this scalability since the penalized ELBO in Eq. (3) decomposes over cells and admits minibatch stochastic optimization on a GPU. Exact posterior inference by sampling, such as MCMC, is far more costly (Blei et al., 2017) and does not reach the scale of a modern atlas.

Neighborhood regularizer. Each $W_{i\cdot}$ is fit from cell i's own counts alone, which might be sequenced shallowly, leading to a noisy $W_{i\cdot}$ estimate even when its neighbors are unambiguous. To circumvent this issue, we let each cell borrow strength from similar cells by pulling each toward the mean embedding of its neighbors on a k_{nn}-nearest-neighbor (kNN) graph,

$$\mathcal{R} = \lambda_g \sum_i \big\| \, m_{W_{i\cdot}} - \mathrm{sg}\big[\, \overline{m}_{\mathcal{G}(i)} \,\big] \, \big\|_2^2, \tag{4}$$

where $\lambda_g > 0$ is the penalty weight, $\mathcal{G}(i)$ the neighbor set of cell i, $\overline{m}_{\mathcal{G}(i)}$ their mean embedding, and $\mathrm{sg}[\cdot]$ stops the gradient. Two choices keep this from degenerating. The graph is built once from

the data and then held fixed, preventing the model from rewiring its own neighborhoods to suit its errors. The target $\overline{m}_{\mathcal{G}(i)}$ is also stop-gradient, keeping neighbors from attracting each other and the embedding from collapsing. In sum, the regularizer $\mathcal{R}$ sharpens the embedding where cell types separate weakly, and pays a cost where they already separate cleanly (Appendix D.2).

Batch integration. When the data span $b = 1, \ldots, B$ batches, a batch dummy in $X_{i\cdot}$ buys each batch a free per-gene shift. That fixed effect is unpooled, however, and grows as $B \times p$, which is wasteful when batch effects trace only a few technical axes shared across batches. When batches must be integrated, we therefore add an optional low-rank random effect to Eq. (1), $\log \mu_{ij}$ += $Z^{\top}_{b(i)\cdot}\delta_{\cdot j}$ with

$$\lambda^b_q \sim \text{HalfCauchy}(1), \qquad \delta_{qj} \mid \lambda^b_q \sim \mathcal{N}\big(0, (\lambda^b_q)^2\big), \qquad Z_{bq} \sim \mathcal{N}(0,1), \tag{5}$$

where $b(i)$ is the batch of cell i and $q \in [Q]$ indexes the low-rank random effect subspace. Note that the random effect is applied to batches rather than cells, and the priors in Eq. (5) mirror the same two roles as in Eq. (2). The unit normal on $Z_{b\cdot}$ anchors the scale gauge, so a batch axis carries its magnitude in δ rather than in Z. The half-Cauchy on λ^b_q then shrinks the axes the batches do not need, so Q sets a ceiling on the subspace rather than its rank. Giving batch structure a home at the batch level helps keep batch effects out of the embedding W. It is paired with a cross-batch variant of $\mathcal{G}(i)$ that takes each cell's nearest neighbors within every batch, computed in a batch-centered space, then keeps the globally nearest of those, so the target averages cells across batches rather than within one (Appendix A.6). The two do complementary work: the random effect makes batches comparable, while the redirected graph makes them overlap. Appendix D.4 details an ablation study comparing different batch-correction strategies.

Reading off the active set. During training, the half-Cauchy prior shrinks the factors the data do not support, so they do not affect the quality of the fit. The active set is then selected after convergence. For factor ℓ, we compare the spread of its fitted score means $m_{W\cdot\ell} \in \mathbb{R}^n$ across cells to the posterior noise, which collectively represent the total variance of $W_{\cdot\ell}$:

$$\text{SNR}_\ell = \frac{\widehat{\text{Var}}(m_{W\cdot\ell})}{\frac{1}{n}\sum_i s^2_{W i\ell}}, \qquad \ell \text{ active} \iff \text{SNR}_\ell > 1, \tag{6}$$

with $\widehat{\mathcal{S}} = \{\ell : \text{SNR}_\ell > 1\}$ denoting the active set and $L_{\text{active}} = |\widehat{\mathcal{S}}|$ its cardinality. When the numerator of SNR_ℓ is small, cells have similar score means $m_{W\cdot\ell}$ in the ℓth dimension, rendering factor ℓ uninformative for characterizing or separating cells. The denominator is the model's own uncertainty about dimension ℓ averaged over cells, and it sets the noise floor the numerator is compared against. Note that $\text{SNR}_\ell > 1$ is equivalent to the reliability $\text{SNR}_\ell/(1+\text{SNR}_\ell) > \frac{1}{2}$, so factor ℓ is kept when more than half of the total variance of $W_{\cdot\ell}$ is signal rather than posterior noise. Because the noise floor is supplied by the fit itself, the cutoff needs no tuning to the dataset. On the HLCA atlas of Section 4, L_{active} remains robust as the ceiling L rises.

3 Theoretical guarantees

As $W\alpha$ is the centerpiece of RIPPLE, the natural target is $W^0\alpha^0$, but it is not identifiable. Since Eq. (1) models $\log(\mu_{ij})$ as a sum of terms, a factor aligned with the columns of X can therefore be absorbed into $X\beta$, and one aligned with the columns of V into $\gamma V^{\top}$, leaving the distribution of Y unchanged. Two different products $W\alpha$ then give Y the same distribution. Without oracle information on β and γ, no estimator is thus consistent for $W^0\alpha^0$.

The identifiable quantity is thus the part of $W\alpha$ that the covariates X and V cannot produce. For notational simplicity, we write $M := W\alpha$ and let $M^{\perp}$ denote the residual of M after least-squares removal of $X\beta$ and $\gamma V^{\top}$. By the Frisch–Waugh–Lovell identity (Lemma 1 of Appendix E), $M^{\perp} \equiv P^{\perp}_X M P^{\perp}_V$ with $P^{\perp}_X$ and $P^{\perp}_V$ projecting off the column spans of X and V. Furthermore, let $L^0 \le L$ be the rank of $M^{0\perp}$, the true latent dimension beyond what the covariates can produce.

Proposition 1 (the design-residual signal is the identifiable target)**.** *Let $C \in \mathbb{R}^{H\times L}$ be a shift matrix such that the shifted parameters stay inside a compact set (Appendix E). Consider two parameter settings θ^0 and θ^C, where θ^0 denotes the truth and θ^C its translate: $W^0 \mapsto W^0 + XC$ and $\beta^0 \mapsto \beta^0 - C\alpha^0$. Both θ^0 and θ^C produce the same distribution of Y at every (n,p) while their signals*

differ by $XC\alpha^0$, so no estimator T can be close to both,

$$\max_{\theta\in\{\theta^0,\theta^C\}} \tfrac{1}{np}\,\mathbb{E}_\theta\|T - W(\theta)\alpha(\theta)\|_F^2 \;\geq\; \tfrac{1}{4np}\|XC\alpha^0\|_F^2.$$

The right-hand side does not vanish as n and p grow, so no estimator is consistent for $M^0 := W^0\alpha^0$. The design-residual signal is the exception since $P_X^\perp W(\theta^C)\alpha(\theta^C)P_V^\perp = M^{0\perp}$ for every C.

Identifiability fixes the estimand. The design-residual signal $M^{0\perp}$ is the largest functional of M any method could certify (Lemma 7), and the same holds for the symmetric substitution on the gene covariates V. Moving a factor into $X\beta$ is the low-rank regressor problem of the interactive-fixed-effects literature (Bai, 2009; Moon & Weidner, 2017), whose rank budget does not apply here because $\beta \in \mathbb{R}^{H\times p}$ is a matrix in RIPPLE rather than a coefficient vector shared across genes. Every guarantee below compares $\widehat{M}^\perp$ with $M^{0\perp}$ under assumptions standard in the factor-model and variational-inference literature, stated in Appendix E along with the proofs and discussion.

Let $M^{0\perp} = \sum_{k\leq L^0} d_k u_k v_k^\top$ be the singular value decomposition (SVD) of $M^{0\perp}$ with $d_1 \geq \cdots \geq d_{L^0} > 0$. We call $d_k u_k v_k^\top$ the kth true factor, where $u_k \in \mathbb{R}^n$ and $v_k \in \mathbb{R}^p$ are the score and loading directions, and d_k^2 its strength. On the fitted side, $\widehat{\mathcal{U}} = \operatorname{span}\{P_X^\perp \widehat{W}_{\cdot\ell} : \ell \in \widehat{\mathcal{S}}\}$ denotes the span of the selected factors' scores with the design partialled out, and $\widehat{M}_{\widehat{\mathcal{S}}}^\perp$ the design-residual signal rebuilt from the selected factors alone.

Theorem 1 (consistency of the factor signal)**.** *Under Assumptions 1–4, for any fixed $\lambda_g \geq 0$,*

$$\tfrac{1}{np}\|\widehat{M}^\perp - M^{0\perp}\|_F^2 \xrightarrow{p} 0 \qquad \textit{and} \qquad \tfrac{1}{np}\|\widehat{M}_{\widehat{\mathcal{S}}}^\perp - M^{0\perp}\|_F^2 \xrightarrow{p} 0.$$

That is, the fitted signal $\widehat{M}^\perp$ is consistent. Furthermore, the factors RIPPLE drops carry no share of the true signal in the limit since $\widehat{M}_{\widehat{\mathcal{S}}}^\perp$ is also consistent.

Beyond signal-level consistency, Theorem 2 shows that RIPPLE's selection loses no true factor. It needs one more assumption, i.e., every true factor is strong enough to be estimated (Assumption 5), the strength condition standard in the factor-model literature (Chen & Li, 2022).

Theorem 2 (no true factor is lost)**.** *Under Assumptions 1–5, as $n, p \to \infty$:*

(a) $\mathbb{P}(L_{\text{active}} \geq L^0) \to 1$*;*

(b) every true factor lies asymptotically in the span of the selected factors,

$$\max_{k\leq L^0} \operatorname{dist}\left(u_k,\, \widehat{\mathcal{U}}\right) \xrightarrow{p} 0,$$

where $\operatorname{dist}(u, \mathcal{V}) = \|u - P_{\mathcal{V}}u\|$ is the distance from a unit vector to a subspace;

(c) $\sum_{k>L^0} \sigma_k^2(\widehat{M}^\perp)/\|\widehat{M}^\perp\|_F^2 \xrightarrow{p} 0$*, where $\sigma_k(\cdot)$ denotes the k-th singular value;*

(d) moreover, under the pervasive form of Assumption 5, some injective assignment π of true factors to selected factors has $|\operatorname{corr}(u_k, \widehat{W}_{\cdot\pi(k)})| \geq \tau$ for every $k \leq L^0$, with $\tau > 0$ a worst-case constant depending only on L and the constants of Appendix E.

Since dropping an informative factor is destructive, we design the SNR selection rule to err toward keeping and show in Theorem 2 that carrying a weak factor is cheap for RIPPLE. In particular, part (a) shows that RIPPLE selects no fewer factors than the truth has, part (b) proves that no true direction is missing from the selected span $\widehat{\mathcal{U}}$, part (c) states that the fitted signal beyond the true rank is asymptotically weightless, and part (d) that each true factor has a distinct selected counterpart correlated with it at a level bounded away from zero.

However, reading factors one at a time asks for more than the product, since a rotation of W and α leaves $W\alpha$ unchanged while mixing the factors. The SVD is the canonical repair: if consecutive strengths are separated, the SVD of $\widehat{M}^\perp$ recovers each true factor $d_k u_k v_k^\top$ (Proposition 2, Appendix E). This canonical basis is chosen by variance, not by biology. Nothing in Eq. (1) makes gene programs orthogonal or ordered by strength, so two programs of similar strength can share one canonical factor. We therefore read the fitted factors, which the sparsity prior scales and selects one at a time, and Section 4 audits that reading directly through reproducibility across seeds and robustness to marker-gene removal.

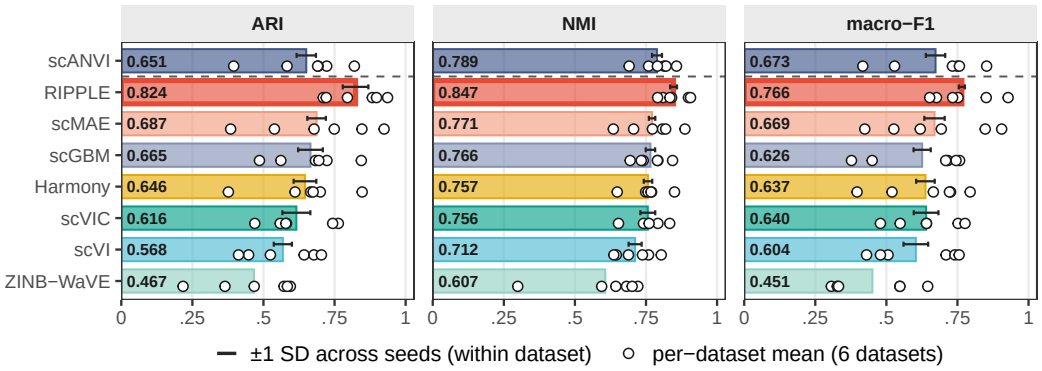


Figure 1: Benchmark performance, averaged over six datasets and ten seeds. Bars denote grand mean at oracle k, the error bar is ± 1 SD across seeds pooled within dataset, and the open circles represent per-dataset means. scANVI alone, above the dashed rule, is trained with cell-type labels.

4 Experiments

4.1 Benchmark: RIPPLE leads on every biology-preserving metric

We benchmark on six public scRNA-seq datasets, covering three tissues and five sequencing protocols: pancreas by inDrop (Baron et al., 2016), CEL-Seq2 (Muraro et al., 2016) and full-length Smart-seq2 (Segerstolpe et al., 2016); mouse brain by STRT-seq (Zeisel et al., 2015); and blood by 10x Chromium, as FACS-purified populations mixed in equal proportion (Zhengmix8eq, Duò et al. (2018)) and as unsorted peripheral blood (PBMC, Zheng et al. (2017)). Their sizes range from $n = 2{,}122$ to $11{,}989$ cells and their annotations from 7 to 14 cell types, with the rarest cell type covering 0.08–9.96% of cells. Their batch structure also differs. Zhengmix8eq and PBMC are single-batch, while Baron and Muraro each carry 4 donor batches. Segerstolpe's ten donors are confounded with disease status and Zeisel is split across 76 small plates, so no method integrates batches on these two datasets. Details on sources and preprocessing are deferred to Appendix B.

RIPPLE and the seven baselines described in Section 1 are trained on their recommended or default settings (Appendices A.8, B.2), where scANVI alone receives cell-type labels and is thus treated as a supervised ceiling rather than a competitor. On each dataset, every method but ZINB-WaVE is run ten times with different random seeds. ZINB-WaVE alone is run once per dataset, since one fit takes 1.5 to 41.4 hours and it ranks last on every metric at the grand mean by a wide margin. We then evaluate each method's embedding by clustering it with k-means at the true number of cell types (*oracle* k), so the comparison scores the embedding rather than a choice of k.

We report three metrics that reflect different aspects of the clustering quality: the adjusted Rand index (ARI) counts co-clustered pairs and corrects for chance; normalized mutual information (NMI) measures the information one partition carries about the other; and macro-F_1 matches clusters to cell types by Hungarian assignment before averaging F_1 over cell types. Figure 1 shows that RIPPLE attains the highest grand mean on all three metrics, at 0.824 ARI, 0.847 NMI and 0.766 macro-F_1, ahead of the strongest unsupervised baseline by 0.137, 0.076 and 0.097, respectively. Per dataset, RIPPLE ranks first among the unsupervised methods on 5 of 6 datasets by ARI, 6 of 6 by NMI and 4 of 6 by macro-F_1 (Appendix B.4). A Friedman test over the six datasets rejects equality of the seven unsupervised methods on all three metrics ($p \leq 0.006$), with RIPPLE holding the best mean rank on all three metrics. Compared with scANVI's embedding, RIPPLE wins on all three metrics at grand average. Per dataset, RIPPLE wins on 4 of 6 datasets, regardless of the metric used. We compare with scANVI's embedding instead of its cell-type classification head here since the embedding serves downstream analysis beyond cell-type clustering, e.g., differential expression of Section 4.2. The embedding is also the only object every method here produces, so it is the only fair comparison (Appendix B.9). The same ordering holds on the four further metrics of Appendix B.6: accuracy, AMI, and the two clusterer-independent scores, label-ASW and cLISI of scIB (Luecken

et al., 2022). Label-ASW is the most informative of them since it does not depend on a clusterer, and RIPPLE leads it at 0.728 against 0.655 for the strongest unsupervised baseline.

RIPPLE's lead also does not depend on the cell and gene covariates, which not every baseline can utilize. An ablation study in Appendix D.3 shows that stripping X and V costs 0.085 in grand mean ARI and still leaves RIPPLE ahead of all seven baselines, at 0.739 against 0.687 of the next-best method. Although not the fastest method, RIPPLE is quite efficient. Among the five methods trained on a GPU, RIPPLE ranks third in speed, averaging 263.3s per fit, against 3.9s for scMAE, 166.5s for scVIC, 268.5s for scVI, and 455.4s for scANVI.

On the two batch datasets, Baron and Muraro, we observe a tradeoff between biology conservation and batch integration (Appendix B.7). The best batch integrator, scVI, reaches a dataset average of 0.477 on integration LISI (iLISI), but scores second-to-last among all unsupervised methods on cell type clustering, at 0.601 ARI. RIPPLE, on the other hand, only integrates batches as far as the biology requires, reaching 0.352 on iLISI (fourth out of seven) while scoring 0.866 on ARI (first out of seven). The same observation holds on other metrics. On a min-max-normalized composite of the two axes, RIPPLE leads the unsupervised methods in bio-conservation (0.985 over seven metrics) while scVI leads in batch integration (0.948 over four metrics). However, scVI's lead on integration comes at a steep cost to bio-conservation, where it scores 0.435, the second-lowest. On the contrary, RIPPLE's lead on bio-conservation comes at a modest cost to integration, scoring 0.650 on batch integration and ranking fourth out of seven.

4.2 Atlas: reading 585k cells with 24 factors

The Human Lung Cell Atlas (HLCA) is the field's consensus map of the lung, harmonizing 14 datasets and 107 donors into 61 annotated cell types, nested into lineages within four tissue compartments and spanning three orders of magnitude in abundance (Sikkema et al., 2023). That combination is what makes it the hard case: a representation must hold cell types apart across studies, protocols and donors, and must do so for populations of a few dozen cells as well as for populations of tens of thousands. To show what RIPPLE recovers at that scale, we fit it to the HLCA core: $584{,}942$ cells over a panel of $3{,}000$ highly variable genes (HVGs). As in Section 4.1, no cell-type label reaches the model during training, and the annotations are used only for evaluation. Appendix C describes the preprocessing pipeline and the HVG selection.

At the reported reference fit of $L = 80$, Figure 2A shows the best non-selected factor at $\mathrm{SNR} = 0.47$ against 1.02 for the last one kept. Furthermore, L_{active} varies by at most one across ceilings L from 60 to 120, at 23, 24, 24 and 23, and never approaches the ceiling itself. In addition, we find that the factors themselves are consistent across L. For any two loading matrices α_1 and α_2, we Hungarian-match their factors (rows) and compute the sign-free cosine similarity, $|\cos|$, between the matched factors. Compared to the reference fit, all three $L = 60, 100, 120$ fits have near identical factors, with mean $|\cos|$ of 0.989, 0.996, and 0.994 respectively. A user may therefore over-specify L and let the prior collapse what the data do not support, provided L is loose enough to leave slack (Figure 2A). The baselines, on the other hand, cannot. Every baseline that learns a latent space of fixed dimension requires a full refit at each candidate L. At the reported setting of $L = 80$, we further verify the robustness of RIPPLE by examining L_{active} fluctuation and the stability of the loading matrix α across ten random seeds. Not only does L_{active} remain at 24 across all seeds, but the mean matched $|\cos|$ exceeds 0.998 in all $\binom{10}{2} = 45$ pairwise comparisons, with no individual factor below 0.99.

A single factor separates each of the 61 annotated cell types from the rest. For each cell type we take the factor that separates it best and score that separation by a 1-vs-rest area under the ROC curve (AUROC), defined in Appendix C. Over the 61 cell types the AUROC runs from 0.814 to 0.999, with a median of 0.959 and 52 above 0.90. Figure 2B shows that this does not come at the expense of the rare cell types. The Spearman correlation between $\log n$ and AUROC is 0.03 ($p = 0.80$). The three rarest cell types, at $n = 28$, 60 and 159, separate at 0.971, 0.918 and 0.958.

We next ask whether a cell type still separates when the canonical markers, i.e., genes that define it, are gone. Here, we analyze three important cell types: ionocyte ($n = 561$), pulmonary neuroendocrine cells (PNECs, $n = 159$), and lymphatic EC mature (LEC mature, $n = 4{,}001$). Ionocytes are rare airway epithelial cells and the airway's highest per-cell expressers of the chloride channel mutated in cystic fibrosis, the most common life-shortening inherited disease in people of European ancestry (Montoro et al., 2018; Plasschaert et al., 2018; Elborn, 2016). PNECs are also rare airway

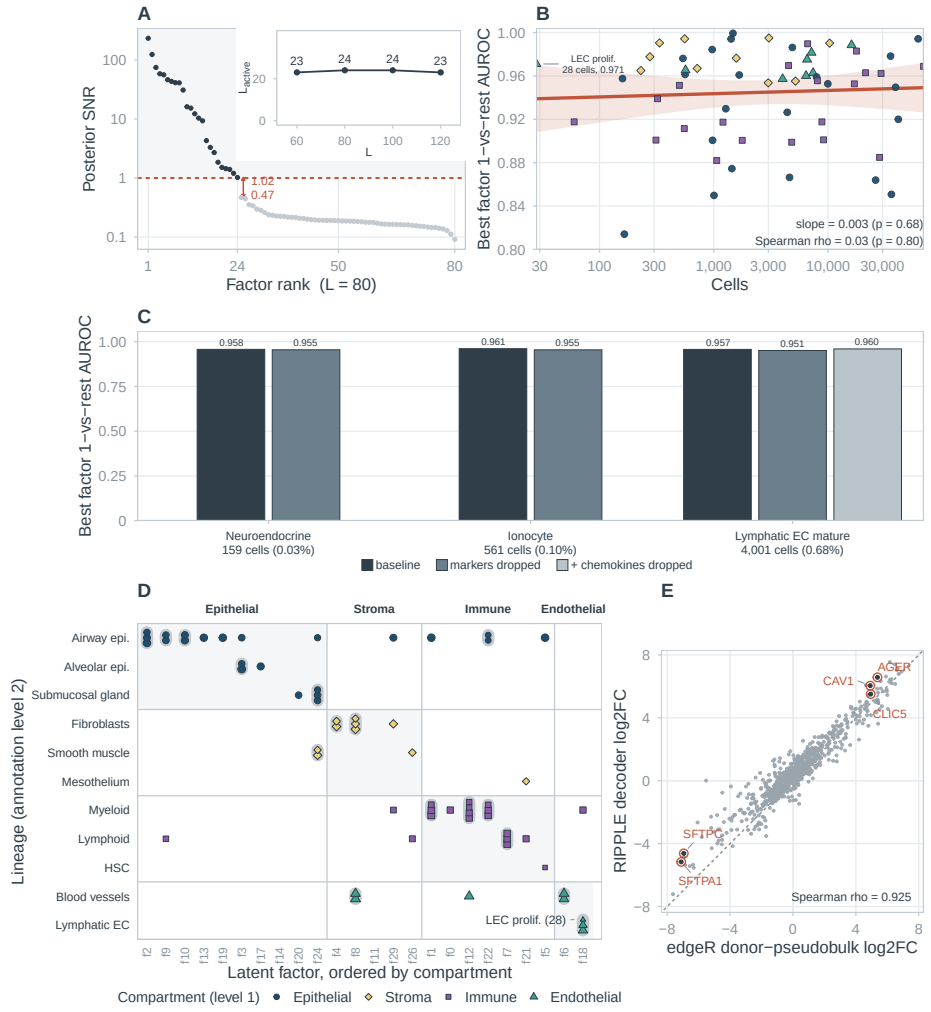


Figure 2: Reading HLCA with 24 factors. (A) The SNR spectrum with the inset showing L_{active} against the ceiling L. (B) Best factor 1-vs-rest AUROC versus the number of cells in each cell type, with a linear fit in red line and its 95% confidence band in light red. Each dot is a cell type. Color and shape indicate the cell type's compartment. (C) Best AUROC before and after a cell type's canonical markers are deleted from the gene panel and the model refit. (D) Factor arrangement. One dot per cell type, placed at the factor it is most active on. The y-axis denotes a cell type's lineage. Color and shape indicate the compartment. (E) Decoder log-fold-change against donor-pseudobulk edgeR for AT1 vs. AT2. One point per gene and AT1/AT2 markers in dark points with a red circle.

epithelial cells and are the origin of small-cell lung cancer, the most aggressive lung cancer (Branchfield et al., 2016; Sutherland et al., 2011; Rudin et al., 2021). On the other hand, LEC mature is chosen because lymphatic vessels are wrapped in immune cells secreting chemokines (Randolph et al., 2017), allowing us to test whether RIPPLE is robust to the removal of both the canonical marker genes and the chemokines. As Figure 2C shows, AUROC barely moves when the canonical markers are removed for all three cell types tested, with AUROC dropping from 0.958 to 0.955 for PNECs, from 0.961 to 0.955 for ionocytes, and from 0.957 to 0.951 for LEC mature. With the chemokines also removed, AUROC for LEC mature remains high at 0.960. These results show that RIPPLE is robust to marker removal (the strongest discriminative signals), and can still separate a cell type from the rest when remnant signals are present in the dataset.

Notably, RIPPLE represents 61 cell types with only $L_{\text{active}} = 24$ factors, so the factors cannot be per-cell-type indicators. Specifically, we read the columns of the score matrix W, which encodes

the activity of each cell on each factor. Figure 2D places each cell type on the column it is most active on. Of the 24 columns, 17 are the primary axis of more than one cell type, and 5 are the primary axis of only one cell type. The remaining two are not the best discriminator of a cell type but they remain top-ranked for some cell types (see Appendix C). That sharing is not arbitrary. As evidenced by the moderately diagonal block structure in Figure 2D, RIPPLE's factor placement mostly follows HLCA annotation's own hierarchy. In particular, 10 factors encode cell types within the same lineage and 11 within the same compartment, the five single-cell-type columns among them trivially; the latter include f_0, which concentrates the depth association while still resolving cell types (Appendix C). Of the six columns holding more than one cell type inside a compartment, all place those cell types at the same pole, indicating a coherent shared program (e.g., factor 7, the NK/CD4/CD8 lymphoid axis). When a factor places its cell types at opposite poles, all six such factors span two compartments, encoding a contrast (e.g., factor 8 sets capillary endothelium against fibroblasts). The remaining cross-compartment factors share a single pole, and we defer their interpretations to Appendix C. Nothing in Eq. (2) ties a factor to a compartment: the prior is exchangeable over factors and the likelihood never sees the annotation. The arrangement is therefore a property of what RIPPLE found rather than of what it was given.

RIPPLE's factors also support a gene-level readout. Since $\log \mu$ is linear in W with the loading matrix α frozen, the decoder turns a contrast between two groups of cells into a per-gene log-fold-change, allowing for differential expression (DE) analysis. Since there is no ground-truth value, we compare RIPPLE against donor-pseudobulk edgeR, the standard DE analysis for atlas data, which sums counts within each donor and cell type into bulk-like profiles and tests them with an NB generalized linear model (Robinson et al., 2010). We analyze four contrasts, each pairing a cell type with its nearest relative inside a compartment: AT1 versus AT2, alveolar versus adventitial fibroblasts, alveolar against monocyte-derived macrophages, and the two capillary endothelial cell types of the alveolus. Figure 2E shows the AT1-versus-AT2 comparison, and the rank concordance with edgeR reaches Spearman $\rho = 0.925$. On the 50 genes edgeR calls most strongly differential, the two methods agree in sign on 98% of them. Over the four contrasts analyzed, Spearman ρ ranges from 0.804 to 0.925, and the top-50 sign agreement is at least 98% in every case. This alignment is structural: RIPPLE's factors span the highest-variance directions of the data, and the strongest DE lies along those directions, so the decoder is sharpest exactly where the DE signal is strongest. Beyond recovering these DE genes, the decoder attributes each change to the gene program its factors identify, which a traditional DE test cannot. However, the decoder DE analysis is not a direct replacement for a traditional DE test since mean-field VB can underestimate the posterior variance, and a calibrated significance test is non-trivial. We therefore view RIPPLE's decoder DE as a complementary tool to the traditional DE test, where traditional tests flag what is differentially expressed, and RIPPLE names the gene program that drives the difference.

5 Discussion

RIPPLE settles two choices that are usually left unprincipled in single-cell embeddings: the latent dimension and the meaning of the coordinates. A half-Cauchy prior on the per-factor loading scales lets the data choose the dimension, and the linear decoder keeps every selected factor readable as a gene program. Under mild conditions, the selection rule loses no true factor. On average across six datasets, RIPPLE outperforms every unsupervised baseline. On HLCA, a compact set of factors separates every cell type regardless of its rarity. These results suggest that interpretability of single-cell embeddings need not be traded for accuracy.

We also identify three interesting future directions. (i) *Testing differential expression*: since mean-field VB can underestimate the posterior variance, a DE test needs recalibrated uncertainty, e.g., Trustworthy Variational Bayes (Liu & Li, 2025). The transfer is not immediate, since its guarantees assume a fixed number of parameters and its calibration refits the whole model. (ii) *Mapping new cells*: RIPPLE has no encoder, so a cell outside the fit has no scores. However, the linear decoder makes a new cell's log-likelihood concave in its scores once the gene-level parameters are fixed at their fitted values. Placing new cells on a fitted reference would then be a small convex problem with no re-training. (iii) *Certifying individual factors*: our theory certifies the selected span and a distinct counterpart of every true factor, while the reading of individual factors rests on empirical audits. Conditions under which the prior's basis provably aligns with the generative gene programs, such as sparse loadings (Rohe & Zeng, 2023), remain open.

AI USE STATEMENT

In this work, we used generative AI tools for developing the SNR selection rule of the RIPPLE model, refining mathematical claims, providing some of the critical ingredients for proving mathematical claims, assisting in the writing of proofs, providing feedback on research methodology and experimental design, refining some of the hypotheses, implementing methods, cleaning datasets, and interpreting some of the empirical results. We have not used generative AI tools to generate synthetic data sets, assist with translation, or support qualitative and thematic data analysis. Additionally, we used generative AI tools to refine some modeling decisions, draft some sentences of the paper, summarize existing literature, improve the paper's readability, and create or revise code and figures. We have reviewed all AI-assisted work. For example, AI-generated code and proofs in Appendix E were checked by multiple authors. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

Every result in this paper regenerates from the repository at `https://anonymous.4open.science/r/ripple-61C3`. It includes the RIPPLE package and necessary scripts to reproduce results shown in Sections 4.1–4.2 and Appendices B–D. The default RIPPLE hyperparameters are stated in Appendix A.8, and Appendix A states the generative model, the ELBO, the optimizer, the graph construction and the SNR selection rule in full. Appendix B.1 describes each benchmark dataset and its processing, all from public sources that the repository downloads. Appendix B.2 gives every baseline's configuration and the pinned software versions, and Appendix B.3 details the evaluation protocol and the ten seeds behind every reported mean. The HLCA panel is built from the public HLCA core as Appendix C.1 describes. The assumptions and complete proofs of Section 3 are in Appendix E.

REFERENCES

Jushan Bai. Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279, 2009. doi: 10.3982/ECTA6135.

Z. D. Bai and Y. Q. Yin. Limit of the smallest eigenvalue of a large dimensional sample covariance matrix. *The Annals of Probability*, 21(3), 1993. doi: 10.1214/aop/1176989118.

Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005. doi: 10.1214/009117905000000233.

Afonso S. Bandeira and Ramon van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *The Annals of Probability*, 44(4), 2016. doi: 10.1214/15-aop1025.

Maayan Baron, Adrian Veres, Samuel L. Wolock, Aubrey L. Faust, Renaud Gaujoux, Amedeo Vetere, Jennifer Hyoje Ryu, Bridget K. Wagner, Shai S. Shen-Orr, Allon M. Klein, Douglas A. Melton, and Itai Yanai. A single-cell transcriptomic map of the human and mouse pancreas reveals inter- and intra-cell population structure. *Cell Systems*, 3(4):346–360.e4, 2016. doi: 10.1016/j.cels.2016.08.011.

Florent Benaych-Georges and Raj Rao Nadakuditi. The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis*, 111:120–135, 2012. doi: 10.1016/j.jmva.2012.04.019.

Christopher M. Bishop. Bayesian PCA. In Michael J. Kearns, Sara A. Solla, and David A. Cohn (eds.), *Advances in Neural Information Processing Systems 11, [NIPS Conference, Denver, Colorado, USA, November 30 - December 5, 1998]*, pp. 382–388. The MIT Press, 1998.

David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017. doi: 10.1080/01621459.2017.1285773.

Kelsey Branchfield, Leah Nantie, Jamie M. Verheyden, Pengfei Sui, Mark D. Wienhold, and Xin Sun. Pulmonary neuroendocrine cells function as airway sensors to control lung immune response. *Science*, 351(6274):707–710, 2016. doi: 10.1126/science.aad7969.

C. M. Carvalho, N. G. Polson, and J. G. Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 2010. doi: 10.1093/biomet/asq017.

Y Chen and X Li. Determining the number of factors in high-dimensional generalized latent factor models. *Biometrika*, 109(3):769–782, 2022. doi: 10.1093/biomet/asab044.

Yunxiao Chen, Xiaoou Li, and Siliang Zhang. Joint maximum likelihood estimation for high-dimensional exploratory item factor analysis. *Psychometrika*, 84(1):124–146, 2019. doi: 10.1007/s11336-018-9646-5.

Chris Cremer, Xuechen Li, and David Duvenaud. Inference suboptimality in variational autoencoders. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1086–1094. PMLR, 2018.

CZI Cell Science Program, Shibla Abdulla, Brian Aevermann, Pedro Assis, Seve Badajoz, Sidney M. Bell, Emanuele Bezzi, Batuhan Cakir, Jim Chaffer, Signe Chambers, J. Michael Cherry, Tiffany Chi, Jennifer Chien, Leah Dorman, Pablo Garcia-Nieto, Nayib Gloria, Mim Hastie, Daniel Hegeman, Jason Hilton, Timmy Huang, Amanda Infeld, Ana-Maria Istrate, Ivana Jelic, Kuni Katsuya, Yang Joon Kim, Karen Liang, Mike Lin, Maximilian Lombardo, Bailey Marshall, Bruce Martin, Fran McDade, Colin Megill, Nikhil Patel, Alexander Predeus, Brian Raymor, Behnam Robatmili, Dave Rogers, Erica Rutherford, Dana Sadgat, Andrew Shin, Corinn Small, Trent Smith, Prathap Sridharan, Alexander Tarashansky, Norbert Tavares, Harley Thomas, Andrew Tolopko, Meghan Urisko, Joyce Yan, Garabet Yeretssian, Jennifer Zamanian, Arathi Mani, Jonah Cool, and Ambrose Carr. CZ CELLxGENE Discover: a single-cell data platform for scalable exploration, analysis and modeling of aggregated data. *Nucleic Acids Research*, 53(D1):D886–D900, 2025. doi: 10.1093/nar/gkae1142.

Angelo Duò, Mark D. Robinson, and Charlotte Soneson. A systematic performance evaluation of clustering methods for single-cell rna-seq data. *F1000Research*, 7:1141, 2018. doi: 10.12688/f1000research.15666.1.

Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. doi: 10.1007/BF02288367.

Ahmed El Alaoui and Michael I. Jordan. Detection limits in the high-dimensional spiked rectangular model. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 410–438. PMLR, 2018.

J. Stuart Elborn. Cystic fibrosis. *The Lancet*, 388(10059):2519–2531, 2016. doi: 10.1016/S0140-6736(16)00576-6.

Zhaoyu Fang, Ruiqing Zheng, and Min Li. scmae: a masked autoencoder for single-cell rna-seq clustering. *Bioinformatics*, 40(1), 2024. doi: 10.1093/bioinformatics/btae020.

C. Gourieroux, A. Monfort, and A. Trognon. Pseudo maximum likelihood methods: Theory. *Econometrica*, 52(3):681–700, 1984. doi: 10.2307/1913471.

Dominic Grün, Lennart Kester, and Alexander van Oudenaarden. Validation of noise models for single-cell transcriptomics. *Nature Methods*, 11(6):637–640, 2014. doi: 10.1038/nmeth.2930.

Laleh Haghverdi, Aaron T L Lun, Michael D Morgan, and John C Marioni. Batch effects in single-cell rna-sequencing data are corrected by matching mutual nearest neighbors. *Nature Biotechnology*, 36(5):421–427, 2018. doi: 10.1038/nbt.4091.

P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(1):26–30, 1935. doi: 10.1112/jlms/s1-10.37.26.

Graham Heimberg, Rajat Bhatnagar, Hana El-Samad, and Matt Thomson. Low dimensionality in gene expression data enables the accurate extraction of transcriptional programs from shallow sequencing. *Cell Systems*, 2(4):239–250, 2016. doi: 10.1016/j.cels.2016.04.001.

Stephanie C. Hicks, F. William Townes, Mingxiang Teng, and Rafael A. Irizarry. Missing data and technical variability in single-cell rna-sequencing experiments. *Biostatistics*, 19(4):562–578, 2018. doi: 10.1093/biostatistics/kxx053.

Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2012. ISBN 9780521548236. doi: 10.1017/cbo9781139020411.

Peter V. Kharchenko, Lev Silberstein, and David T. Scadden. Bayesian approach to single-cell differential expression analysis. *Nature Methods*, 11(7):740–742, 2014. doi: 10.1038/nmeth.2967.

Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.

Vladimir Yu. Kiselev, Tallulah S. Andrews, and Martin Hemberg. Challenges in unsupervised clustering of single-cell rna-seq data. *Nature Reviews Genetics*, 20(5):273–282, 2019. doi: 10.1038/s41576-018-0088-9.

Ilya Korsunsky, Nghia Millard, Jean Fan, Kamil Slowikowski, Fan Zhang, Kevin Wei, Yuriy Baglaenko, Michael Brenner, Po-ru Loh, and Soumya Raychaudhuri. Fast, sensitive and accurate integration of single-cell data with harmony. *Nature Methods*, 16(12):1289–1296, 2019. doi: 10.1038/s41592-019-0619-0.

Peter Langfelder and Steve Horvath. WGCNA: an R package for weighted correlation network analysis. *BMC Bioinformatics*, 9(1):559, 2008. doi: 10.1186/1471-2105-9-559.

Jan Lause, Philipp Berens, and Dmitry Kobak. Analytic Pearson residuals for normalization of single-cell RNA-seq UMI data. *Genome Biology*, 22(1):258, 2021. doi: 10.1186/s13059-021-02451-7.

Jiaming Liu and Meng Li. Bend to mend: Toward trustworthy variational Bayes with valid uncertainty quantification, 2025.

Romain Lopez, Jeffrey Regier, Michael B. Cole, Michael I. Jordan, and Nir Yosef. Deep generative modeling for single-cell transcriptomics. *Nature Methods*, 15(12):1053–1058, 2018. doi: 10.1038/s41592-018-0229-2.

László Lovász and Michael Plummer. *Matching Theory*. American Mathematical Society, 2009. ISBN 9781470415754. doi: 10.1090/chel/367.

Malte D. Luecken and Fabian J. Theis. Current best practices in single-cell rna-seq analysis: a tutorial. *Molecular Systems Biology*, 15(6):e8746, 2019. doi: 10.15252/msb.20188746.

Malte D. Luecken, M. Büttner, K. Chaichoompu, A. Danese, M. Interlandi, M. F. Mueller, D. C. Strobl, L. Zappia, M. Dugas, M. Colomé-Tatché, and Fabian J. Theis. Benchmarking atlas-level data integration in single-cell genomics. *Nature Methods*, 19(1):41–50, 2022. doi: 10.1038/s41592-021-01336-8.

Evan Z. Macosko, Anindita Basu, Rahul Satija, James Nemesh, Karthik Shekhar, Melissa Goldman, Itay Tirosh, Allison R. Bialas, Nolan Kamitaki, Emily M. Martersteck, John J. Trombetta, David A. Weitz, Joshua R. Sanes, Alex K. Shalek, Aviv Regev, and Steven A. McCarroll. Highly parallel genome-wide expression profiling of individual cells using nanoliter droplets. *Cell*, 161(5):1202–1214, 2015. doi: 10.1016/j.cell.2015.05.002.

Thomas P. Minka. Automatic choice of dimensionality for PCA. In Todd K. Leen, Thomas G. Dietterich, and Volker Tresp (eds.), *Advances in Neural Information Processing Systems 13, Papers from Neural Information Processing Systems (NIPS) 2000, Denver, CO, USA*, pp. 598–604. MIT Press, 2000.

Daniel T. Montoro, Adam L. Haber, Moshe Biton, Vladimir Vinarsky, Brian Lin, Susan E. Birket, et al. A revised airway epithelial hierarchy includes CFTR-expressing ionocytes. *Nature*, 560 (7718):319–324, 2018. doi: 10.1038/s41586-018-0393-7.

Hyungsik Roger Moon and Martin Weidner. Dynamic linear panel regression models with interactive fixed effects. *Econometric Theory*, 33(1):158–195, 2017. doi: 10.1017/S0266466615000328.

Mauro J. Muraro, Gitanjali Dharmadhikari, Dominic Grün, Nathalie Groen, Tim Dielen, Erik Jansen, Leon van Gurp, Marten A. Engelse, Francoise Carlotti, Eelco J.P. de Koning, and Alexander van Oudenaarden. A single-cell transcriptome atlas of the human pancreas. *Cell Systems*, 3 (4):385–394.e3, 2016. doi: 10.1016/j.cels.2016.09.002.

J. Neyman and Elizabeth L. Scott. Consistent estimates based on partially consistent observations. *Econometrica*, 16(1):1–32, 1948. doi: 10.2307/1914288.

Phillip B Nicol and Jeffrey W Miller. Model-based dimensionality reduction for single-cell rna-seq using generalized bilinear models. *Biostatistics*, 26(1), 2025. doi: 10.1093/biostatistics/kxaf024.

Alexei Onatski, Marcelo J. Moreira, and Marc Hallin. Asymptotic power of sphericity tests for high-dimensional data. *The Annals of Statistics*, 41(3), 2013. doi: 10.1214/13-aos1100.

Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617–1642, 2007.

Amelia Perry, Alexander S. Wein, Afonso S. Bandeira, and Ankur Moitra. Optimality and sub-optimality of pca i: Spiked random matrix models. *The Annals of Statistics*, 46(5), 2018. doi: 10.1214/17-aos1625.

Juho Piironen and Aki Vehtari. Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics*, 11(2), 2017. doi: 10.1214/17-EJS1337SI.

Lindsey W. Plasschaert, Rapolas Žilionis, Rayman Choo-Wing, Virginia Savova, Judith Knehr, Guglielmo Roma, Allon M. Klein, and Aron B. Jaffe. A single-cell atlas of the airway epithelium reveals the CFTR-rich pulmonary ionocyte. *Nature*, 560(7718):377–381, 2018. doi: 10.1038/s41586-018-0394-6.

Krzysztof Polański, Matthew D Young, Zhichao Miao, Kerstin B Meyer, Sarah A Teichmann, and Jong-Eun Park. Bbknn: fast batch alignment of single cell transcriptomes. *Bioinformatics*, 36(3): 964–965, 2020. doi: 10.1093/bioinformatics/btz625.

R. Price. A useful theorem for nonlinear devices having Gaussian inputs. *IRE Transactions on Information Theory*, 4(2):69–72, 1958. doi: 10.1109/TIT.1958.1057444.

Gwendalyn J. Randolph, Stoyan Ivanov, Bernd H. Zinselmeyer, and Joshua P. Scallan. The lymphatic system: Integral roles in immunity. *Annual Review of Immunology*, 35(1):31–52, 2017. doi: 10.1146/annurev-immunol-041015-055354.

Aviv Regev, Sarah A. Teichmann, Eric S. Lander, Ido Amit, Christophe Benoist, Ewan Birney, Bernd Bodenmiller, Peter Campbell, Piero Carninci, Menna Clatworthy, Hans Clevers, Bart Deplancke, Ian Dunham, James Eberwine, Roland Eils, Wolfgang Enard, Andrew Farmer, Lars Fugger, Berthold Göttgens, Nir Hacohen, Muzlifah Haniffa, Martin Hemberg, Seung Kim, Paul Klenerman, Arnold Kriegstein, Ed Lein, Sten Linnarsson, Emma Lundberg, Joakim Lundeberg, Partha Majumder, John C. Marioni, Miriam Merad, Musa Mhlanga, Martijn Nawijn, Mihai Netea, Garry Nolan, Dana Pe'er, Anthony Phillipakis, Chris P. Ponting, Stephen Quake, Wolf Reik, Orit Rozenblatt-Rosen, Joshua Sanes, Rahul Satija, Ton N. Schumacher, Alex Shalek, Ehud Shapiro, Padmanee Sharma, Jay W. Shin, Oliver Stegle, Michael Stratton, Michael J. T. Stubbington, Fabian J. Theis, Matthias Uhlen, Alexander van Oudenaarden, Allon Wagner, Fiona Watt, Jonathan Weissman, Barbara Wold, Ramnik Xavier, and Nir Yosef. The Human Cell Atlas. *eLife*, 6:e27041, 2017. doi: 10.7554/eLife.27041.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In Eric P. Xing and Tony Jebara (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 1278–1286. PMLR, 2014.

Davide Risso, Fanny Perraudeau, Svetlana Gribkova, Sandrine Dudoit, and Jean-Philippe Vert. A general and flexible method for signal extraction from single-cell rna-seq data. *Nature Communications*, 9(1), 2018. doi: 10.1038/s41467-017-02554-5.

Mark D. Robinson, Davis J. McCarthy, and Gordon K. Smyth. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26(1):139–140, 2010. doi: 10.1093/bioinformatics/btp616.

Karl Rohe and Muzhe Zeng. Vintage factor analysis with varimax performs statistical inference. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(4):1037–1060, 2023. doi: 10.1093/jrsssb/qkad029.

Charles M. Rudin, Elisabeth Brambilla, Corinne Faivre-Finn, and Julien Sage. Small-cell lung cancer. *Nature Reviews Disease Primers*, 7(1):3, 2021. doi: 10.1038/s41572-020-00235-0.

Åsa Segerstolpe, Athanasia Palasantza, Pernilla Eliasson, Eva-Marie Andersson, Anne-Christine Andréasson, Xiaoyan Sun, Simone Picelli, Alan Sabirsh, Maryam Clausen, Magnus K. Bjursell, David M. Smith, Maria Kasper, Carina Ämmälä, and Rickard Sandberg. Single-cell transcriptome profiling of human pancreatic islets in health and type 2 diabetes. *Cell Metabolism*, 24(4): 593–607, 2016. doi: 10.1016/j.cmet.2016.08.020.

Lisa Sikkema, Ciro Ramírez-Suástegui, Daniel C. Strobl, Tessa E. Gillett, Luke Zappia, Elo Madissoon, Nikolay S. Markov, Laure-Emmanuelle Zaragosi, Yuge Ji, Meshal Ansari, Marie-Jeanne Arguel, Leonie Apperloo, Martin Banchero, Christophe Bécavin, Marijn Berg, Evgeny Chichelnitskiy, Mei-i Chung, Antoine Collin, Aurore C. A. Gay, Janine Gote-Schniering, Baharak Hooshiar Kashani, Kemal Inecik, Manu Jain, Theodore S. Kapellos, Tessa M. Kole, Sylvie Leroy, Christoph H. Mayr, Amanda J. Oliver, Michael von Papen, Lance Peter, Chase J. Taylor, Thomas Walzthoeni, Chuan Xu, Linh T. Bui, Carlo De Donno, Leander Dony, Alen Faiz, Minzhe Guo, Austin J. Gutierrez, Lukas Heumos, Ni Huang, Ignacio L. Ibarra, Nathan D. Jackson, Preetish Kadur Lakshminarasimha Murthy, Mohammad Lotfollahi, Tracy Tabib, Carlos Talavera-López, Kyle J. Travaglini, Anna Wilbrey-Clark, Kaylee B. Worlock, Masahiro Yoshida, Yuexin Chen, James S. Hagood, Ahmed Agami, Peter Horvath, Joakim Lundeberg, Charles-Hugo Marquette, Gloria Pryhuber, Christos Samakovlis, Xin Sun, Lorraine B. Ware, Kun Zhang, Maarten van den Berge, Yohan Bossé, Tushar J. Desai, Oliver Eickelberg, Naftali Kaminski, Mark A. Krasnow, Robert Lafyatis, Marko Z. Nikolic, Joseph E. Powell, Jayaraj Rajagopal, Mauricio Rojas, Orit Rozenblatt-Rosen, Max A. Seibold, Dean Sheppard, Douglas P. Shepherd, Don D. Sin, Wim Timens, Alexander M. Tsankov, Jeffrey Whitsett, Yan Xu, Nicholas E. Banovich, Pascal Barbry, Thu Elizabeth Duong, Christine S. Falk, Kerstin B. Meyer, Jonathan A. Kropski, Dana Pe'er, Herbert B. Schiller, Purushothama Rao Tata, Joachim L. Schultze, Sarah A. Teichmann, Alexander V. Misharin, Martijn C. Nawijn, Malte D. Luecken, and Fabian J. Theis. An integrated cell atlas of the lung in health and disease. *Nature Medicine*, 29(6):1563–1577, 2023. doi: 10.1038/s41591-023-02327-2.

Joshua M. Stuart, Eran Segal, Daphne Koller, and Stuart K. Kim. A gene-coexpression network for global discovery of conserved genetic modules. *Science*, 302(5643):249–255, 2003. doi: 10.1126/science.1087447.

Kate D. Sutherland, Natalie Proost, Inge Brouns, Dirk Adriaensen, Ji-Ying Song, and Anton Berns. Cell of origin of small cell lung cancer: Inactivation of Trp53 and Rb1 in distinct cell types of adult mouse lung. *Cancer Cell*, 19(6):754–764, 2011. doi: 10.1016/j.ccr.2011.04.019.

Valentine Svensson. Droplet scrna-seq is not zero-inflated. *Nature Biotechnology*, 38(2):147–150, 2020. doi: 10.1038/s41587-019-0379-5.

Fuchou Tang, Catalin Barbacioru, Yangzhou Wang, Ellen Nordman, Clarence Lee, Nanlan Xu, Xiaohui Wang, John Bodeau, Brian B. Tuch, Asim Siddiqui, Kaiqin Lao, and M. Azim Surani. mrna-seq whole-transcriptome analysis of a single cell. *Nature Methods*, 6(5):377–382, 2009. doi: 10.1038/nmeth.1315.

F. William Townes, Stephanie C. Hicks, Martin J. Aryee, and Rafael A. Irizarry. Feature selection and dimension reduction for single-cell rna-seq based on a multinomial model. *Genome Biology*, 20(1):295, 2019. doi: 10.1186/s13059-019-1861-6.

Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2018. ISBN 9781108415194. doi: 10.1017/9781108231596.

Yixin Wang and David M. Blei. Frequentist consistency of variational bayes. *Journal of the American Statistical Association*, 114(527):1147–1161, 2019. doi: 10.1080/01621459.2018.1473776.

R. W. M. Wedderburn. Quasi-likelihood functions, generalized linear models, and the Gauss–Newton method. *Biometrika*, 61(3):439–447, 1974. doi: 10.1093/biomet/61.3.439.

Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972. doi: 10.1007/BF01932678.

F. Alexander Wolf, Philipp Angerer, and Fabian J. Theis. SCANPY: large-scale single-cell gene expression data analysis. *Genome Biology*, 19(1):15, 2018. doi: 10.1186/s13059-017-1382-0.

Jiankang Xiong, Fuzhou Gong, Liang Ma, and Lin Wan. scvic: deep generative modeling of heterogeneity for scrna-seq data. *Bioinformatics Advances*, 4(1), 2024. doi: 10.1093/bioadv/vbae086.

Chenling Xu, Romain Lopez, Edouard Mehlman, Jeffrey Regier, Michael I Jordan, and Nir Yosef. Probabilistic harmonization and annotation of single-cell transcriptomics data with deep generative models. *Molecular Systems Biology*, 17(1), 2021. doi: 10.15252/msb.20209620.

Amit Zeisel, Ana B. Muñoz-Manchado, Simone Codeluppi, Peter Lönnerberg, Gioele La Manno, Anna Juréus, Sueli Marques, Hermany Munguba, Liqun He, Christer Betsholtz, Charlotte Rolny, Gonçalo Castelo-Branco, Jens Hjerling-Leffler, and Sten Linnarsson. Cell types in the mouse cortex and hippocampus revealed by single-cell rna-seq. *Science*, 347(6226):1138–1142, 2015. doi: 10.1126/science.aaa1934.

Fengshuo Zhang and Chao Gao. Convergence rates of variational posterior distributions. *The Annals of Statistics*, 48(4), 2020. doi: 10.1214/19-aos1883.

Grace X. Y. Zheng, Jessica M. Terry, Phillip Belgrader, Paul Ryvkin, Zachary W. Bent, Ryan Wilson, Solongo B. Ziraldo, Tobias D. Wheeler, Geoff P. McDermott, Junjie Zhu, Mark T. Gregory, Joe Shuga, Luz Montesclaros, Jason G. Underwood, Donald A. Masquelier, Stefanie Y. Nishimura, Michael Schnall-Levin, Paul W. Wyatt, Christopher M. Hindson, Rajiv Bharadwaj, Alexander Wong, Kevin D. Ness, Lan W. Beppu, H. Joachim Deeg, Christopher McFarland, Keith R. Loeb, William J. Valente, Nolan G. Ericson, Emily A. Stevens, Jerald P. Radich, Tarjei S. Mikkelsen, Benjamin J. Hindson, and Jason H. Bielas. Massively parallel digital transcriptional profiling of single cells. *Nature Communications*, 8(1):14049, 2017. doi: 10.1038/ncomms14049.

A Model and inference details

A.1 Full generative model

Written top-down, with indices $i \in [n]$, $j \in [p]$, $\ell \in [L]$, $h \in [H]$, $g \in [G]$, and the optional batch subspace indexed by $q \in [Q]$ over batches $b \in [B]$:

$$
\begin{aligned}
\lambda_\ell &\sim \mathrm{HalfCauchy}(1), \\
\alpha_{\ell j} \mid \lambda_\ell &\sim \mathcal{N}\big(0, \lambda_\ell^2\big), \\
W_{i\ell} &\sim \mathcal{N}(0, 1), \\
\big[\, \lambda_q^b \,\big] &\sim \mathrm{HalfCauchy}(1), \\
\big[\, \delta_{qj} \mid \lambda_q^b \,\big] &\sim \mathcal{N}\big(0, (\lambda_q^b)^2\big), \\
\big[\, Z_{bq} \,\big] &\sim \mathcal{N}(0, 1), \\
\beta_{hj} &\sim \mathcal{N}\big(0, \sigma_\beta^2\big), \\
\xi_j &\sim \mathcal{N}\big(\mu_\xi, \sigma_\xi^2\big), \qquad \theta_j = e^{\xi_j}, \\
p(\gamma_{ig}) &\propto 1, \\
\log \mu_{ij} &= X_{i\cdot}^\top \beta_{\cdot j} + \gamma_{i\cdot}^\top V_{j\cdot} + W_{i\cdot}^\top \alpha_{\cdot j} + \big[\, Z_{b(i)\cdot}^\top \delta_{\cdot j} \,\big] + O_i, \\
Y_{ij} &\sim \mathrm{NB}(\mu_{ij}, \theta_j),
\end{aligned}
\tag{7}
$$

where $[\,\cdot\,]$ marks the opt-in batch subspace of Eq. (5), off by default. The prior hyperparameters are fixed rather than learned: $\sigma_\beta = 5$ and $\sigma_\xi = 1$. The log-dispersion prior location μ_ξ is empirical-Bayes by default, set at initialization to the method-of-moments mean $\overline{\log \theta^{\mathrm{MoM}}}$ with $\theta_j^{\mathrm{MoM}} = \bar{Y}_j^2/(\mathrm{Var}(Y_{\cdot j}) - \bar{Y}_j)$.

In the mean–dispersion parametrization, $\mathbb{E}[Y_{ij}] = \mu_{ij}$ and $\mathrm{Var}(Y_{ij}) = \mu_{ij} + \mu_{ij}^2/\theta_j$, with log-density:

$$\log \mathrm{NB}(y \mid \mu, \theta) = \log\Gamma(y+\theta) - \log\Gamma(\theta) - \log\Gamma(y+1) + y\log\mu + \theta\log\theta - (y+\theta)\log(\mu+\theta). \quad (8)$$

A.2 Covariate designs used in our experiments

Note that both designs X and V are user-supplied (Section 2). Let $b(i) \in [B]$ denote the batch of cell i. Throughout our experiments, we use

$$\begin{aligned} X_{i\cdot} &= \big[\,1,\ \mathbb{1}\{b(i)=2\},\ \ldots,\ \mathbb{1}\{b(i)=B\},\ r_i - \bar{r}\,\big] \in \mathbb{R}^H, \\ V_{j\cdot} &= \big[\,1,\ z_j - \bar{z},\ \log(1+\bar{y}_j) - \overline{\log(1+\bar{y}_j)}\,\big] \in \mathbb{R}^G, \end{aligned} \quad (9)$$

where $r_i = p^{-1}\sum_j \mathbb{1}\{Y_{ij} > 0\}$ is the detection rate of cell i, $z_j = n^{-1}\sum_i \mathbb{1}\{Y_{ij} = 0\}$ the fraction of cells in which gene j is zero, and $\bar{y}_j = n^{-1}\sum_i Y_{ij}$ its mean count. The first batch level is absorbed into the intercept, so $H = B + 1$ when batch is included as a covariate, and $G = 3$. Both designs are centered, not scaled. The offset is fixed rather than fitted, at $O_i = \log(d_i/\bar{d})$ for library size $d_i = \sum_j Y_{ij}$. Dropping batch from $X_{i\cdot}$ is supported and removes the fixed-effect batch path entirely. We treat a representation as batch-corrected only when both that path and the random effect $Z\delta$ are removed, since removing one alone leaves a mean that is neither fitted nor corrected.

A.3 The variational family

Every parameter takes an independent Gaussian variational distribution with free mean $m_\bullet$ and variance $s_\bullet^2$ of the same shape. The loading scales λ and the gene-covariate loadings γ are optimized point estimates and receive no variational variance.

parameter	posterior	shape	role
$W_{i\cdot}$	$\mathcal{N}(m_{W_{i\cdot}}, \mathrm{diag}(s^2_{W_{i\cdot}}))$	L per cell	embedding
α	$\mathcal{N}(m_\alpha, \mathrm{diag}(s^2_\alpha))$	$L \times p$	loadings
β	$\mathcal{N}(m_\beta, \mathrm{diag}(s^2_\beta))$	$H \times p$	cell-covariate coefficients
ξ	$\mathcal{N}(m_\xi, \mathrm{diag}(s^2_\xi))$	p	log-dispersion
Z	$\mathcal{N}(m_Z, \mathrm{diag}(s^2_Z))$	$B \times Q$	batch scores (opt-in)
δ	$\mathcal{N}(m_\delta, \mathrm{diag}(s^2_\delta))$	$Q \times p$	batch loadings (opt-in)
λ	point estimate	L	loading scales
$\gamma_{i\cdot}$	point estimate	G per cell	gene-covariate loadings

A.4 The ELBO in full

The bound is optimized on minibatches of cells, with every per-cell term rescaled to the full data. Let $\mathcal{B}$ be a minibatch of cell indices, T the number of Monte-Carlo samples, and $w_{\mathrm{KL}} \in [0,1]$ a KL-annealing weight. The estimator optimized in practice is

$$\widehat{\mathcal{L}} = \frac{n}{|\mathcal{B}|T}\sum_{t=1}^{T}\sum_{i\in\mathcal{B}}\sum_{j=1}^{p}\log p\big(Y_{ij} \mid W_{i\cdot}^{(t)}, \Psi^{(t)}\big) - w_{\mathrm{KL}}\Big[\tfrac{n}{|\mathcal{B}|}\textstyle\sum_{i\in\mathcal{B}}\mathrm{KL}_{W_{i\cdot}} + \mathrm{KL}_{\mathrm{glob}} - \log p(\lambda)\Big] - \mathcal{R}(\mathcal{B}), \quad (10)$$

where $\Psi^{(t)}$ collects the reparameterized global draws and $\mathrm{KL}_{\mathrm{glob}} = \mathrm{KL}(q(\beta)) + \mathrm{KL}(q(\xi)) + \mathrm{KL}_\alpha(q(\alpha))$, together with the batch-subspace terms when it is enabled. The likelihood and the per-cell embedding KL carry the factor $n/|\mathcal{B}|$, so the minibatch is an unbiased estimate of the full-data bound. The global KLs are not rescaled.

The graph penalty carries the same upscaling as the per-cell KL,

$$\mathcal{R}(\mathcal{B}) = \lambda_g \frac{n}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \left\| m_{W_{i\cdot}} - \mathrm{sg}\left[\overline{m}_{\mathcal{G}(i)} \right] \right\|_2^2,$$

so its weight relative to the likelihood does not depend on the minibatch size. It sits outside the probabilistic bound and is *not* annealed.

All KL terms are analytic. For the embeddings,

$$\mathrm{KL}_{W_{i\cdot}} = \frac{1}{2} \sum_{\ell} \left(s_{W_{i\ell}}^2 + m_{W_{i\ell}}^2 - 1 - \log s_{W_{i\ell}}^2 \right).$$

For a factorized Gaussian $q = \mathcal{N}(m, \mathrm{diag}(s^2))$ against an isotropic prior $\mathcal{N}(0, \sigma_0^2 I)$, $\mathrm{KL} = \frac{1}{2} \sum \left((s^2 + m^2)/\sigma_0^2 - 1 + \log \sigma_0^2 - \log s^2 \right)$, applied to β with $\sigma_0 = \sigma_\beta$ and to ξ with prior $\mathcal{N}(\mu_\xi, \sigma_\xi^2)$ (with m replaced by $m - \mu_\xi$). The loadings use the conditional prior of Eq. (2), giving the per-row prior variance λ_ℓ^2:

$$\mathrm{KL}_\alpha = \frac{1}{2} \sum_{\ell=1}^{L} \sum_{j=1}^{p} \left(\frac{s_{\alpha_{\ell j}}^2 + m_{\alpha_{\ell j}}^2}{\lambda_\ell^2} - 1 + 2 \log \lambda_\ell - \log s_{\alpha_{\ell j}}^2 \right), \tag{11}$$

and the same form for δ with λ_q^b. When the batch subspace is enabled, $\mathrm{KL}(q(Z)) = \frac{1}{2} \sum (s_Z^2 + m_Z^2 - 1 - \log s_Z^2)$ joins $\mathrm{KL}_{\mathrm{glob}}$ and is likewise not rescaled by $n/|\mathcal{B}|$, since Z is indexed by batch rather than by cell.

Since λ is point-estimated rather than given a posterior, the half-Cauchy enters as the summed log-prior $\log p(\lambda) = \sum_\ell -\log(1 + \lambda_\ell^2)$ rather than as a KL term. Its gradient with respect to $\log \lambda_\ell$ is $-2\lambda_\ell^2/(1 + \lambda_\ell^2) \in (-2, 0]$, bounded in magnitude by 2 however large λ_ℓ grows. This bound is what delivers the asymptotically unbiased shrinkage claimed in Section 2: writing $\widehat{R}_\ell = \sum_j \left(m_{\alpha_{\ell j}}^2 + s_{\alpha_{\ell j}}^2 \right)$ for row ℓ's fitted second moment (the same quantity as in Lemma 5 of Appendix E), the scale's stationarity condition — available for the same reason, the graph penalty being free of λ_ℓ as well as of s, and $\log \lambda$ carrying no guard — $\widehat{R}_\ell/\lambda_\ell^3 - p/\lambda_\ell - 2\lambda_\ell/(1 + \lambda_\ell^2) = 0$ gives $\lambda_\ell^2 = (\widehat{R}_\ell/p)\big(1 + O(1/p)\big)^{-1}$, the loadings' mean square up to a vanishing relative correction, since the prior contributes $O(1)$ against the $O(p)$ pull of Eq. (11). However, a light-tailed prior on the scale has an unbounded gradient there, so its relative shift grows with the factor's strength rather than staying $O(1/p)$ uniformly.

A.5 Variational details

There is no closed-form coordinate ascent: the NB likelihood and the graph term couple the parameters nonlinearly, so every parameter is updated by the same gradient step. Write ϕ for the free parameters, namely the variational means and standard deviations of Appendix A.3 together with the point estimates λ and γ. Each standard deviation is carried on the log scale, $s_\bullet = \exp(\log s_\bullet)$, which enforces positivity without a constraint and makes the entropy contributions $-\log s_\bullet^2$ of Eq. (11) and its siblings linear in the parameter.

The three parts of Eq. (10) are differentiated in three different ways. All Gaussian parameters are reparameterized (Kingma & Welling, 2014; Rezende et al., 2014), $u^{(t)} = m_u + s_u \odot \varepsilon_{\mathrm{rp}}^{(t)}$ with $\varepsilon_{\mathrm{rp}}^{(t)} \sim \mathcal{N}(0, I)$, so the expected log-likelihood alone is Monte-Carlo estimated and its gradient is pathwise: for any parameter u and any draw,

$$\frac{\partial}{\partial m_u} = \frac{\partial}{\partial u}, \qquad \frac{\partial}{\partial \log s_u} = s_u \, \varepsilon_{\mathrm{rp}}^{(t)} \odot \frac{\partial}{\partial u}. \tag{12}$$

The KL gradients are analytic, and the graph penalty is differentiated by autodiff.

Every likelihood gradient factors through the natural parameter $\eta_{ij} = \log \mu_{ij}$. Differentiating Eq. (8) at $\mu_{ij} = e^{\eta_{ij}}$,

$$\frac{\partial}{\partial \eta_{ij}} \log \mathrm{NB}\big(Y_{ij} \mid e^{\eta_{ij}}, \theta_j\big) = Y_{ij} - (Y_{ij} + \theta_j) \frac{\mu_{ij}}{\mu_{ij} + \theta_j}, \qquad \frac{\partial^2}{\partial \eta_{ij}^2} = -(Y_{ij} + \theta_j) \frac{\theta_j \, \mu_{ij}}{(\mu_{ij} + \theta_j)^2}. \tag{13}$$

The second derivative is negative for every $Y_{ij} \geq 0$, so the log-likelihood is concave in η whatever the counts. Since η is affine in every block of Eq. (1), the remaining chain rules are one factor each: $\partial\eta_{ij}/\partial W_{i\ell} = \alpha_{\ell j}$, $\partial\eta_{ij}/\partial\alpha_{\ell j} = W_{i\ell}$, $\partial\eta_{ij}/\partial\beta_{hj} = X_{ih}$ and $\partial\eta_{ij}/\partial\gamma_{ig} = V_{jg}$.

The dispersions are the one block that does not factor through η, since θ_j enters the log-density directly as well as through the mean. With $\theta_j = e^{\xi_j}$ and ψ_0 the digamma function:

$$\frac{\partial}{\partial \xi_j} \log \mathrm{NB}\big(Y_{ij} \mid \mu_{ij}, \theta_j\big) = \theta_j \Big[\psi_0(Y_{ij}+\theta_j) - \psi_0(\theta_j) + \log\theta_j + 1 - \log(\mu_{ij}+\theta_j) - \frac{Y_{ij}+\theta_j}{\mu_{ij}+\theta_j}\Big]. \quad (14)$$

The KL gradients need no sampling. For the embeddings, and for the loadings under the conditional prior of Eq. (2),

$$\frac{\partial\,\mathrm{KL}_{W_{i\cdot}}}{\partial m_{W_{i\ell}}} = m_{W_{i\ell}}, \quad \frac{\partial\,\mathrm{KL}_{W_{i\cdot}}}{\partial \log s_{W_{i\ell}}} = s^2_{W_{i\ell}} - 1, \qquad \frac{\partial\,\mathrm{KL}_{\alpha}}{\partial m_{\alpha_{\ell j}}} = \frac{m_{\alpha_{\ell j}}}{\lambda^2_\ell}, \quad \frac{\partial\,\mathrm{KL}_{\alpha}}{\partial \log s_{\alpha_{\ell j}}} = \frac{s^2_{\alpha_{\ell j}}}{\lambda^2_\ell} - 1, \quad (15)$$

with the same pair for β and ξ at their own prior scales and means, and for δ at λ^b_q. Each pair has the same fixed point: a mean pulled toward the prior mean, and a scale that stops where the posterior variance matches the prior's. Only α (and δ, when the batch subspace is enabled) carries a learned scale in its denominator, which is what couples the loadings to λ_ℓ and drives the shrinkage.

The graph term touches the embedding means alone. Differentiating $\mathcal{R}(\mathcal{B})$ through the stop-gradient,

$$\frac{\partial \mathcal{R}(\mathcal{B})}{\partial m_{W_{i\cdot}}} = 2\lambda_g \frac{n}{|\mathcal{B}|}\Big(m_{W_{i\cdot}} - \mathrm{sg}\left[\overline{m}_{\mathcal{G}(i)}\right]\Big), \quad (16)$$

and a cell receives no gradient through its appearance in a neighbor's target. The penalty is therefore a pull of each cell toward a fixed point, recomputed each step but never differentiated, rather than a mutual attraction that could collapse the embedding.

Optimization details are as follows.

- **Optimizer.** AdamW with two learning-rate groups, the per-cell parameters $\{m_W, \log s_W, \gamma\}$ at $5\times$ the global rate. Weight decay applies to global parameters only, never to the per-cell ones. Gradients are clipped to norm 5.
- **KL annealing.** Linear warm-up $w_{\mathrm{KL}} = \min(1, \mathrm{epoch}/n_{\mathrm{warm}})$ with $n_{\mathrm{warm}} = \max\big(10, \min(50, \lfloor 0.02\,\mathrm{epochs}\rfloor)\big)$.
- **Learning-rate schedule.** Linear warm-up then cosine annealing to a floor.
- **Early stopping.** A plateau check on the clean full-batch ELBO, recomputed each epoch, with an adaptive threshold proportional to $|\mathrm{ELBO}|$ fixed once at warm-up completion.
- **Initialization.** Dispersions come from the method of moments, and scores and loadings from PCA on centered $\log(1+Y)$, with scores rescaled to unit standard deviation per factor and loadings rescaled by the score standard deviations so the product reproduces the PCA reconstruction. This gives high-variance factors larger loading norms at the start, an informative asymmetry for the prior to act on. The method-of-moments log-dispersions are clipped to $[-4, 4]$ before the empirical-Bayes prior location is taken as their mean.
- **Numerical guards.** $\log s_W$ is clamped to $[-6, 2]$. This bound enters both the ELBO and the denominator of Eq. (6), so a factor whose posterior variance has been driven to zero cannot produce an unbounded signal-to-noise ratio.
- **Weight decay as an implicit prior.** AdamW's decay reaches every global parameter, including $\log\lambda$ and the log standard deviations, so the optimized objective is Eq. (10) plus a small quadratic term the bound itself does not contain. At the default rate its effect over a full run is under 0.2% of the scale of $\log\lambda$, but the estimator is not exactly the stated one.
- **Factor selection.** None during training. The active set is read off once after convergence by Eq. (6). There is no pruning step and no selection threshold in the configuration.

A.6 BUILDING THE NEIGHBORHOOD GRAPH

The graph enters Eq. (4) only through the neighbor set $\mathcal{G}(i)$, so the two variants below differ in that set and in nothing else. Both are built once from the counts, before training, and then held fixed, and neither uses cell-type labels.

Default. The neighbor search runs in a principal-component embedding of the counts: let $\mathcal{P} = \mathrm{PCA}_{n_{\mathrm{pca}}}\big(\log(1+Y)\big)$ collect the top n_{pca} principal components of the log counts, centered across cells, and write $\mathcal{P}_{i\cdot} \in \mathbb{R}^{n_{\mathrm{pca}}}$ for cell i's coordinates in it. Each cell takes its k_{nn} nearest neighbors there, itself excluded:

$$\mathcal{G}(i) = \big\{\text{the } k_{\mathrm{nn}} \text{ cells } i' \neq i \text{ with smallest } \|\mathcal{P}_{i\cdot} - \mathcal{P}_{i'\cdot}\|_2\big\}. \tag{17}$$

Cross-batch. When batch correction is enabled, the same pull acquires a neighbor set that is required to span batches. Write $\mathcal{C}_b = \{i : b(i) = b\}$ for the cells of batch b. The set is built in three steps. First, $\log(1+Y)$ is centered *within* each batch, which removes the additive part of the batch shift before any distance is taken, and the same embedding is then recomputed on the residual:

$$\tilde{Y}_{i\cdot} = \log(1+Y_{i\cdot}) - \frac{1}{|\mathcal{C}_{b(i)}|}\sum_{i'\in\mathcal{C}_{b(i)}} \log(1+Y_{i'\cdot}), \qquad \mathcal{P} = \mathrm{PCA}_{n_{\mathrm{pca}}}\big(\tilde{Y}\big). \tag{18}$$

A Euclidean neighbor in this space is then a cell of the same cell type in a possibly different batch, rather than merely a cell from the same batch. Second, with a per-batch degree $k_{\mathrm{nn,per}} = \max\big(1, \mathrm{round}(k_{\mathrm{nn}}/B)\big)$, each cell takes its $k_{\mathrm{nn,per}}$ nearest neighbors *within every batch* and the union is kept, so the candidate set is guaranteed to reach across batches:

$$\mathcal{G}_b(i) = \big\{\text{the } k_{\mathrm{nn,per}} \text{ cells } i' \in \mathcal{C}_b \setminus \{i\} \text{ with smallest } \|\mathcal{P}_{i\cdot} - \mathcal{P}_{i'\cdot}\|_2\big\}, \qquad \mathcal{G}^+(i) = \bigcup_{b=1}^{B} \mathcal{G}_b(i). \tag{19}$$

Third, the candidates are capped at their globally nearest members, which bounds the degree on many-batch data and raises the share of neighbors of the same cell type, since nearer candidates are likelier to match:

$$\mathcal{G}(i) = \big\{\text{the } \min\big(2k_{\mathrm{nn}}, |\mathcal{G}^+(i)|\big) \text{ members of } \mathcal{G}^+(i) \text{ with smallest } \|\mathcal{P}_{i\cdot} - \mathcal{P}_{i'\cdot}\|_2\big\}. \tag{20}$$

Polański et al. (2020) introduce this within-batch neighbor search in BBKNN, a fast approximation to the mutual-nearest-neighbor matching of Haghverdi et al. (2018) that drops the reciprocity check. Since the target of Eq. (4) then averages cells of the same cell type drawn from other batches, the pull moves matching cell types together across batches while leaving the cell type structure intact.

Masked mean. Under Eq. (17) every cell has exactly k_{nn} neighbors, but Eq. (20) leaves $|\mathcal{G}(i)|$ varying from cell to cell: a batch holding fewer than $k_{\mathrm{nn,per}}$ cells contributes less than its quota, so $|\mathcal{G}^+(i)|$ need not reach $Bk_{\mathrm{nn,per}}$, and the cap binds only on those cells whose candidate set exceeds $2k_{\mathrm{nn}}$. The mean in Eq. (4) is therefore taken over each cell's own $|\mathcal{G}(i)|$, which is what makes the pull comparable across cells: a cell in a small batch is drawn toward its neighbors as strongly as a cell in a large one, rather than proportionally to how many neighbors it happens to have.

A.7 FACTOR SELECTION IN DETAIL

Margin. Across the six benchmark datasets and ten seeds, the largest *inactive* SNR_ℓ is 0.92 (PBMC) and the smallest *active* one is 2.33 (Baron), so the selected set is unchanged by any cut in $[0.92, 2.33)$. The remaining four datasets sit far above the upper end: Zeisel 5.39, Muraro 9.07, Zhengmix8eq 14.03, Segerstolpe 20.24. That insensitivity margin, rather than a tuned constant, is what makes the unit cutoff safe in finite samples.

Censoring. The rule returns $L_{\mathrm{active}} = |\widehat{\mathcal{S}}|$, which Theorem 2 bounds below by the number of true factors when every true factor is strong enough (no upper bound is claimed), and which cannot exceed the ceiling L. The ceiling binds at the shipped $L = 20$ on Zeisel, Muraro and Segerstolpe, where $L_{\mathrm{active}} = L$ on every seed. However, that is expected rather than pathological, since the number of latent dimensions needed to separate cell types is not the number of cell types, and Appendix D.1 shows the count plateaus once the ceiling is raised.

Rounding. For downstream clustering, dropping an informative factor is irreversible, while carrying an extra noise factor costs almost nothing. A factor sitting at $\mathrm{SNR}_\ell \approx 1$ is therefore kept.

Relation to the horseshoe. The regularized-horseshoe shrinkage coefficient $\kappa_\ell = 1/(1 + \mathrm{signal}_\ell)$ has the canonical selection rule $\kappa_\ell < \frac{1}{2}$. Under the identification $\mathrm{signal}_\ell \leftrightarrow \mathrm{SNR}_\ell$ this coincides with $\mathrm{SNR}_\ell/(1+\mathrm{SNR}_\ell) > \frac{1}{2}$ and hence with $\mathrm{SNR}_\ell > 1$. We note the correspondence as an analogy, not an identity: the fit here places a half-Cauchy on the scale and estimates it by MAP, and the rule thresholds an empirical reliability of the embedding rather than $\kappa(\lambda_\ell)$.

A.8 Default hyperparameters

Defaults are the keyword arguments of `fit_ripple`, which the benchmark's configuration file mirrors key for key, and a test asserts the two agree.

symbol	meaning	default
L	latent dimension ceiling	20
T	Monte-Carlo samples per step	4
σ_β	prior std of β	5
σ_ξ	prior std of ξ	1
μ_ξ	prior location of ξ	empirical Bayes
Q	batch-subspace ceiling	0 (off)
λ_g	graph penalty weight	100
k_{nn}	graph neighbors	5
n_{pca}	PCA dimensions for the graph space	50
$\lvert\mathcal{B}\rvert$	minibatch size	256
–	learning rate, globals	10^{-3}
–	learning rate, per-cell parameters	5×10^{-3}
–	weight decay (globals only)	10^{-4}
–	gradient-norm clip	5
–	maximum epochs	15,000
–	LR warm-up epochs, then cosine to 10^{-5}	100
–	ELBO-plateau patience	50
–	ELBO-plateau threshold fraction	2×10^{-4} †

† Baron is the single per-dataset departure from this table, at 5×10^{-5}. A smaller fraction makes the ELBO-plateau test more conservative, so the fit trains longer before stopping. It does not give Baron an unusual amount of training: its ten fits run 815–976 epochs, inside the 530–1,340 range the other five datasets reach under the default. Every other value here is used unchanged wherever it applies; the batch-subspace ceiling Q follows the batch-correction flag, which is enabled on the two datasets whose batches separate.

B Benchmark details

B.1 Datasets

	Tissue	Protocol	n	p	C	Batches	Rarest
Baron	Human pancreas	inDrop	8,569	15,117	14	4 donors	0.08%
Muraro	Human pancreas	CEL-Seq2	2,122	16,062	9	4 donors	0.14%
Segerstolpe	Human pancreas	Smart-seq2	2,209	18,894	14	10 donors	0.09%
Zeisel	Mouse brain	STRT-seq	3,005	2,800	7	76 plates	3.26%
Zhengmix8eq	Human blood	10x Chromium	3,994	10,600	8	—	9.96%
PBMC	Human blood	10x Chromium	11,989	3,345	9	—	0.73%

Table 1: The six benchmark datasets after filtering. C is the number of annotated cell types and "Rarest" the share of cells in the smallest one. Zhengmix8eq and PBMC carry no batch variable. Segerstolpe and Zeisel carry one but are excluded from the batch metrics, for the reasons given below.

Table 1 lists the six datasets used in Section 4.1. They were chosen to vary along the axes that make this problem hard, not to agree with one another: three tissues, five protocols, a $5.6\times$ spread in cell count, number of cell types from 7 to 14, and rarest cell type ranging from 2 cells to a tenth of the data. Balance and difficulty do not travel together here. Zhengmix8eq is the most balanced, its eight populations FACS-purified and mixed in equal proportion, yet five of them are T-cell subsets that no unsupervised method separates cleanly. Segerstolpe is the hardest by the score methods actually reach on it, averaging 0.517 ARI across the seven unsupervised methods against Muraro's 0.792, and it carries 14 cell types of which three are annotated as unclassified or co-expressing and retained as given.

Data preprocessing. Preprocessing is light and no step uses cell-type labels. Gene j is kept when $\sum_i \mathbb{1}\{Y_{ij} > 0\} \geq 10$. However, Zeisel takes the stricter $\sum_i \mathbb{1}\{Y_{ij} > 10\} \geq 10$, which reproduces the panel of Risso et al. (2018). Cells are dropped where the source annotation records no call, following each dataset's own vocabulary: Muraro loses its `unknown` and `unclear` cells and Segerstolpe its `not applicable` ones, while PBMC additionally requires 200 detected genes per cell, the filter its source pipeline applies. No mitochondrial, ribosomal or highly-variable-gene selection is performed on any dataset.

Batch structure. Batch metrics are reported on Baron and Muraro alone, on the following criterion: a dataset qualifies when its batch factor is a nuisance a method may remove without significantly removing biology. Both Baron and Muraro split into four donors balanced across cell types and so qualify. On the remaining four we do not integrate batch, for reasons that differ by dataset. Segerstolpe's ten donors are six healthy and four type-2 diabetic, so donor is confounded with disease and mixing donors would erase the contrast the data was collected to show. Zeisel's $B = 76$ plates hold a median of 39 cells each, and are the sampling design itself rather than a nuisance factor. Zhengmix8eq is a single sequencing run, and the PBMC file pools two 10x runs whose indicator its source pipeline discards, so neither has a batch variable to use.

Three properties of the six datasets bear on the comparison. Muraro's counts are UMI-collision-corrected transcript estimates and are not integers. RIPPLE, scVI and scANVI evaluate the negative binomial likelihood on them as they are. ZINB-WaVE, scGBM and scVIC require integer input and round or ceiling first, as does scMAE, which carries no count likelihood. Each method therefore sees the input its implementation accepts, and the proposed method is on the unrounded side of that split. The library size behind the mean offset is computed before gene filtering on the five datasets prepared in `R` and after it on PBMC, so that offset is whole-transcriptome depth in the first case and within-panel depth in the second. The file we call PBMC is the pooled 8k and 4k peripheral-blood runs distributed with `scvi-tools` rather than the smaller 2.7k dataset often given that name, which is why it is the largest here at $11{,}989$ cells.

B.2 Baselines and configurations

	Model class	Labels
RIPPLE	NB factor model, no encoder	no
scVI	NB/ZINB autoencoder	no
scANVI	scVI plus a cell-type latent variable	**yes**
Harmony	PCA, then linear correction	no
ZINB-WaVE	ZINB matrix factorization	no
scMAE	Masked autoencoder (MSE)	no
scGBM	Poisson bilinear model	no
scVIC	ZINB autoencoder, mixture prior	count only

Table 2: The eight methods. Only scANVI trains on cell-type labels. scVIC uses the *number* of cell types to size its mixture prior, the same count k-means receives as oracle k.

Table 2 summarizes the eight methods shown in Section 4.1. Every reported figure is a mean over the ten seeds: $\{123, 321, 456, 654, 789, 987, 555, 777, 888, 999\}$. Throughout, *oracle* k denotes the number of annotated cell types in a dataset (column C in Table 1), which is the cluster count every method's embedding is scored at.

Of the baselines, only scANVI trains on the full cell-type labels, and we treat it as the supervised ceiling hereafter. scVIC sits between the two regimes. It places a Gaussian mixture over the latent space and sizes that mixture at the number of cell types, so it consumes oracle k in the same way k-means does, as a count and not as an assignment. Its data loader also receives the label vector, though the model is documented to read only the count. We could not confirm that from the legacy environment the method requires.

The batch dummies of Eq. (9) enter X for RIPPLE and ZINB-WaVE on every dataset carrying a batch index. Batch *integration* is a separate matter, and is enabled only where Appendix B.1 finds batch to be a nuisance, which is Baron and Muraro: there scVI and scANVI take a batch key, scVIC conditions its decoder on a batch index, Harmony corrects principal components after the fact, and RIPPLE turns on its random batch effect and cross-batch graph. When no batch is supplied, Harmony reduces to plain PCA. Batch metrics are reported on those two datasets alone. scMAE cannot receive batch at all: the published architecture takes the expression matrix and nothing else, so supplying batch would mean adding a mechanism its authors never defined. scGBM exposes a batch factor and fits per-batch gene intercepts, but supplying batches makes its fit diverge on these data. Its internal decomposition returns undefined values and halts before producing an embedding, on Baron and Muraro ($B = 4$) as well as Zeisel ($B = 76$), so we run it at the package default of a single batch throughout. Its scores are those of a single-batch fit.

We run all methods at their default or recommended settings, RIPPLE included: $\lambda_g = 100$, $k_{\mathrm{nn}} = 5$ and $L = 20$ are the shipped defaults and are used unchanged on all six datasets, and no per-dataset configuration file overrides any of them. Two per-dataset settings do exist and are disclosed where they arise: the batch-correction flag and its subspace ceiling $Q = 5$, enabled on Baron and Muraro alone by the criterion above and matched by the batch key given to scVI and scANVI, and Baron's ELBO-plateau tolerance, a convergence setting recorded in Appendix A.8. All runs used pinned versions on a single machine: `scvi-tools` 1.4.2, `harmonypy` 0.2.0, `scikit-learn` 1.8.0 and `PyTorch` 2.10.0 in the primary environment, `zinbwave` 1.32.0 and `scGBM` 1.2.0 under R 4.5.3, and `scvi` 0.6.5 under Python 3.7 for scVIC, which requires a legacy environment. The scMAE model is a minimal port of the authors' released code at `https://zenodo.org/records/10465991`.

B.3 Evaluation protocol

Every method is scored through the same routine. The embedding is clustered by k-means at oracle k, which is again the number of distinct cell-type labels shown in Table 1, with 20 restarts from a fixed seed and the best inertia retained. Let $C_{cq} = |\{i : y_i = c,\ \hat{y}_i = q\}|$ be the contingency between true cell type c and cluster q. Clusters are matched to cell types by the permutation

$$\pi^\star = \arg\max_{\pi} \sum_{c=1}^{C} C_{c\,\pi(c)}, \tag{21}$$

solved as a linear sum assignment problem. Accuracy is $n^{-1} \sum_c C_{c\,\pi^\star(c)}$, and macro-$F_1$ the unweighted mean over the C cell types of the per-cell-type F_1 that $\pi^\star$ induces.

The remaining clustering metrics are taken from `scikit-learn` at their standard definitions, with arithmetic averaging for both NMI and AMI. Two further metrics read the embedding without clustering it at all, and so cannot be confounded by the choice of clusterer: cLISI, the label-purity of a cell's neighborhood, and label-ASW, the silhouette width of the true cell types. Both are computed with `scib-metrics` (Luecken et al., 2022) on an exact Euclidean 90-nearest-neighbor graph. We use `scib-metrics` to compute batch metrics on Baron and Muraro.

B.4 Per-dataset results

Figures 3 to 5 resolve the grand averages of Section 4.1 into the six datasets they average over. Each dot is a mean across ten seeds and each interval is a 95% percentile bootstrap confidence interval of that mean, drawn from 2,000 resamples of the seeds. Table 3 prints exact numbers of Figures 3 to 5.

The figures and the table answer different questions. The intervals show how far run-to-run variation reaches, while the printed values make the win counts quoted in Section 4.1 auditable. However, six of the eighteen contests are settled by less than 5% of a panel's vertical span, and the narrowest, NMI

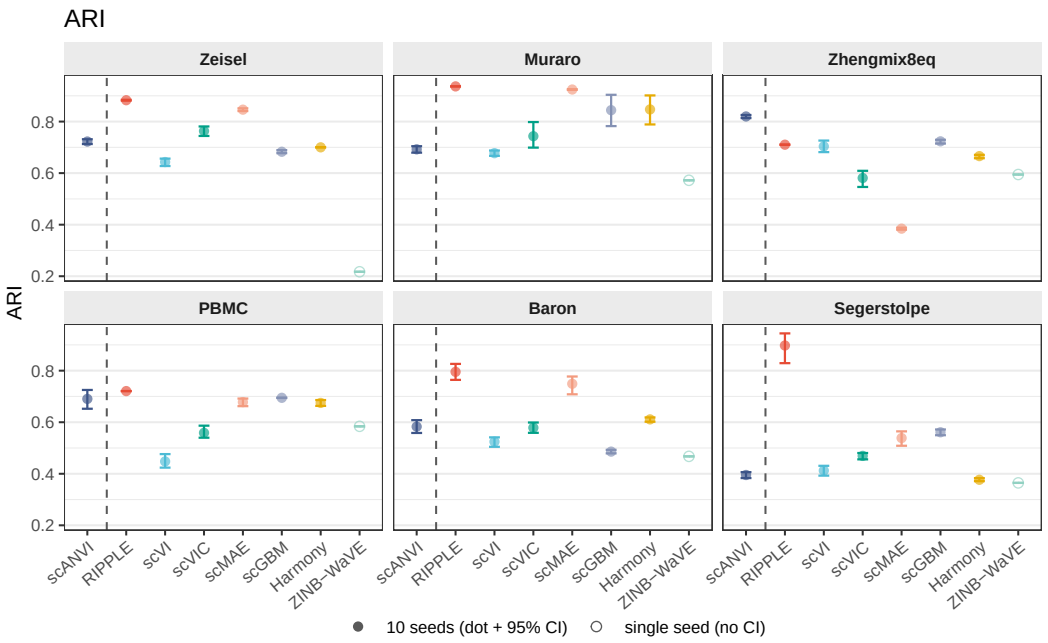


Figure 3: ARI per dataset. Each dot is a mean over ten seeds and each bar a 95% percentile bootstrap confidence interval of that mean. ZINB-WaVE ran one seed and is drawn as an open circle without an interval. scANVI is supervised and lies left of the dashed line.

on PBMC, by 0.0007—a gap no reader can resolve on the figure, and the one that places RIPPLE first on 6 of 6 datasets instead of 5 of 6.

What separates the methods is consistency more than peak performance. Across the six datasets scMAE spans 0.385 to 0.924 ARI and Harmony 0.377 to 0.847, while RIPPLE spans 0.710 to 0.936. scMAE reaches within 0.012 of RIPPLE on Muraro and then falls to 0.385 on Zhengmix8eq, just over half of RIPPLE's 0.710, and it is that collapse rather than any single win that separates the two on the average. A method that is occasionally best and occasionally poor is harder to deploy on an unlabeled dataset than one that is reliably near the top, since nothing in an unsupervised setting reveals which regime the data are in.

B.5 Runtime

Wall-clock time per fit is given in Figure 6, by dataset and on a logarithmic scale, since the methods span six orders of magnitude. RIPPLE, scVI, scANVI, scVIC and scMAE were fitted on one NVIDIA RTX 4070 Super, and Harmony, scGBM and ZINB-WaVE on an AMD Ryzen 9 5900X. The two groups are separated in the figure and should not be read against each other. Averaged over the six datasets, RIPPLE takes 263 s per fit against 268 s for scVI and 455 s for scANVI, its own range running from 31 s on Zeisel to 748 s on Baron. However, ZINB-WaVE stands apart at 44,295 s on average and 149,115 s on Baron, which is what confines it to a single seed.

RIPPLE carries one variational parameter row per cell, which invites the objection that it should scale poorly in n. It does not, because the cost is dominated by the likelihood over all np entries rather than by the per-cell parameters. Fit time tracks np at 3 to 7.5 seconds per million entries across the six datasets, and the two extremes bracket it: Zeisel is the smallest problem at 8.4M entries and the fastest fit at 31 s, while Baron is the largest at 130M and the slowest at 748 s. Measured per cell instead, RIPPLE costs 9.9 s per thousand cells on PBMC, the largest dataset at 11,989 cells, against 87.3 s per thousand on Baron, which has fewer cells but 4.5× the genes.

The comparison with the amortized models is therefore not one-sided. scVI is cheaper per cell than RIPPLE on the two widest datasets and costlier on the two narrowest, and neither ordering holds

	Baron	Muraro	Segerstolpe	Zeisel	Zhengmix8eq	PBMC	Grand
ARI							
scANVI†	0.583	0.692	0.395	0.723	0.820	0.690	0.651
RIPPLE	**0.795**	**0.936**	**0.898**	**0.883**	0.710	**0.721**	0.824
scMAE	0.749	0.924	0.539	0.846	0.385	0.678	0.687
scGBM	0.486	0.844	0.561	0.683	**0.723**	0.695	0.665
Harmony	0.611	0.847	0.377	0.700	0.666	0.675	0.646
scVIC	0.579	0.743	0.469	0.763	0.581	0.559	0.616
scVI	0.524	0.677	0.412	0.643	0.704	0.448	0.568
ZINB-WaVE	0.468	0.572	0.365	0.217	0.595	0.584	0.467
NMI							
scANVI†	0.790	0.814	0.691	0.759	0.857	0.820	0.789
RIPPLE	**0.834**	**0.899**	**0.907**	**0.840**	**0.810**	**0.791**	0.847
scMAE	0.819	0.887	0.706	0.808	0.635	0.773	0.771
scGBM	0.733	0.843	0.694	0.740	0.792	0.790	0.766
Harmony	0.768	0.850	0.649	0.749	0.770	0.757	0.757
scVIC	0.792	0.833	0.761	0.758	0.653	0.741	0.756
scVI	0.737	0.804	0.688	0.646	0.759	0.637	0.712
ZINB-WaVE	0.644	0.721	0.594	0.299	0.681	0.701	0.607
macro-F_1							
scANVI†	0.529	0.732	0.417	0.750	0.853	0.757	0.673
RIPPLE	**0.676**	**0.851**	**0.652**	**0.928**	0.752	0.733	0.766
scMAE	0.619	0.848	0.526	0.905	0.424	0.693	0.669
scGBM	0.450	0.709	0.377	0.716	**0.758**	**0.745**	0.626
Harmony	0.520	0.795	0.397	0.726	0.722	0.665	0.637
scVIC	0.548	0.777	0.479	0.752	0.641	0.641	0.640
scVI	0.479	0.741	0.430	0.710	0.756	0.506	0.604
ZINB-WaVE	0.327	0.547	0.332	0.306	0.645	0.548	0.451

Table 3: Per-dataset means over ten seeds, and the grand mean over datasets. Bold marks the best unsupervised method on each dataset. scANVI (†) is supervised and excluded from that comparison.

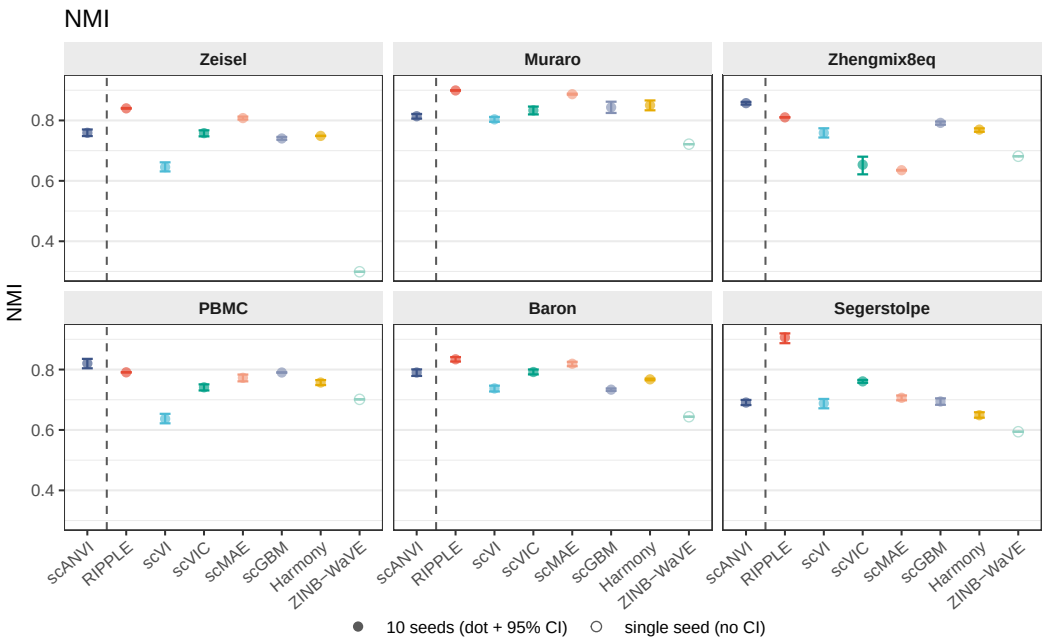


Figure 4: NMI per dataset, read as in Figure 3.

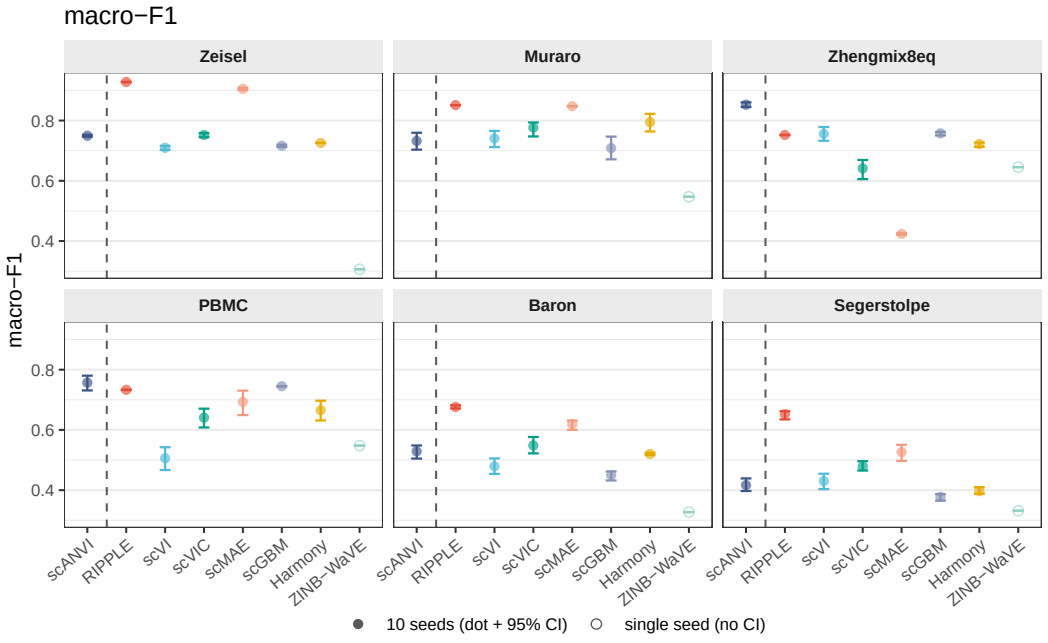


Figure 5: Macro-F_1 per dataset, read as in Figure 3.

throughout. What the encoder gains is amortization at fit time, which is a different quantity from the total cost of a fit, and on these six datasets the two effects roughly cancel.

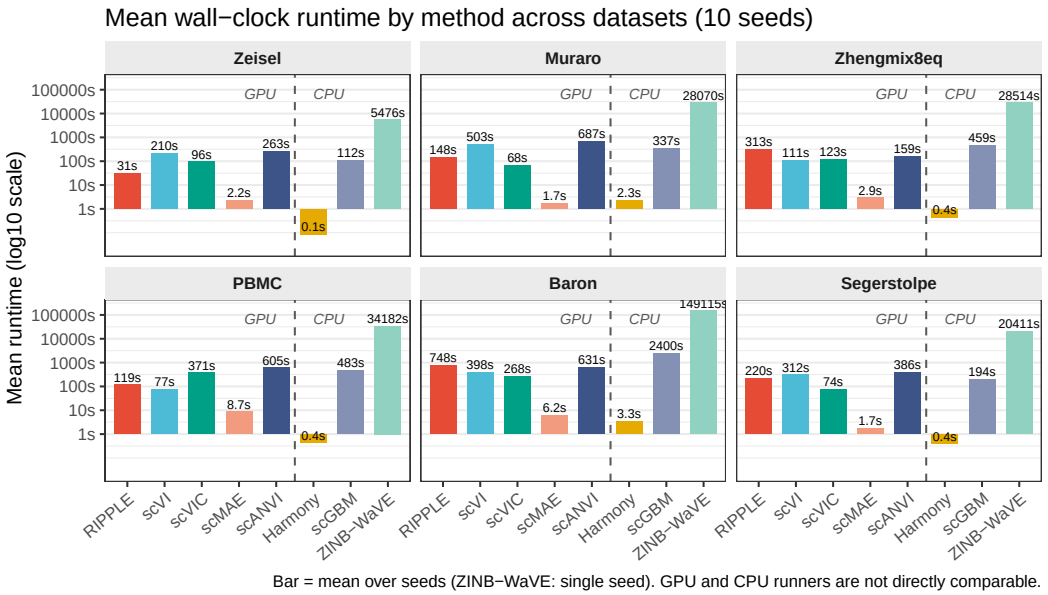


Figure 6: Wall-clock time per fit, by dataset, on a logarithmic scale. Bars are means over seeds and carry no interval. The rule separates the methods fitted on GPU from those fitted on CPU, which are not comparable with one another.

B.6 The wider metric battery

Section 4.1 reports three of the eleven metrics we compute. Figures 7 and 8 give the remaining eight averaged over datasets, in the same three-mark form as Figure 1. The ordering does not change: RIPPLE leads every biology-preserving metric, and the batch metrics separate the batch specialists from the rest.

The eight metrics do not carry equal information. cLISI is saturated on this method set, at 0.9998 for RIPPLE against 0.9996 for the best baseline and 0.945 for the worst, so it separates almost nothing and should not be read as corroboration. The label-agreement metrics spread much further, and RIPPLE's lead over the strongest baseline is largest on the two that weight assignment directly, at $+0.137$ ARI and $+0.127$ accuracy, and smaller on the information-theoretic pair, at $+0.075$ NMI and $+0.076$ AMI. AMI tracks NMI to within 0.008 everywhere, which is expected once k is fixed at oracle k for every method, so reporting both adds little.

The batch metrics order the methods differently. Among the unsupervised methods, scVI, scVIC and Harmony lead iLISI, batch-ASW and graph connectivity. On kBET the ordering is the same at the top, with scVI at 0.254 and Harmony at 0.229 against RIPPLE's 0.126. The gaps are wide where the metric is sensitive: on iLISI, the leaders reach 0.45 to 0.48 against RIPPLE's 0.35, while on graph connectivity every unsupervised method sits between 0.886 and 0.986, which makes that metric nearly as uninformative here as cLISI in the clustering metrics.

The battery splits by what it measures. Accuracy, AMI, cLISI and label-ASW join ARI, NMI and macro-F_1 on the biology-preserving side, of which cLISI and label-ASW are clusterer-independent and read the embedding directly (Luecken et al., 2022). Integration LISI (iLISI), kBET, batch-ASW and graph connectivity measure batch mixing, and are defined only on Baron and Muraro for the reason given in Appendix B.1. Graph connectivity sits with the batch metrics rather than the biological ones, following the taxonomy of Luecken et al. (2022).

B.7 The bio-batch tradeoff

Figure 9 places each unsupervised method on two composite axes, where upper-right is best on both axes. For a metric t on dataset d, write $v_{td}(a)$ for method a's seed-mean. Each cell is min-max

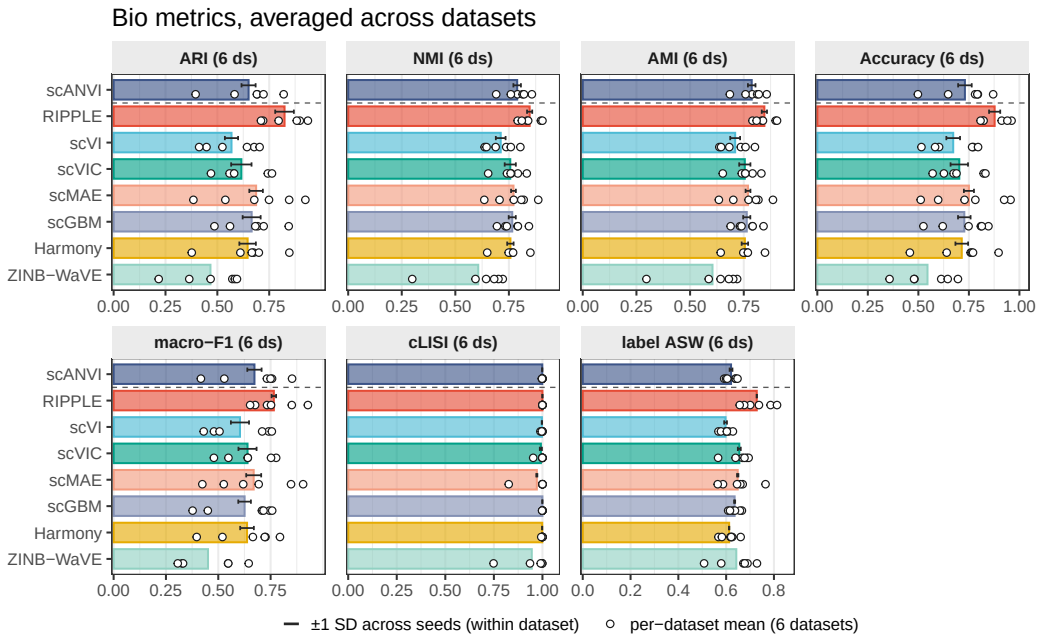


Figure 7: Biology-preserving metrics averaged over all six datasets, in the three-mark form of Figure 1: bar for the grand mean, whisker above for ± 1 SD across seeds pooled within dataset, circles below for the per-dataset means.

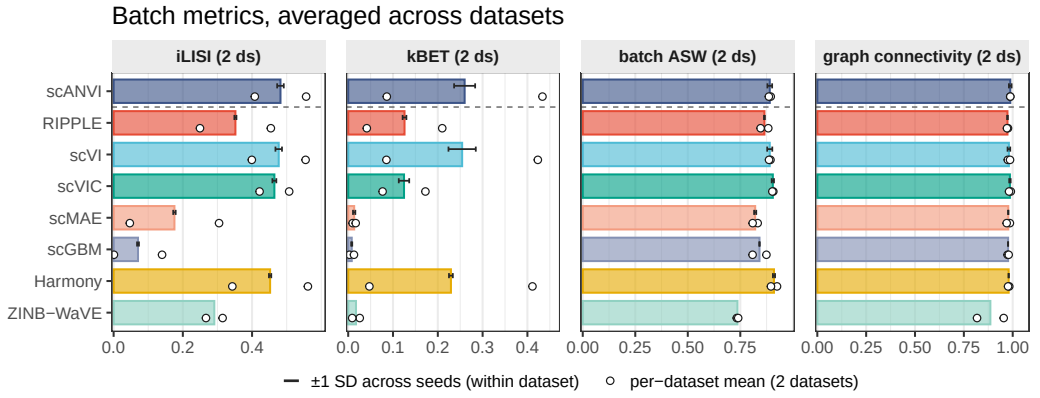


Figure 8: Batch metrics, read as in Figure 7. These average over Baron and Muraro only, as the facet labels record.

normalized across the method set $\mathcal{M}$ and the axis is the unweighted mean over its cells,

$$\tilde{v}_{td}(a) = \frac{v_{td}(a) - \min_{a' \in \mathcal{M}} v_{td}(a')}{\max_{a' \in \mathcal{M}} v_{td}(a') - \min_{a' \in \mathcal{M}} v_{td}(a')}, \qquad \text{score}(a) = \frac{1}{|\mathcal{T}|\,|\mathcal{D}|} \sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} \tilde{v}_{td}(a), \quad (22)$$

over $|\mathcal{T}| = 7$ metrics and $|\mathcal{D}| = 6$ datasets for the biological axis, and 4 over 2 for the batch axis. Normalizing within a dataset couples the methods through the endpoints, so a score is meaningful only relative to this method set and not as an absolute quantity. ZINB-WaVE is worst on nearly every cell and is therefore placed at the lower-left corner of Figure 9.

The two axes trade off across the seven methods. scVI integrates batches best at 0.948 but places sixth of seven on biology at 0.435. In contrast, RIPPLE leads the biological axis at 0.985 but places

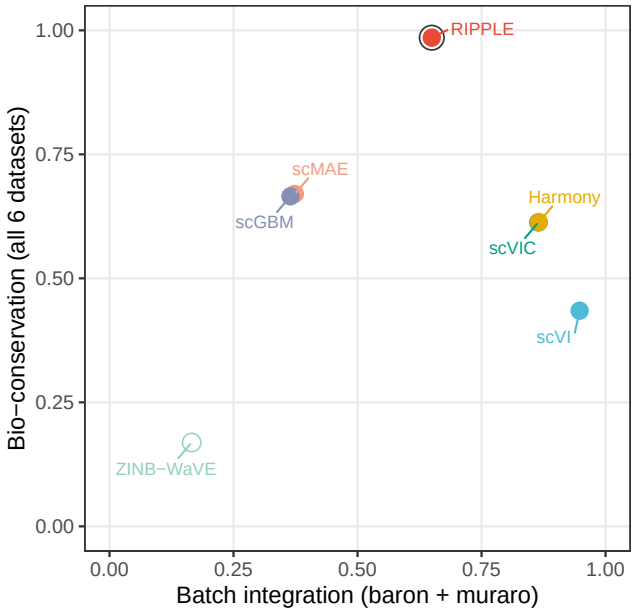


Figure 9: The two composite axes, over the seven unsupervised methods. Both are means of min-max normalized metrics and are meaningful only relative to this method set. The biological axis averages seven metrics over six datasets, the batch axis four over two.

fourth on integration at 0.650. Harmony and scVIC sit with scVI on the mixing side, and scMAE and scGBM on RIPPLE's side of the tradeoff, though both trail it by more than 0.31 on biology. Several of these pairs are separated by less than 0.005, so only the extremes of each axis should be read as ordered. Batch integration is an opt-in flag in RIPPLE, not a design goal, and we read this figure as positioning, not as a performance claim.

B.8 Uncertainty and hypothesis testing

Two things vary in this benchmark and only one of them is estimation error. Seeds vary within a dataset and datasets vary from one another, but every method is scored on the same six datasets, so dataset difficulty is a block effect that cancels in any comparison between two methods. Since the claim in Section 4.1 concerns a grand average, it generalizes over datasets, which makes the dataset the unit of replication and the seed a repeated measurement of it. Inference is therefore run at the dataset level, and seed-level variation is reported as an interval rather than as a test.

Table 4 gives a Friedman test over the six datasets, blocked by dataset, across the seven unsupervised methods. Equality is rejected on all three metrics, at $\chi^2 = 21.64$ and $p = 0.0014$ for ARI, $\chi^2 = 23.43$ and $p = 0.0007$ for NMI, and $\chi^2 = 18.29$ and $p = 0.0056$ for macro-F_1. RIPPLE holds the best mean rank on each, at 1.17, 1.00 and 1.50 of a possible seven. Figure 10 shows the same ranks with the range each method attains. scANVI is excluded from the pool, since a method trained on the labels is not competing on the same terms.

Nemenyi's critical difference, the post-hoc that follows the omnibus, is 3.68 mean ranks for this design. It separates RIPPLE from scVI and ZINB-WaVE on ARI and NMI, and from ZINB-WaVE on macro-F_1. Pairing two methods across the six datasets instead is a different procedure and a weaker one here: a dataset-level signed-rank test at $N = 6$ attains a smallest two-sided p of 0.031, reached only by a clean sweep, against Holm's tightest threshold of 0.008 over six competitors, so no comparison of that kind can clear correction at any effect size. The remaining pairwise argument is therefore carried by effect sizes.

Those effect sizes are differences in seed means against RIPPLE, each with a 95% percentile bootstrap interval over the ten paired seed differences, drawn from 10,000 resamples. Percentile intervals are used in preference to bias-corrected and accelerated ones since at ten seeds the acceleration es-

	ARI	NMI	macro-F_1
RIPPLE	1.17	1.00	1.50
scMAE	3.17	3.17	3.00
scGBM	3.33	3.83	4.17
Harmony	4.00	4.00	4.00
scVIC	4.67	4.00	4.00
scVI	5.33	5.50	4.83
ZINB-WaVE	6.33	6.50	6.50
Friedman χ^2	21.64	23.43	18.29
p	0.0014	0.0007	0.0056

Table 4: Mean rank across the six datasets, over the seven unsupervised methods (1 = best), with the Friedman omnibus beneath. Nemenyi's critical difference at $\alpha = 0.05$ is 3.68 mean ranks for this design.

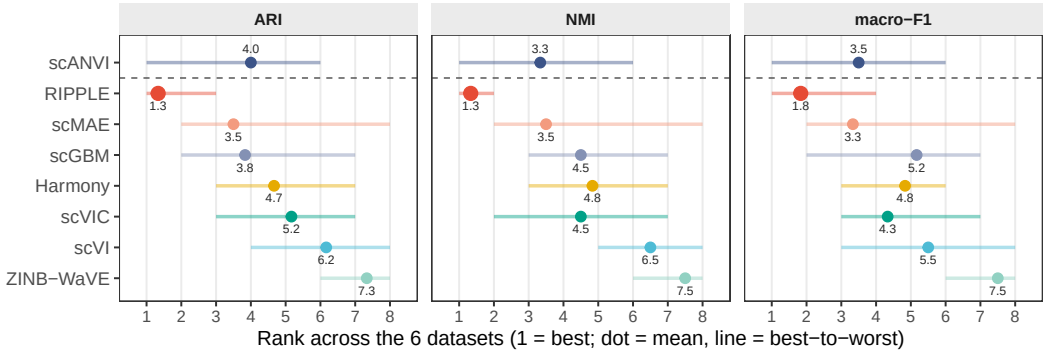


Figure 10: Mean rank across the six datasets, with the range each method attains. Ranks here run over all eight methods, including supervised scANVI. The Friedman test of Table 4 ranks the seven unsupervised methods only.

timate is noisier than the correction it would apply. ZINB-WaVE ran a single seed disjoint from the ten, so its differences carry no interval.

Table 5 collects the contests decided by less than 5% of their panel's span, over the eighteen unsupervised contests and the one against the supervised ceiling, the ones a reader cannot adjudicate from Figures 3 to 5. Five of the seven exclude zero, including the 0.0007 NMI gap on PBMC that carries the 6-of-6 count. Two do not. RIPPLE trails scGBM on macro-F_1 on Zhengmix8eq by a margin the ten seeds do not resolve, and we count that dataset as a loss rather than claim a tie. The other is the 0.0303 ARI gap over scANVI on PBMC, which the interval does not resolve. Counting that dataset as a tie, RIPPLE is ahead of scANVI on four of six by ARI, tied on one and behind on Zhengmix8eq.

Metric	Dataset	Rival	Δ	95% CI
ARI	Muraro	scMAE	$+0.0121$	$[+0.0108, +0.0136]$
ARI	Zhengmix8eq	scGBM	-0.0130	$[-0.0192, -0.0040]$
NMI	PBMC	scGBM	$+0.0007$	$[+0.0003, +0.0011]$
macro-F_1	Muraro	scMAE	$+0.0036$	$[+0.0030, +0.0043]$
macro-F_1	Zeisel	scMAE	$+0.0227$	$[+0.0194, +0.0260]$
macro-F_1	Zhengmix8eq	scGBM	-0.0055	$[-0.0102, +0.0016]^*$
ARI	PBMC	scANVI†	$+0.0303$	$[-0.0037, +0.0692]^*$

Table 5: The contests settled by less than 5% of a dotplot panel's vertical span, with RIPPLE's margin over its nearest rival and a 95% bootstrap interval. Asterisks mark the two intervals that span zero.

No seed-level p-value is quoted. An exact signed-rank test on ten seeds has a smallest attainable two-sided p of 0.00195, and 27 of 30 ARI comparisons, 27 of 30 for NMI and 26 of 30 for macro-F_1 sit exactly on that floor, reporting only that all ten differences share a sign. RIPPLE's within-dataset seed SD falls as low as 0.0004, which would render any nonzero gap significant. An interval says what such a p-value cannot, which is how large a gap is and how precisely it is known.

B.9 scANVI'S PREDICTED-LABEL HEAD

scANVI is trained on every cell's label and carries a classifier head that predicts them directly. Scored against those same labels, with no cell held out, it reaches 0.974–0.996 ARI and 0.987–0.997 accuracy.

The head is also a narrower object than the space it sits on. It emits a cell-type assignment and nothing more, so no other downstream analysis can be performed with it. Furthermore, no other method here predicts cell types at all, making a fair comparison infeasible. Section 4.1 scores scANVI's latent embedding instead, clustered by k-means at oracle k exactly as every other method's is, which reaches 0.395–0.820 ARI. A supervised classifier recovering its own training labels is a statement about the classifier, while the embedding comparison asks whether the geometry a method induces separates cell types.

B.10 OPTIMIZATION BIMODALITY

The variational objective is not convex, and on two datasets the fit reaches two distinct optima across seeds. On Segerstolpe eight of ten seeds land near 0.94 ARI and two near 0.72, which is the wide interval that dataset carries in Figure 3 and 76% of RIPPLE's pooled ARI seed variance. Baron supplies the remaining 24%, and it splits about as evenly as ten seeds can: five between 0.723 and 0.791, five between 0.828 and 0.880, and nothing in the 0.037 between. The other four datasets contribute nothing measurable, at seed SDs of 0.0029 and below.

What splits on Baron is the partition, not the representation. The same five seeds occupy the bottom five on ARI, NMI, AMI and accuracy alike, at Spearman ≥ 0.95 against ARI in each case, so it is one reproducible event rather than a quirk of a single metric. Label-ASW, which scores the embedding without clustering it, spans 0.004 across those same ten seeds and is if anything inversely ranked ($\rho = -0.24$). The embedding is stable, and k-means lands in two places on it, and ARI registers that most sharply because it is quadratic in pair agreement.

A label-free rule removes the bad basin. Fitting ten times and keeping the run whose embedding is most concordant with the others—the medoid of the restarts by pairwise ARI—reaches 0.946 on Segerstolpe on every seed, and lifts Baron from a 0.795 single-fit mean to 0.837. Five restarts suffice, at 0.946 and again every seed in the good basin. The mean pairwise ARI across restarts is itself the label-free diagnostic for when the rule is needed: 0.86 on Baron and 0.90 on Segerstolpe, against 0.99 on Zeisel, whose ten seeds are unimodal and where the rule is correspondingly inert—medoid and single-fit means agree to four decimals at 0.883. The rule costs a constant factor in compute and is available to any method here, so the benchmark reports single-fit numbers throughout.

C ATLAS DETAILS

C.1 THE PANEL AND THE REFERENCE FIT

We fit the HLCA core as published, with 61 cell types annotated at `ann_finest_level` (Sikkema et al., 2023). Preprocessing is deliberately light, and no step uses a cell-type label. We keep HLCA's own cross-study filter, `min_genes` $= 200$, which removes 2 cells and leaves the $584{,}942$ we fit, and we add no atlas-wide mitochondrial threshold, as HLCA itself imposes none. Library size is computed on the full transcriptome before any gene is dropped, so the offset O_i measures sequencing depth rather than depth within the panel.

Gene panel. We keep genes detected in at least ten cells, then select $3{,}000$ highly variable genes with `scanpy` (Wolf et al., 2018), by Pearson residuals on raw counts (Lause et al., 2021), passing `batch_key` $=$ `dataset` and $\theta = 100$. The residual clip is passed explicitly as $\sqrt{n}$ over all cells. This matters: `scanpy` assigns the clip once in its batched path, from whichever batch it visits

first, and since batches are visited in sorted order the panel would otherwise be a function of the dataset *names*. Selection runs on all $584{,}942$ cells, so no population is over-weighted relative to its abundance in the atlas.

Labels and covariates. Evaluation uses `ann_finest_level`, the 61 cell types. The cell design X_i carries an intercept, 13 dataset dummies and the centered detection rate. The gene design V_j carries three columns: an intercept, the fraction of cells in which gene j is zero, and $\log(1+\bar{y}_{\cdot j})$ for its mean count, the last two centered across genes.

The fit. The reference fit sets the ceiling at $L = 80$ with a batch size of $1{,}024$ and an epoch cap of $3{,}000$. Every other hyperparameter follows the default described in Appendix A.8. The neighborhood graph is the default variant at $\lambda_g = 100$ with $k_{\mathrm{nn}} = 5$ on 50 principal components. The batch subspace is off, so batch enters only through the fixed dataset dummies in X_i. Every fit runs on one RTX A6000. The ten seeds are the ten the benchmark uses, 123, 321, 456, 654, 789, 987, 555, 777, 888 and 999, of which the first is the reference fit that every read-off in Section 4.2 is taken from. The ceilings $L = 60$, 100 and 120 are each fitted once, at that same first seed, so the sweep varies the ceiling alone.

C.2 Matching factors across fits

Two fits are compared through their loadings, where factor ℓ is the row $\alpha_{\ell\cdot} \in \mathbb{R}^p$ carrying its gene program. We take the active rows of each fit, normalize them to unit ℓ_2 norm across genes, and form the matrix of absolute cosines $C_{k\ell} = |\langle \tilde{\alpha}^{(1)}_{k\cdot}, \tilde{\alpha}^{(2)}_{\ell\cdot} \rangle|$. Rows are then matched one to one by the assignment maximizing $\sum_k C_{k\sigma(k)}$, and we report the mean of C over the matched pairs. The absolute value is what makes the comparison *sign-free*: a factor and its negation are the same axis, since only the product $W\alpha$ enters the likelihood. Where two fits select different numbers of factors, the assignment is rectangular and the surplus rows of the larger set go unmatched. We report the mean over matched pairs throughout, and the floor sits close to it, and no matched pair anywhere in the ceiling sweep falls below $|\cos| = 0.97$.

C.3 Best-factor 1-vs-rest AUROC

Section 4.2 scores how well one factor separates one cell type. For cell type c, let $y_i = \mathbb{1}\{\text{cell } i \text{ has cell type } c\}$, and let $W_{i\ell}$ be cell i's score on active factor ℓ. Writing $A_\ell(c)$ for the area under the ROC curve of the scores $W_{\cdot\ell}$ against y, the reported statistic is

$$\mathrm{AUROC}^{\star}(c) = \max_{\ell \in \widehat{\mathcal{S}}} \max\big\{A_\ell(c),\, 1 - A_\ell(c)\big\}, \tag{23}$$

the best any single factor achieves at either sign. The inner maximum is needed since a factor and its negation are the same axis, and it makes Eq. (23) a maximum over $2L_{\mathrm{active}} = 48$ statistics. No label enters the fit, and the scores are standardized per column beforehand, which leaves the AUROC unchanged since it depends only on ranks.

Over the 61 cell types the statistic has median 0.959 and minimum 0.814, with 52 above 0.90 and 39 above 0.95. Splitting by abundance leaves the median almost unchanged—0.958 for the nine cell types below 500 cells, 0.961 for the nineteen between 500 and $2{,}000$, and 0.959 for the thirty-three above—which is the stratified form of the null rank correlation with abundance that Section 4.2 reports ($\rho = 0.03$).

C.4 Marker sets and marker removal

The marker-removal test deletes a cell type's canonical markers from the gene panel and refits the model from scratch. Marker sets are taken from the papers that described the populations and are listed in Table 6.

Two of the nine canonical ionocyte markers, CFTR and ATP6V1B1, are not in the $3{,}000$-gene panel, since Pearson-residual selection did not retain them. Seven are deleted, and the two absent genes were never available to the fit in the first place. Each refit selects its own factor set, so the comparison is between the best factor of one fit and the best factor of another, not between two factors of the same basis.

Cell type	Markers deleted	AUROC
Neuroendocrine	ASCL1, CALCA, CHGA, GRP, INSM1	$0.958 \to 0.955$
Ionocyte	FOXI1, ASCL3, ATP6V1G3, ATP6V0D2, BSND, CLCNKA, CLCNKB	$0.961 \to 0.955$
Lymphatic EC mature	CCL21, FLT4, PROX1, PDPN, MMRN1	$0.957 \to 0.951$
+ chemokines	*the five above, plus* CSF2, CCL3, CCL3L3, CCL4, CCL4L2, CCL20, CCL22, CXCL1, CCR7, LAMP3, LAIR2, HLA-DQA2	$0.957 \to 0.960$

Table 6: Marker sets, and the best-factor AUROC before and after deletion and refitting.

What the lymphatic arm does and does not show. Lymphatic vessels are wrapped in immune cells that secrete chemokines (Randolph et al., 2017), so deleting those chemokines alongside the canonical markers asks whether the isolating axis survives losing its surroundings as well as its markers. It does, at 0.960. The axis is not, however, best described as a lymphatic one. The same column isolates lymphatic endothelium in all three arms, and mast cells score higher on it throughout, at 0.9896, 0.9897 and 0.9898 against the vessel's 0.957, 0.951 and 0.960. Its surviving top loadings after deletion are mast-cell tryptases and granule genes, TPSAB1, TPSB2, TPSD1, CPA3 and MS4A2, with no lymphatic gene among them. Mast cells are enriched along lymphatic vessels (Randolph et al., 2017), so the factor is best read as a perilymphatic program that both populations load on. The arm therefore establishes that the separation survives deleting the genes the annotation runs on, which is what Section 4.2 claims of it, and not that the axis encodes the vessel rather than its niche.

C.5 How the factors divide the cell types

Figure 2D places each cell type on the factor it is most active on. Activity is read from the score matrix W as follows. Each column is standardized across cells to give $z_{i\ell}$, and averaged within each cell type, $\bar{z}_{c\ell} = n_c^{-1} \sum_{i:\, c_i = c} z_{i\ell}$ over the n_c cells of that cell type. Cell type c is then placed at $\arg\max_\ell |\bar{z}_{c\ell}|$. Signs are fixed once per factor, orienting each so that its largest-magnitude loading is positive, which flips 12 of the 24 and leaves the poles comparable across factors. Factors keep their column index in the $L = 80$ fit, so the 24 selected labels f_ℓ are neither contiguous nor bounded by 24.

Under this placement, 22 of the 24 columns carry at least one cell type, with 17 carrying several and 5 exactly one. The two remaining columns, f_{11} and f_{14}, are the primary axis of no cell type, which is not the same as being idle. Each is a strong discriminator that is merely outranked. Column f_{11} separates pericytes from the rest at an AUROC of 0.976, while f_{29} reaches 0.995, and f_{14} separates bronchial submucosal-gland serous cells at 0.946, while f_{20} reaches 0.994. Neither is technical, at η^2 on cell type of 0.50 and 0.44 and no correlation above 0.39 with any quality covariate. What distinguishes them is that they carry more *within*-cell-type variation than most, 50% and 56% of their variance against a median of 36%, which is the signature of a graded state spanning several cell types rather than of one cell type's identity. A column of that kind is real signal that the argmax rule cannot award to anybody. Ten columns keep every cell type they carry inside one lineage and eleven inside one compartment. Two rules hold exactly. Of those eleven, the six carrying more than one cell type all place them at the same pole, so co-located cell types co-activate a common program rather than being told apart by it. Second, every two-pole column spans two compartments, in all six cases, f_8 setting the fibroblast matrix program against the alveolar capillary endothelium being the clearest. The converse fails: spanning compartments does not make a column a contrast, since five of the eleven cross-compartment columns hold their cell types at one pole.

Those five are where the arrangement is most interesting. Four are read most naturally as shared anatomical niches rather than as cell-type axes. Two are unambiguous. Column f_{12} places systemic venous endothelium beside three dendritic-cell populations and perivascular interstitial macrophages, all at the positive pole, and loads on LYVE1, CD93, STAB1, DNASE1L3 and CD1A. LYVE1$^+$/STAB1$^+$ macrophages sit on the vasculature and dendritic cells traffic through venules, so the column encodes the vascular cuff rather than any one of its residents. Column f_{24} places both serous and the mucous submucosal gland populations with smooth muscle and with ionocytes, loading on the serous exocrine program STATH, FDCSP, MUC7, AZGP1, TCN1 and PIP. Gland

acini are wrapped in smooth muscle, and ionocytes are the airway's ion-transport cell. A third is plausible: f_{18} carries mast cells at $\bar{z} = 4.3$ together with all three lymphatic endothelial states near 1.8, on tryptases and granule genes, which is the perilymphatic program of Appendix C.4. A fourth, f_{21}, puts B cells with mesothelium and mixes the mesothelial PRG4 with the follicular chemokines CXCL13 and CCL19 and with CD19, which would suit a sub-pleural lymphoid aggregate. We report it as speculative.

The fifth is not a niche, and it illustrates the limit of the placement rule. Column f_{29} is carried by pericytes at $\bar{z} = -3.6$, on mural genes, with hillock-like cells and MT-positive alveolar macrophages parked alongside at -1.0 and -1.7. Activity is a magnitude, so a cell type with no strongly dedicated factor lands on whichever column its weak activity happens to maximize. Ten of the 61 placements sit below $|\bar{z}| = 1.5$, and those placements should not be read as programs.

One column deserves separate comment. Column f_0 is where technical variation collects: its scores correlate 0.75 with log sequencing depth and 0.64 with detected-gene count, and both associations strengthen when factor and covariate are centered within cell type, to 0.77 and 0.72, so neither is a by-product of some cell types carrying more RNA. No other factor exceeds 0.40 with depth, at a median of 0.11 and a range of 0.03 to 0.40, so the association is concentrated rather than diffuse. It is not a low-quality-cell axis, since mitochondrial fraction is uncorrelated with it either way. It is not purely technical either: cell type explains $\eta^2 = 0.69$ of its variance, above the median factor, it separates aerocyte capillary endothelium at 0.957, and its loadings are an antimicrobial and IL-17/22 program. We read f_0 as the column where sequencing depth and an inflammatory program collect together, and note that the other twenty-three are comparatively free of the depth association.

C.6 The remaining decoder contrasts

Contrast	n_1	n_2	Donors	Genes	Spearman	Top-50 sign
AT1 vs. AT2	7,937	61,429	51	816	0.925	98%
Alveolar vs. adventitial fibroblasts	5,182	10,321	35	663	0.829	100%
Alveolar vs. monocyte-derived Mϕ	68,486	28,223	62	945	0.891	98%
EC aerocyte vs. EC general capillary	7,157	16,048	30	627	0.804	100%

Table 7: Decoder differential expression against donor-pseudobulk edgeR. "Genes" is the number surviving edgeR's expression filter, over which the rank concordance is computed.

Table 7 gives all four contrasts, of which Figure 2E shows the first. The decoder log-fold-change is the log ratio of the fitted mean between the two groups, taken with the per-cell size-factor offset set to zero so that it measures a library-normalized rate, which is what the pseudobulk anchor measures. The dataset dummies are zeroed as well, leaving the detection-rate covariate in place. The anchor sums raw counts within each donor and group, filters by expression, normalizes by trimmed mean of M-values and tests a paired $\sim$donor $+$ condition design with a negative-binomial generalized linear model (Robinson et al., 2010). Rank concordance is computed over the genes that survive edgeR's own expression filter, between 627 and 945 of the $3{,}000$ in the panel, rather than over the whole panel.

The sign-agreement column reports the 50 genes with the largest absolute edgeR log-fold-change among those edgeR calls significant at a false discovery rate of 0.05. Ranking by effect size rather than by significance is the more conservative of the two choices, since ranking by false discovery rate gives complete sign agreement on every contrast. Fold changes are compressed throughout, the ordinary least-squares slope against edgeR running from 0.55 to 0.93. That follows from the decoder living in the row space of 24 factors, since a gene whose difference lies outside that space cannot carry its full amplitude.

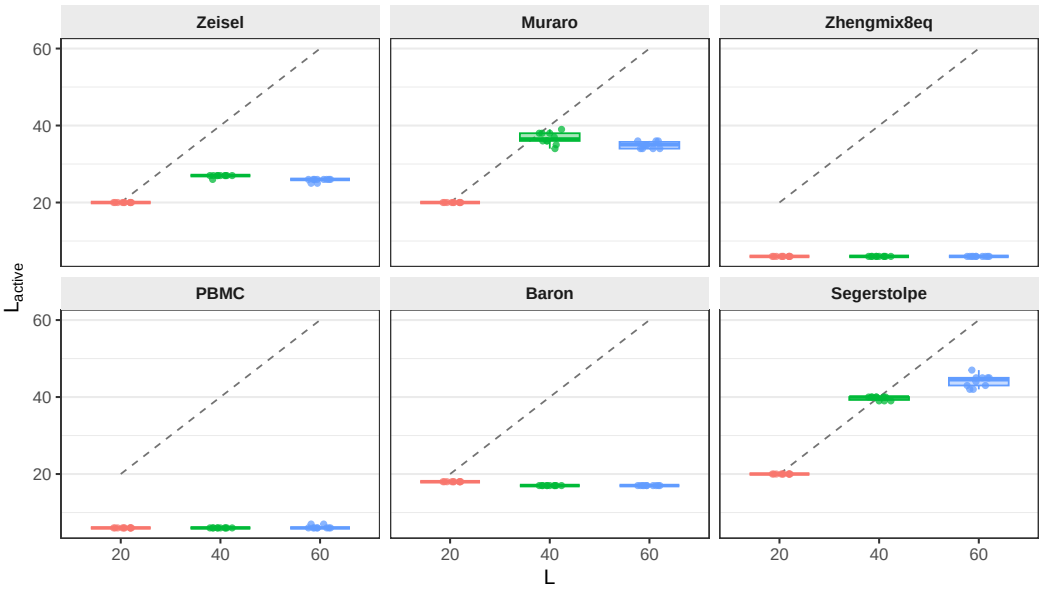


Figure 11: Ablation of the embedding dimension ceiling L. Each boxplot is associated with 10 dots, one per seed.

	L		
	20 (default)	40	60
Zeisel	0.883 (0.003)	0.836 (0.019)	0.828 (0.005)
Muraro	0.936 (0.000)	0.937 (0.001)	0.937 (0.000)
Zhengmix8eq	0.710 (0.002)	0.711 (0.002)	0.711 (0.001)
PBMC	0.721 (0.001)	0.721 (0.000)	0.720 (0.001)
Baron	0.795 (0.051)	0.747 (0.043)	0.781 (0.056)
Segerstolpe	0.898 (0.091)	0.883 (0.054)	0.886 (0.052)

Table 8: ARI across $L \in \{20, 40, 60\}$ for all 6 datasets, 10 seeds each. Values in the table are mean ARI with standard deviation in parentheses.

D Ablations

D.1 Latent dimension ceiling

We refit all six datasets in Section 4.1 at ceiling $L \in \{20, 40, 60\}$, ten seeds each, and ask two questions: whether L_{active} settles below the ceiling L as the ceiling rises, and whether over-specifying it costs anything downstream. Comparisons are paired Wilcoxon tests over the ten shared seeds.

The selected count stops tracking the ceiling once the ceiling has slack. As shown in Figure 11, L_{active} is nearly constant across the range of L on Zhengmix8eq, PBMC, and Baron. In particular, $L_{\text{active}} = 6$ on Zhengmix8eq at every L across all 10 seeds. On PBMC, $L_{\text{active}} = 6$ on every seed at $L = 20, 40$, with two seeds rising to 7 at $L = 60$. For Baron, the average L_{active} ranges from 17 to 18 as L varies, but never tracks the ceiling L. At the default $L = 20$, L_{active} for Zeisel is pinned at 20 on all ten seeds. When the ceiling L is raised, average L_{active} ranges from 25.8 to 26.9, so L_{active} is not tracking the ceiling once given room. On Muraro, a similar pattern holds: L_{active} is 20 at $L = 20$, then rises to 36.7 at $L = 40$ and 34.9 at $L = 60$. Segerstolpe is the only dataset where L_{active} tracks the 2 lower settings of L, with average $L_{\text{active}} = 20$ for $L = 20$ and 39.7 for $L = 40$. When L is raised to 60, L_{active} stops tracking the ceiling, with an average $L_{\text{active}} = 44.1$.

Over-specifying is harmless on five of six. Table 8 shows that oracle-k ARI is unchanged from $L = 20$ to $L = 60$ on Muraro, Zhengmix8eq, PBMC and Segerstolpe ($|\Delta\text{ARI}| \leq 0.015$). On Baron, ARI dips at $L = 40$ then recovers at $L = 60$, but the differences are not significant (Wilcoxon $p = 0.16$ for $L = 20$ vs $L = 40$, and $p = 0.49$ for $L = 20$ vs $L = 60$). The dip at $L = 40$ is due to the k-means bimodality on Baron, where it reaches 1/10 good-basin seed compared to 5/10 and 3/10 for $L = 20$ and $L = 60$, respectively.

The Zeisel exception, and it is not a loss of signal. On Zeisel over-specifying L to 40 or 60 results in a loss of ARI. We analyzed two candidate explanations for this loss, but both are ruled out. Unlike Baron's, this loss is not due to an optimization bimodality since Zeisel's ARI is unimodal at every L. The decrease in ARI is also not due to the extra latent dimensions when L is raised. When constraining the k-means to the top d latent factors ranked by their SNR, ARI is flat between 0.807 and 0.833 for every d from 10 to 26, so more selected factors do not harm the k-means ARI.

The loss is not in the embedding but in the clusterer's read of it. On both clusterer-independent metrics the $L = 60$ embedding is better on 10/10 seeds: label-ASW rises from 0.657 to 0.665 and cLISI from 0.9998 to 0.9999, each $p = 0.002$. Zeisel's embedding therefore separates the cell types at least as well at the larger ceiling, and it is k-means at oracle k that finds a worse partition of it. Theorem 2(c) shows that the fitted signal beyond the true rank is weightless, which label-ASW and cLISI confirm; no claim is made about a particular clusterer applied afterwards.

D.2 Neighborhood regularizer strength

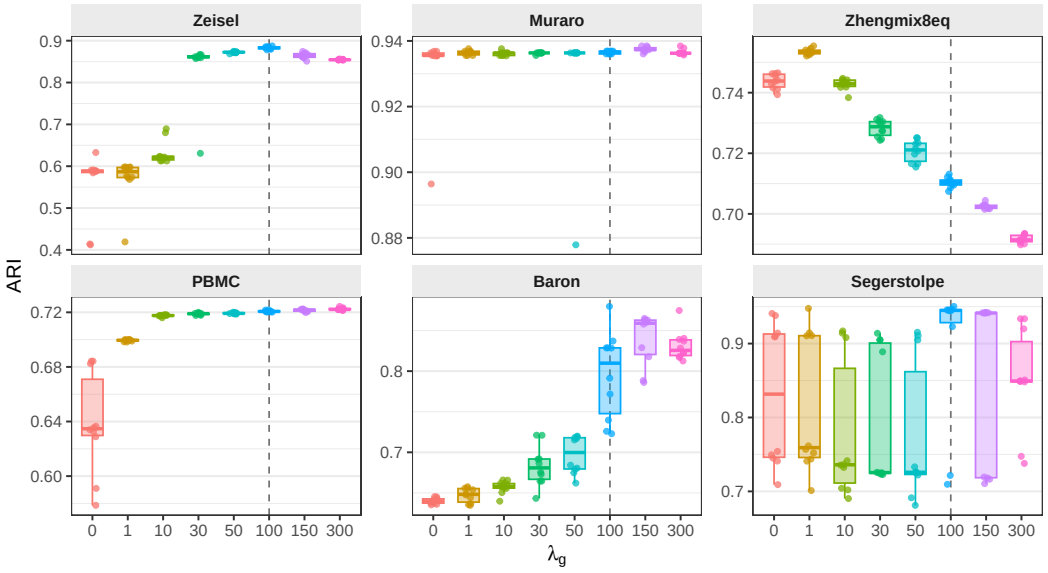


Figure 12: Ablation of the neighborhood regularizer strength λ_g. Each boxplot is associated with 10 dots, one per seed. The vertical dashed line indicates the default $\lambda_g = 100$.

We refit all six datasets at $\lambda_g \in \{0, 1, 10, 30, 50, 100, 150, 300\}$, ten seeds each, holding every other setting at its shipped value. Note that $\lambda_g = 0$ disables the graph regularizer in Eq. (4). Figure 12 shows the per-dataset dose–response.

The graph helps in proportion to how weakly the model does without it. Turning it off ($\lambda_g = 0$) costs the three weakest datasets the most. On average, Zeisel falls from 0.883 to 0.558, Baron from 0.795 to 0.640, and PBMC from 0.721 to 0.639. In contrast, RIPPLE already performs well on Muraro and Segerstolpe without the graph regularizer, so there is far less for λ_g to recover: turning it off costs Baron 0.155 ARI but Muraro only 0.005, whose mean ARI moves within 0.007 across the entire grid, and on Segerstolpe the $\lambda_g = 0$ versus $\lambda_g = 100$ contrast is not significant despite a mean that moves from 0.83 to 0.90—the spread there is the k-means bimodality, not the regularizer.

On the other hand, ARI is monotonically decreasing once $\lambda_g > 1$ on Zhengmix8eq, where the graph regularizer only helps when the regularization is weak ($\lambda_g = 1$). This is expected since Zhengmix8eq is a FACS-purified mixture whose populations are separated by construction, so the graph regularizer is not needed to separate them.

On the two bimodal datasets, λ_g moves the probability of the good basin rather than its level. On Baron, the better mode (ARI ≈ 0.8) does not exist below the default $\lambda_g = 100$. As λ_g increases past 100, RIPPLE's ARI falls in the good basin more often, with $\lambda_g = 300$ being the most consistent. Thus, a stronger graph regularizer is needed for Baron, with $\lambda_g = 100$ being the tipping point for more consistent strong performance. On Segerstolpe, the large seed-to-seed variance is again due to the bimodality of the k-means clustering. With the same medoid restart fix described in Appendix B.10, the medoid-selected ARI is stable at $\lambda_g = 100$ (0.946) and $\lambda_g = 150$ (0.942), and RIPPLE reaches ARI > 0.9 more often than smaller λ_g.

D.3 Cell and gene covariates

In this ablation, we investigate the effect of covariates on RIPPLE's performance. We refit all six datasets over the ten seeds along the three arms detailed in Table 9. In particular, all hyperparameters are kept constant, and the ladder moves the two design matrices (X and V) and the offset (O) only.

Arm	Mean structure	What it prices
Base	$\log \mu = X\beta + \gamma V^\top + W\alpha + O$	—
Offset only	$\log \mu = W\alpha + O$	the two designs
No covariates	$\log \mu = W\alpha$	the depth offset as well

Table 9: The covariate ablation ladder. Each arm removes one more term from the mean, so the step between two arms prices the term that separates them. All three are fitted over the ten canonical seeds on all six datasets; the base arm is the benchmark run itself.

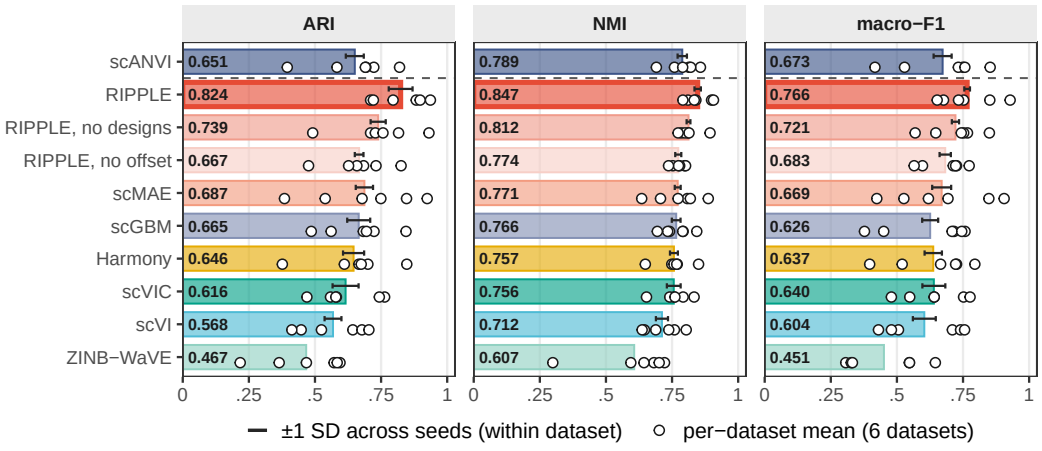


Figure 13: The benchmark figure of Section 4.1 with the two ablation arms added, read exactly as Figure 1 is.

Stripping the designs costs little, and the lead survives it. Dropping both designs but keeping the depth offset costs 0.085 ARI averaged over the six datasets, from 0.824 to 0.739, which still outperforms every unsupervised baseline (next best 0.687) and the supervised scANVI (0.651). On NMI the same arm scores 0.812 against 0.771 for the best unsupervised baseline and 0.789 for scANVI. The cost is concentrated rather than uniform. Segerstolpe loses 0.406 ARI, Zeisel 0.067, Baron 0.038 and Muraro 0.006, while Zhengmix8eq is unchanged (-0.0002, $p = 1.00$) and PBMC improves slightly ($+0.008$). Segerstolpe is the dataset whose ten donors are confounded with disease status, so it is the one where a cell-covariate design has the most to remove. Figure 13 places the arm

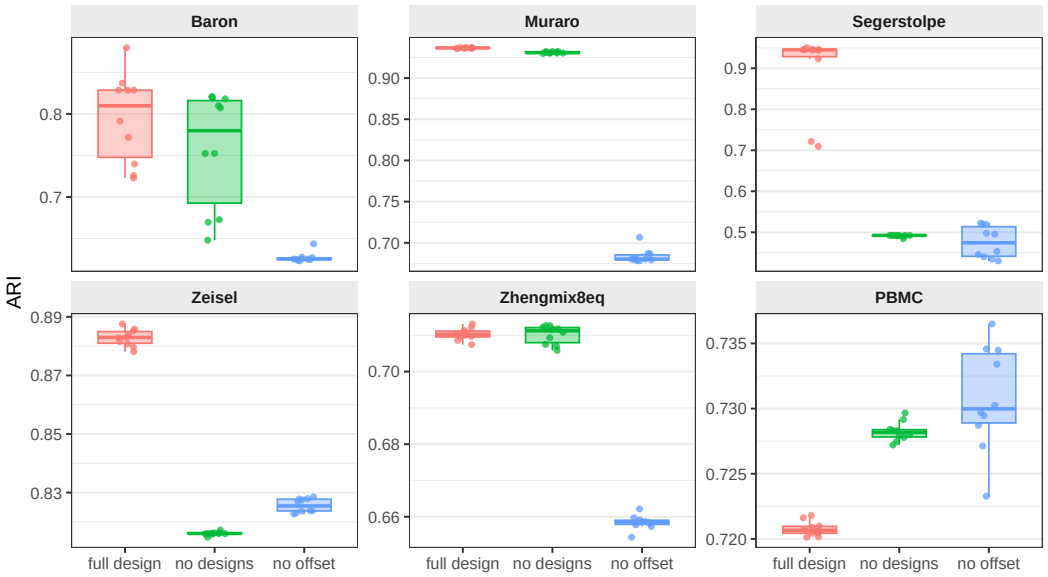


Figure 14: Oracle-k ARI along the ladder of Table 9, per dataset. Boxes are the ten canonical seeds with the individual seeds overlaid.

against the benchmark: with both designs gone RIPPLE still outperforms every baseline on all three metrics, and the margin it keeps over the strongest of them, 0.052 ARI, is wider than the margin scMAE keeps over scGBM.

The depth offset is the load-bearing term, and what it costs is a partition, not a representation. Removing it as well costs a further 0.247 ARI on Muraro and 0.130 on Baron, and the mechanism is visible in the embedding: the largest absolute correlation between a latent coordinate and log library size rises from 0.40 to 0.95 on Muraro and from 0.36 to 0.83 on Baron, so a factor is spent on sequencing depth once nothing else absorbs it. Restoring the offset puts it back, and the count moves with it: L_{active} rises from 6 to 8.9 on PBMC and from 6 to 8 on Zhengmix8eq once the offset is gone, the two datasets with the most slack under the ceiling.

What that costs depends on which metric is asked. Averaged over the six datasets the offset step takes 0.072 from ARI but only 0.039 from NMI and 0.039 from macro-F_1, so ARI falls about twice as far as either. The ranking in Figure 13 follows: with the offset gone as well, RIPPLE still outperforms every unsupervised baseline on NMI (0.774 against 0.771) and on macro-F_1 (0.683 against 0.669, and above the supervised scANVI's 0.673), while on ARI it trails scMAE by 0.021 and beats the other five.

We read that gap as the signature of a partition that has been refined rather than scrambled. A factor tracking library size splits a cell type into a deeply and a shallowly sequenced half; the split breaks a large number of within-type pairs, which ARI counts quadratically, while each half still maps to one true cell type, so the mutual information is largely preserved and Hungarian matching gives macro-F_1 back the larger half. The damage concentrates where depth capture is strongest, on Muraro (-0.247 ARI against -0.093 NMI) and Baron (-0.130 against -0.040), and the per-dataset spread in Figure 14 is where to read it. What the offset protects, on this reading, is not the representation but the clustering of it.

The covariate-free arm is also where the theory is simplest. With intercept-only designs $X = \mathbf{1}_n$, $V = \mathbf{1}_p$ the projections reduce to row and column centering and every statement reads for the doubly centered product. The arm fitted here drops the intercepts as well; there Assumption 3 of Appendix E fails, the first part of Theorem 1 transfers unchanged with $\widehat{M}^{\perp} = \widehat{W}\widehat{\alpha}$ and $M^{0\perp} = W^0\alpha^0$, and the selection statements are not asserted. One definition changes: the strengths d_k are read off $\eta^0 - O\mathbf{1}_p^{\top}$

rather than η^0, the offset no longer being removed by V's intercept column. The design-residual frame of §E.1.2 is the price of the two designs, and the ladder of Table 9 is what that price buys.

D.4 BATCH CORRECTION

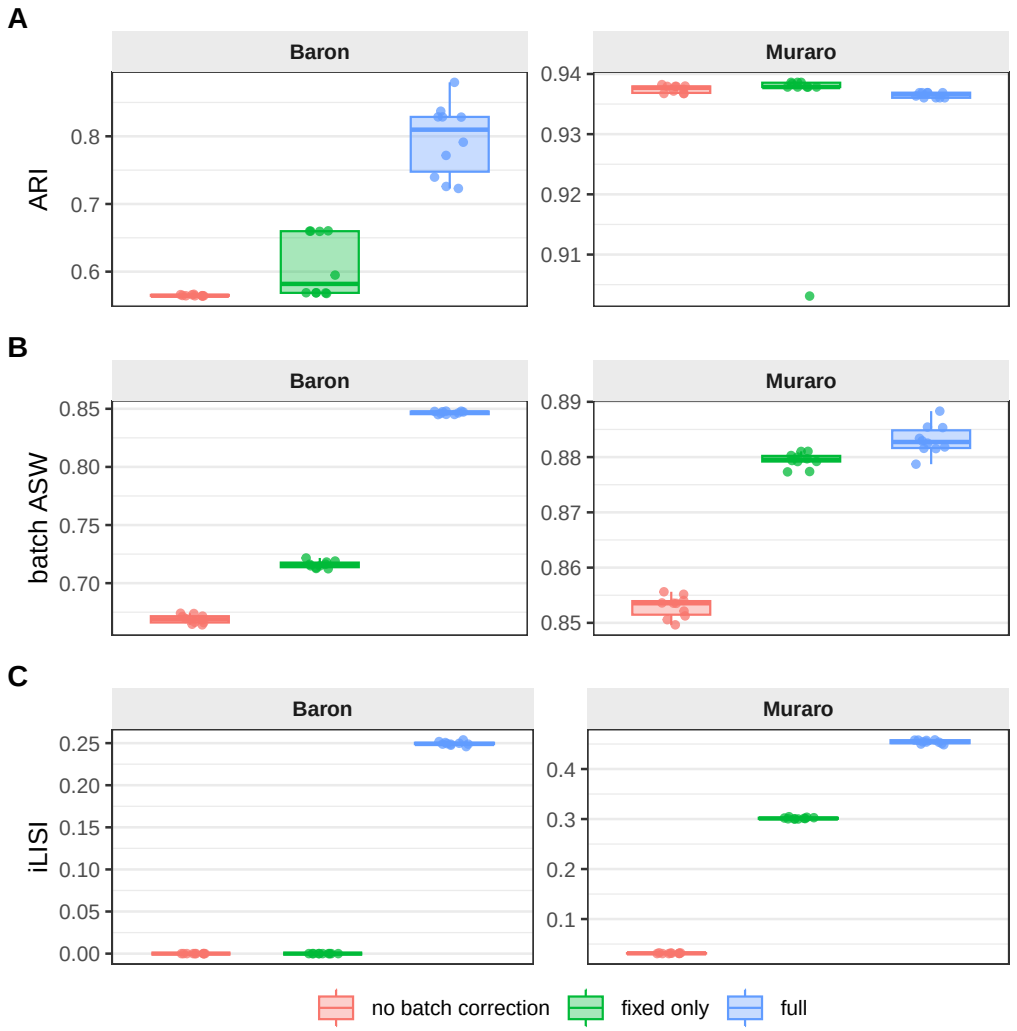


Figure 15: Batch correction ablation. Each boxplot is associated with 10 dots, one per seed. Full corresponds to the batch integration described in Section 2 with cross-batch graph, random batch subspace, and fixed batch dummies. Fixed only uses the batch dummies alone, and no batch correction means RIPPLE does not consume any batch information.

Correction buys batch mixing, and costs biology only where there is none left to gain. We fit RIPPLE on two datasets with known batch effects, Baron and Muraro, and investigate the effect of the batch correction components described in Section 2. Here we ablate the batch correction components with successive increment of integration power: (i) no batch correction, (ii) fixed batch dummies only, and (iii) full batch correction with cross-batch graph, random batch subspace, and fixed batch dummies. Figure 15 shows two metrics: ARI representing the biological signal preservation and batch-ASW representing the batch mixing.

The evidence is clear on batch-ASW (Figure 15B). Each increment of batch correction improves the mixing, and the full arm is best on both datasets. On Baron the largest gain comes from fixed-only to full, lifting the average from 0.716 to 0.847. On Muraro it comes earlier, from no correction to fixed-only, at 0.853 to 0.880. Muraro starts better mixed than Baron (0.853 against 0.669), so the cross-batch graph and random subspace have less left to correct.

On ARI (Figure 15A) the two datasets differ. Baron improves from 0.565 to 0.608 with fixed dummies and then to 0.795 with full correction, the largest step again the last one. Muraro is already at 0.937 without correction, dips to 0.935 with dummies, where one seed lands in a worse basin, and returns to 0.936 with the full arm. Where there is no biological signal left to recover, the full arm buys batch mixing at a small cost in ARI.

E THEORY

This appendix proves the results of Section 3. The per-cell latents grow in number with n, an incidental-parameters problem (Neyman & Scott, 1948), and every argument below treats the global parameters and the per-cell latents as jointly estimated from the same counts; nothing is conditioned on oracle knowledge of either. Statements from the main text are restated (unnumbered) where they are proved. Throughout, $C, c > 0$ are constants depending only on the fixed parameter ranges and on (L, H, G, σ_β), changing from line to line. The graph weight λ_g is fixed; §E.4 displays it explicitly and never absorbs it, and §E.5 onward absorbs it into C where stated.

E.1 SETTING

E.1.1 MODEL AND ESTIMATOR

Counts $Y \in \mathbb{Z}_{\geq 0}^{n \times p}$ over n cells and p genes, independent entries $Y_{ij} \sim \mathrm{NB}(\mu_{ij}^0, \theta_j^0)$ in the mean–dispersion parametrization, with

$$\log \mu_{ij}^0 = \eta_{ij}^0 = X_{i\cdot}^\top \beta_{\cdot j}^0 + \gamma_{i\cdot}^{0\top} V_{j\cdot} + W_{i\cdot}^{0\top} \alpha_{\cdot j}^0 + O_i, \qquad M^0 := W^0 \alpha^0.$$

X, V, O are known; L is the over-provisioned ceiling. The estimator is the one the paper deploys: mean-field variational Bayes with free Gaussian factors per scalar latent ($W, \alpha, \beta, \xi = \log\theta$) and point estimates ($\lambda, \gamma$), maximizing the penalized evidence lower bound

$$\mathcal{L}(\phi) = \sum_{i,j} \mathbb{E}_q\big[\log \mathrm{NB}(Y_{ij} \mid e^{\eta_{ij}}, e^{\xi_j})\big] - \mathrm{KL}(\phi) - \lambda_g \|(I - A) m_W\|_F^2,$$

where $\mathrm{KL}(\phi)$ collects the closed-form Kullback–Leibler terms of Eq. (3) (with the half-Cauchy log-prior of the point-estimated scales), and A is the row-normalized adjacency of the kNN graph, built once from the data and then fixed. The optimization runs over a fixed bounded feasible set Φ, compact once the variance ranges are closed at zero, where only the collapsed rows are attained (Lemma 4); the properties of Φ that the proofs consume are these. Entrywise boxes on all variational means and point estimates: $|m_{W_{i\ell}}| \leq C_W$, $|m_{\alpha_{\ell j}}| \leq C_\alpha$, $|m_{\beta_{hj}}| \leq C_\beta$, and analogous boxes for γ and m_ξ. The scale box $\lambda_\ell \in [0, \lambda_{\max}]$ with $\lambda_{\max}^2 \geq C_\alpha^2 + s_{\max}^2$, so that it contains $\widehat{R}_\ell / p$ for every feasible loading row, and in particular $C_\alpha^2 + 1$. Variance boxes with no positive lower bound and, for the W, α and β blocks, with caps at least 1: the score variances $s_{W_{i\ell}}^2$ range over $(0, s_{W,\max}^2]$ with $s_{W,\max}^2 \geq 1$, so that the value 1 of a collapsed column is admissible, and the loading and coefficient variances $s_{\alpha_{\ell j}}^2$, $s_{\beta_{hj}}^2$ over $(0, s_{\max}^2]$ with $s_{\max}^2 \geq 1$, so that $(n+p)^{-1}$ is admissible for every n, p; the dispersion mixing variances $s_{\xi_j}^2$ are capped by a small constant, $(0, 1/16]$ suffices. A box $|\bar{\eta}_{ij}| \leq C_\eta$ on the mean predictor $\bar{\eta}_{ij} = X_{i\cdot}^\top m_{\beta\cdot j} + \gamma_{i\cdot}^\top V_{j\cdot} + m_{W_{i\cdot}}^\top m_{\alpha\cdot j} + O_i$. And the closure at *collapsed rows* $(\lambda_\ell, m_{\alpha_{\ell\cdot}}, s_{\alpha_{\ell\cdot}}^2) = (0, 0, 0)$, on which the variational law of $\alpha_{\ell\cdot}$ is the point mass at zero (so $\alpha_{\ell\cdot} = 0$ there) and the row-KL has value 0 — the configuration by which the sparsity prior switches a factor off, which must be admissible for the maximum to be attained (Lemma 4). We write $\widehat{\phi}$ for the fitted configuration, $\widehat{W} = m_{\widehat{W}}$, $\widehat{\alpha} = m_{\widehat{\alpha}}$, $\widehat{\eta} = \bar{\eta}(\widehat{\phi})$, $\widehat{M} = \widehat{W}\widehat{\alpha}$, and $\widehat{\nu} = (\widehat{m}_{\xi_j}, \widehat{s}_{\xi_j})_{j \leq p}$ for the fitted dispersion block.

The estimator, and the two consequences of maximality that are used. The configuration analyzed is a maximizer, $\widehat{\phi} \in \arg\max_{\Phi} \mathcal{L}$ (Assumption 4), which exists by Lemma 4; the variational expectations are taken exactly (no Monte Carlo). Exactly two consequences of maximality are used. *(i) The basic inequality against one explicit comparison configuration.* Let $(\beta^0, \gamma^0, W^0, \alpha^0)$ be a parametrization of the truth inside the boxes (Assumption 2) and $\mathcal{A} := \{\ell : \alpha^0_{\ell\cdot} \neq 0\}$ its live rows. Given the fitted dispersion block $\widehat{\nu}$, let $\widetilde{\phi}(\widehat{\nu})$ have $m_\beta = \beta^0$, $\gamma = \gamma^0$, $m_{W\cdot\ell} = W^0_{\cdot\ell}$ and $m_{\alpha\ell\cdot} = \alpha^0_{\ell\cdot}$ for $\ell \in \mathcal{A}$, so that $\bar{\eta}(\widetilde{\phi}(\widehat{\nu})) = \eta^0$; every variational variance of the blocks β, $W_{\cdot\mathcal{A}}$ and $\alpha_{\mathcal{A}\cdot}$ equal to $(n+p)^{-1}$; $\widetilde{\lambda}^2_\ell := \widetilde{R}_\ell/p$ with $\widetilde{R}_\ell := \|\alpha^0_{\ell\cdot}\|^2 + p(n+p)^{-1}$ for $\ell \in \mathcal{A}$; for $\ell \notin \mathcal{A}$ the row collapsed, with $m_{W\cdot\ell} = 0$ and $s^2_{W\cdot\ell} = 1$; and the dispersion block copied from the fit, $q(\xi_j) = \mathcal{N}(\widehat{m}_{\xi_j}, \widehat{s}^2_{\xi_j})$ for every j. It lies in Φ by the definition above: the means are in the boxes and $|\bar{\eta}(\widetilde{\phi}(\widehat{\nu}))| = |\eta^0| \leq C_\eta$ (Assumption 2), $(n+p)^{-1}$ lies in every W, α and β variance box because these have no positive lower bound, $\widetilde{\lambda}^2_\ell \leq C^2_\alpha + (n+p)^{-1} \leq C^2_\alpha + 1 \leq \lambda^2_{\max}$, the collapsed rows and the value $s^2_W = 1$ of their columns are admissible, and $\widehat{\nu}$ is feasible because $\widehat{\phi}$ is. Hence

$$\mathcal{L}(\widehat{\phi}) \;\geq\; \mathcal{L}\big(\widetilde{\phi}(\widehat{\nu})\big).$$

The first part of Theorem 1 uses this inequality and nothing else about $\widehat{\phi}$, and the second part adds (ii); the comparison configuration depends on the fit through its dispersion block only. *(ii) Stationarity in the score variances*: $\partial\mathcal{L}/\partial s^2_{W_{i\ell}} = 0$ at every (i,ℓ). The maximizer is interior in this block, the collapsed rows being the stationary boundary $s^2_{W_{i\ell}} = 1$ (Remark 1). Lemma 5(b), the second part of Theorem 1, Theorem 2 and the selected-part clause of Proposition 2 use (ii) through the identity it yields for the denominator of the selection statistic. Both consequences are checkable on a fit: (i) by evaluating $\mathcal{L}$ at $\widetilde{\phi}(\widehat{\nu})$, (ii) by reading the gradient in $\log s_W$.

Why (i) is worth naming. The deployed optimizer stops the gradient through the neighbor mean (Eq. (4)), so the trained point is stationary for a field carrying the asymmetric operator $(I - A)$ in place of the symmetric $(I - A)^\top(I - A)$ of the analyzed penalty; a linear field is a gradient field only if its matrix is symmetric, so the trained point is not in general a stationary point of $\mathcal{L}$. The two penalty forms agree in value, so $\mathcal{L}$ and its maximizer are well defined, and what Assumption 4 asserts of the trained point is that it is at least as good as $\widetilde{\phi}(\widehat{\nu})$. The stop-gradient acts on the score means alone and never reaches the variance block, whose gradient is the ordinary one, so (ii) is an ordinary first-order condition for the trained point too. Annealing schedules and numerical guards are outside the scope of the analysis, which concerns the mathematical estimator the paper defines.

The selection rule is Eq. (6): with $\widehat{V}_\ell$ the empirical variance across cells of the fitted score means and $\frac{1}{n}\sum_i s^2_{W_{i\ell}}$ the average fitted variance, $\mathrm{SNR}_\ell = \widehat{V}_\ell \big/ \frac{1}{n}\sum_i s^2_{W_{i\ell}}$ and $\widehat{\mathcal{S}} = \{\ell : \mathrm{SNR}_\ell > 1\}$, $L_{\mathrm{active}} = |\widehat{\mathcal{S}}|$.

E.1.2 Centered quantities and the estimand

Let $P^\perp_X$ and $P^\perp_V$ project off the column spans of X and V. For a matrix $Z \in \mathbb{R}^{n\times p}$, the *design-residual spectrum* of Z is

$$d_k(Z) := \min_{\beta,\gamma} \sigma_k\big(Z - X\beta - \gamma V^\top\big), \qquad \mathrm{rank}_{X,V}(Z) := \#\{k : d_k(Z) > 0\},$$

where σ_k denotes the k-th singular value. Then $d_k(Z)$ is the strength of the k-th direction of variation in Z that no assignment of covariate effects can produce, and $\mathrm{rank}_{X,V}(Z)$ counts those directions.

Lemma 1 (design-residual spectrum)**.** *For every Z and every k, $d_k(Z) = \sigma_k(P^\perp_X Z P^\perp_V)$, and the minimum is attained simultaneously in k at the single argument $Z - P^\perp_X Z P^\perp_V$; hence* $\mathrm{rank}_{X,V}(Z) = \mathrm{rank}(P^\perp_X Z P^\perp_V)$.

Proof. Projections contract, so for every $S = X\beta + \gamma V^\top$ and every k, $\sigma_k(P^\perp_X Z P^\perp_V) = \sigma_k(P^\perp_X(Z - S)P^\perp_V) \leq \sigma_k(Z - S)$, which bounds the minimum below. The bound is attained inside the class, since $Z - P^\perp_X Z P^\perp_V = P_X Z + P^\perp_X Z P_V$ has a first term with columns in $\mathrm{col}(X)$ and a second with rows in $\mathrm{col}(V)^\top$, and is therefore of the form $X\beta + \gamma V^\top$. □

The truth's quantities are read off this spectrum. Write $M^{0\perp} := P_X^\perp M^0 P_V^\perp$ for the truth's design-residual signal and $\widehat{M}^\perp := P_X^\perp \widehat{M} P_V^\perp$ for the fit's, and let $M^{0\perp} = \sum_k d_k u_k v_k^\top$ be the former's singular value decomposition, $d_1 \geq d_2 \geq \dots$. Since $X\beta^0$ dies on the left, $\gamma^0 V^\top$ dies on the right, and the offset $O\mathbf{1}_p^\top$ dies on the right through V's intercept column,

$$M^{0\perp} = P_X^\perp \, \eta^0 \, P_V^\perp, \qquad d_k := d_k(\eta^0), \qquad L^0 := \operatorname{rank}_{X,V}(\eta^0) \leq L,$$

so each of these is a functional of η^0, and hence of the conditional law of Y given the designs; the last two identifications are Lemma 1. We write d_k for $d_k(\eta^0)$ throughout; the functional $d_k(\cdot)$ keeps its argument. The unprojected M^0 is not: the substitution $W \mapsto W + XC$, $\beta \mapsto \beta - C\alpha$ leaves η^0 — and therefore that law — exactly unchanged while moving M^0 by $XC\alpha^0 \neq 0$. All statements below are conditional on the realized (X, V), which §E.1 takes as known; when these are built from Y, as in our experiments, the estimand is likewise (X, V)-measurable (the theorems, however, take the designs as fixed; Assumption 3).

Separately, let $P_1^\perp = I_n - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$ denote row centering and $\widetilde{W}_{\cdot\ell} = P_1^\perp m_{\widehat{W}_{\cdot\ell}}$, so that $\widehat{V}_\ell = \|\widetilde{W}_{\cdot\ell}\|_F^2/n$ is the deployed selection statistic of Eq. (6), which we leave in its own gauge.

Partialling the design out is not a convention we impose on the method but one the method already carries, on a single column: the selection statistic is an empirical *variance*, and k-means is invariant to translating a coordinate, so score means are pure gauge — they are absorbed by the intercept path of the design. Since $\mathbf{1}_n$ is a column of X and $\mathbf{1}_p$ a column of V, $P_1^\perp$ *is* $P_X^\perp$ for $X = [\mathbf{1}]$; what follows runs the same argument on all $H + G$ design columns instead of one.

The *canonical factors* of the estimand are the score vectors $\sqrt{n}\, u_k$ (unit empirical variance per factor) with loadings $(d_k/\sqrt{n})\, v_k^\top$, $k \leq L^0$; d_k^2 is factor k's *strength*. They are functionals of $M^{0\perp}$, not columns of the generative W^0, which Proposition 1 shows is not identifiable.

Subspaces and distances. Write $U^0 := [u_1, \dots, u_{L^0}]$, $D^0 := \operatorname{diag}(d_1, \dots, d_{L^0})$ and $V^0 := [v_1, \dots, v_{L^0}]$, so that $M^{0\perp} = U^0 D^0 (V^0)^\top$; let $\mathcal{U}^0 := \operatorname{col}(M^{0\perp})$, with projector $P_{\mathcal{U}^0} = U^0 (U^0)^\top$, and let

$$\widehat{\mathcal{U}} := \operatorname{span}\left\{ P_X^\perp m_{\widehat{W}_{\cdot\ell}} : \ell \in \widehat{\mathcal{S}} \right\}, \qquad \widehat{M}_{\widehat{\mathcal{S}}}^\perp := P_X^\perp \, \widehat{W}_{\cdot\widehat{\mathcal{S}}} \, \widehat{\alpha}_{\widehat{\mathcal{S}}\cdot} \, P_V^\perp = \sum_{\ell \in \widehat{\mathcal{S}}} \left(P_X^\perp m_{\widehat{W}_{\cdot\ell}} \right) \left(\widehat{\alpha}_{\ell\cdot} P_V^\perp \right),$$

with $\widetilde{P}$ the orthogonal projector onto $\widehat{\mathcal{U}}$, so $\dim \widehat{\mathcal{U}} \leq L_{\text{active}}$. For a unit vector u and a subspace $\mathcal{V}$, $\operatorname{dist}(u, \mathcal{V}) := \|u - \Pi_{\mathcal{V}} u\|$; for $u, w \in \mathbb{R}^n$ with nonzero centered parts, $\operatorname{corr}(u, w) := \langle P_1^\perp u, P_1^\perp w \rangle / (\|P_1^\perp u\| \|P_1^\perp w\|)$ is the empirical correlation across cells. Two facts are used throughout. Because $\mathbf{1}_n \in \operatorname{col}(X)$, $P_X^\perp = P_X^\perp P_1^\perp$, so $P_X^\perp m_{\widehat{W}_{\cdot\ell}} = P_X^\perp \widetilde{W}_{\cdot\ell}$; and $u_k \in \operatorname{col}(P_X^\perp)$ gives $P_X^\perp u_k = u_k = P_1^\perp u_k$, whence

$$\operatorname{corr}\left(u_k, \widehat{W}_{\cdot\ell}\right) = \frac{u_k^\top \widetilde{W}_{\cdot\ell}}{\|\widetilde{W}_{\cdot\ell}\|},$$

the cosine of the angle between the true score direction and the centered fitted column.

E.2 Assumptions

Assumption 1 (model and ceiling)**.** Y has independent entries $Y_{ij} \sim \operatorname{NB}(e^{\eta_{ij}^0}, \theta_j^0)$ with $\eta^0 = X\beta^0 + \gamma^0 V^\top + W^0 \alpha^0 + O\mathbf{1}_p^\top$ as in the model above, generated with at most L factors, and $L^0 = \operatorname{rank}(P_X^\perp \eta^0 P_V^\perp)$ satisfies $1 \leq L^0 \leq L$.

Assumption 2 (boundedness and asymptotics)**.** Some parametrization $(\beta^0, \gamma^0, W^0, \alpha^0, \xi^0)$ of the truth lies in the boxes defining Φ, in particular $|\eta_{ij}^0| \leq C_\eta$ and $\theta_j^0 \in [\theta_{\min}, \theta_{\max}]$; and $n, p \to \infty$ jointly, at any aspect ratio, with $r \log(np) = o(n \wedge p)$, where $r := L + H + G$ is fixed.

Assumption 3 (designs)**.** $X \in \mathbb{R}^{n \times H}$ contains $\mathbf{1}_n$ and $V \in \mathbb{R}^{p \times G}$ contains $\mathbf{1}_p$; both are fixed (non-random), of full column rank, with H, G fixed and bounded rows, $\max_i \|X_{i\cdot}\| \leq C$ and $\max_j \|V_{j\cdot}\| \leq C$. The gene-covariate loadings are taken in the canonical representative $X^\top \gamma = 0$, where $\gamma \in \mathbb{R}^{n \times G}$ stacks the $\gamma_{i\cdot}^\top$ as rows.

Assumption 4 (estimator)**.** $\widehat{\phi}$ maximizes $\mathcal{L}$, evaluated with exact variational expectations, over Φ. The proofs use two consequences only:

(i) $\mathcal{L}(\widehat{\phi}) \geq \mathcal{L}(\widetilde{\phi}(\widehat{\nu}))$ for the comparison configuration of §E.1;

(ii) $\partial\mathcal{L}/\partial s^2_{W_{i\ell}} = 0$ for every (i, ℓ).

Assumption 5 (strength of the weakest true factor)**.** $d^2_{L^0}/\omega_{n,p} \to \infty$, where $\omega_{n,p} := r\,(n + p)\log(np)$ is the rate of Theorem 1. Its pervasive form, written Assumption 5^+ and used by Theorem 2(d) only: $d^2_{L^0} \geq c_h\, np$ for a fixed $c_h > 0$ and all large n, p. Since $\omega_{n,p} = o(np)$ under Assumptions 2, 5^+ implies Assumption 5.

Assumption 6 (gap; Proposition 2 only)**.** For every $k \leq L^0$, $\min(d_{k-1} - d_k,\ d_k - d_{k+1}) \geq \delta\, d_k$ for a fixed $\delta \in (0, 1]$, with $d_0 = \infty$ and $d_{L^0+1} = 0$.

Discussion of the assumptions. Assumption 1 is load-bearing: the guarantees are about the model the paper fits; its clause $1 \leq L^0 \leq L$ is the analogue of $1 \leq K^* \leq K_{\max}$ in Chen & Li (2022). Assumption 2 is regularity, bounded log-means and dispersions, which is what the data look like after standard gene filtering; the growth condition has one role, to make $\omega_{n,p} = o(np)$ (it also closes the net step of the empirical-process lemma behind Theorem 1). Assumption 3 asks for no conditioning of the design Grams, which no proof uses; the proofs use the bounded rows (in the net step, where moving a coefficient by ϵ must move every predictor entry by $O(\epsilon)$) and full column rank (for the uniqueness clause of Lemma 7 only). No separation between the factor signal and the covariate spaces is assumed, and none is available: a condition $\|P_X^{\perp} M\|_F^2 \geq c_X \|M\|_F^2$ over the feasible rank-$2L$ class fails for every design, since $M = ab^\top$ with $a \in \mathrm{col}(X)$ lies in the class with $P_X^{\perp} M = 0$; the statements are made for the design-residual signal, which needs no such condition. The designs are fixed: independence of the entries of Y given (X, V) is what every concentration step uses, so a covariate computed from Y, such as the deployed detection rate, is outside the theorems as stated (a plug-in argument, using that the detection rate of a cell concentrates at rate $p^{-1/2}$ uniformly over cells up to a $\sqrt{\log n}$ factor, is the natural extension and is not made here). Assumption 4 is the standard idealization of variational-inference theory, which analyzes the exact optimizer over the family (Wang & Blei, 2019; Zhang & Gao, 2020), and of the joint-likelihood factor literature, which analyzes the constrained maximizer (Chen & Li, 2022; Chen et al., 2019); its consequence (ii) is the price of a variational denominator in the selection statistic, which those analyses do not have, and exact expectations idealize away the Monte-Carlo estimate of the expected log-likelihood, an optimization device (§E.1 discusses the deployed optimizer's stop-gradient). Assumption 5 is the strength condition of Chen & Li (2022), whose Corollary 1 requires $(N \vee J)\log(N \wedge J) = o(\sigma^2_{K^*+1}(M^*))$: their $N \vee J$ is our $n + p$ up to a factor of two, their smallest nonzero singular value of the signal is our exact post-design gap d_{L^0}, and $\log(np)$ stands where they have $\log(n \wedge p)$, the same order for polynomial aspect ratios. As theirs, it concerns the weakest factor of the estimand, and the results that use it are silent about data in which some factor of $M^{0\perp}$ falls below the floor. Assumption 5^+ is pervasiveness (a non-vanishing fraction of cells participate, with loadings touching a non-vanishing fraction of genes at non-vanishing log-fold-change); why Theorem 2(d) needs it is explained in Remark 7. Assumption 6 prices canonical-factor identification and enters nowhere else. Notably absent: any assumption on the neighborhood graph, and any accuracy condition on the fitted dispersions.

E.3 Preliminary lemmas

Throughout, $C, c > 0$ denote constants depending only on the boxes of Assumption 2 and on (L, H, G, σ_β), changing from line to line, and $\omega_{n,p} := r(n + p)\log(np)$ with $r = L + H + G$.

Lemma 2 (negative-binomial structure)**.** *Write $\ell_\theta(y; \eta) = \log \mathrm{NB}(y \mid e^\eta, \theta)$ and $\sigma_\theta(\eta) = e^\eta/(e^\eta + \theta)$. Then:*

(i) $\partial^2_\eta \ell_\theta = -(y + \theta)\sigma(1 - \sigma)$, so ℓ_θ is concave in η with $0 \leq -\partial^2_\eta \ell \leq (y + \theta)/4$ and, on any fixed compact η-range and θ-set, $-\mathbb{E}_Y[\partial^2_\eta \ell] \geq c_{\min} > 0$ (a range-dependent constant);

(ii) for every $\theta' > 0$, $\mathbb{E}_Y[\partial_\eta \ell_{\theta'}(Y; \eta)]$ vanishes iff $\eta = \eta^0$ (the quasi-score is unbiased at the true mean whatever dispersion is used*; Gourieroux et al., 1984; Wedderburn, 1974), and $\eta \mapsto \mathbb{E}_Y[\ell_{\theta'}(Y; \eta)]$ is strictly concave with unique maximum at η^0;*

(iii) Y_{ij} *is sub-exponential, uniformly over the boxes;*

(iv) the dispersion profile $g_j(\xi') := \mathbb{E}_Y[\ell_{\theta_j^0}(Y;\eta^0) - \ell_{e^{\xi'}}(Y;\eta^0)] = \mathrm{KL}(\mathrm{NB}(\mu^0,\theta_j^0)\|\mathrm{NB}(\mu^0,e^{\xi'})) \geq 0$ *satisfies* $g_j(\xi') \leq C(\xi'-\xi_j^0)^2$ *for* ξ' *in a fixed compact neighborhood of the box and* $g_j(\xi') \leq C(1+|\xi'-\xi_j^0|)$ *for all* $\xi' \in \mathbb{R}$*;*

(v) the mixture derivatives are globally tame: for $|\eta|$ *in a fixed compact range and* all $\xi \in \mathbb{R}$, $|\partial_\xi \ell_{e^\xi}(y;\eta)| \leq C(1+y)$.

Proof. (i) Differentiate $\ell = c(y,\theta) + y\eta + \theta\log\theta - (y+\theta)\log(e^\eta+\theta)$ twice; $\sigma(1-\sigma) \leq 1/4$; the population lower bound uses $\mathbb{E}[y+\theta] = \mu^0 + \theta \geq \theta$ and $\sigma(1-\sigma) = \theta e^\eta/(e^\eta+\theta)^2$ bounded below on compacts. (ii) $\mathbb{E}[\partial_\eta \ell_{\theta'}] = \mu^0 - (\mu^0+\theta')\sigma_{\theta'}(\eta)$, strictly decreasing in η and zero at $\eta = \eta^0$; integrate. (iii) The NB probability generating function is finite beyond 1, uniformly on the boxes; standard equivalences (Vershynin, 2018) give the sub-exponential property. (iv) Nonnegativity is Kullback's inequality. The local quadratic bound: $g_j(\xi_j^0) = g_j'(\xi_j^0) = 0$ and $|g_j''| \leq C$ on compacts by dominated convergence under (iii). The global linear bound: with the identity $c(y,\theta) = \sum_{k=0}^{y-1}\log\frac{\theta+k}{k+1}$, for $\theta \leq 1$ one has $|c(y,\theta)| \leq |\log\theta| + 1 + \log(1+y)$ and the remaining terms of $-\ell$ are $\leq C(1+y)$, while for $\theta \geq 1$ the $\log\theta$-growth cancels exactly: $\theta\log\theta - (y+\theta)\log(e^\eta+\theta) = -y\log\theta - \theta\log(1+e^\eta/\theta) - y\log(1+e^\eta/\theta)$, whose last two terms are bounded by $e^{C_\eta}(1+y)$, and $c(y,\theta) - y\log\theta \in [-\log\Gamma(y+1), y\log 2]$; take expectations using (iii). (v) With $\theta = e^\xi$, collecting terms,

$$\partial_\xi \ell = \theta\big[\psi_0(y+\theta) - \psi_0(\theta)\big] - \theta\log\big(1+e^\eta/\theta\big) + \frac{\theta\,(e^\eta - y)}{e^\eta+\theta};$$

the digamma telescope gives $\theta[\psi_0(y+\theta)-\psi_0(\theta)] = \theta\sum_{k<y}(\theta+k)^{-1} \in [0,y]$, while $\theta\log(1+e^\eta/\theta) \in [0,e^\eta]$ and $\theta(e^\eta-y)/(e^\eta+\theta) \in [-y,e^\eta]$; every term is bounded by $C(1+y)$ uniformly in $\xi \in \mathbb{R}$ on a compact η-range. □

Lemma 3 (smoothing bounds)**.** *For every feasible* ϕ*, with* $\bar g_{ij}(m;\nu_j) := \mathbb{E}_{\xi\sim\mathcal{N}(m_{\xi_j},s^2_{\xi_j})}[\ell_{e^\xi}(Y_{ij};m)]$ *the dispersion-mixture plug-in objective:*

$$\mathbb{E}_q[\ell_{e^{\xi_j}}(Y_{ij};\eta_{ij})] \leq \bar g_{ij}(\bar\eta_{ij};\nu_j) \qquad \textit{(Jensen; mean-field independence gives } \mathbb{E}_q\eta_{ij} = \bar\eta_{ij}),$$

and, downward, $\mathbb{E}_q[\ell] \geq \bar g_{ij}(\bar\eta_{ij};\nu_j) - \frac{1}{8}\mathbb{E}_{q(\xi_j)}[Y_{ij}+e^{\xi_j}]\,\mathrm{Var}_q(\eta_{ij})$. *Consequently* $\mathcal{L}(\phi) \leq \sum_{ij}\bar g_{ij}(\bar\eta_{ij};\nu_j)$ *(all penalty terms are nonnegative).*

Proof. Concavity of ℓ in η at each dispersion draw, Jensen for the upper bound, exact second-order Taylor with the global curvature bound of Lemma 2(i) for the lower. □

Lemma 4 (existence)**.** $\sup_\Phi \mathcal{L}$ *is attained (on the closed set including collapsed rows).*

Proof. Compactify the variance boxes at 0; the entropy terms send $\mathcal{L} \to -\infty$ along any vanishing variance except matched row collapses, where the row-KL has limit 0; an upper semicontinuous function on a compact set attains its supremum, and the maximizer has positive variances except on collapsed rows. □

Remark 1 (the maximizer is interior in the score-variance block). Fix (i,ℓ) and all coordinates other than $s^2 := s^2_{W_{i\ell}}$. By Price's theorem (Price, 1958) the likelihood contributes $\frac{1}{2}\mathbb{E}_q[\partial^2_{W_{i\ell}}\sum_j \ell_{ij}] = -\frac{1}{2}\widehat F_{i\ell}$ to $\partial\mathcal{L}/\partial s^2$, with $\widehat F_{i\ell} \geq 0$ as in Lemma 5; the prior KL contributes $-\frac{1}{2}(1-1/s^2)$; every other term is s^2-free. Hence $\partial\mathcal{L}/\partial s^2 = \frac{1}{2s^2} - \frac{1}{2} - \frac{1}{2}\widehat F_{i\ell}$, which tends to $+\infty$ as $s^2 \downarrow 0$ and is negative for $s^2 > (1+\widehat F_{i\ell})^{-1}$, in particular for every $s^2 > 1$. Over $(0, s^2_{W,\max}]$ with $s^2_{W,\max} \geq 1$ a maximizer in this coordinate is therefore interior and stationary, or sits at $s^2_{W,\max}$ with one-sided derivative ≥ 0, which the display allows only if $\widehat F_{i\ell} = 0$ and $s^2_{W,\max} = 1$; and $\widehat F_{i\ell} = 0$ holds exactly on a collapsed row, the integrand $\alpha^2_{\ell j}(Y_{ij}+e^{\xi_j})\sigma_{ij}(1-\sigma_{ij})$ being positive unless $\mathbb{E}_q\alpha^2_{\ell j} = 0$ for all j. So Assumption 4(ii) holds at every (i,ℓ) for the maximizer, the collapsed rows being the stationary boundary $s^2 = 1$, and Φ needs no positive lower bound on s^2_W. In the derivative display

$\widehat{F}_{i\ell}$ is evaluated at the current s^2, so the identity $s^2 = (1+\widehat{F}_{i\ell})^{-1}$ is an implicit equation; the sign argument is unaffected, since $\widehat{F}_{i\ell} > 0$ on every alive row at every s^2. The deployed implementation clamps $\log s_W$ to $[-6, 2]$; this is a property of the optimizer, not of Φ. Where the lower clamp binds ($\widehat{F}_{i\ell} > e^{12} - 1$) the fitted variance exceeds its stationary value and the deployed SNR_ℓ falls below its stationary counterpart, but $S_\ell = \frac{1}{n}\sum_i \max\{e^{-12}, (1+\widehat{F}_{i\ell})^{-1}\} \le \max\{e^{-12}, (1+c_\star\widehat{R}_\ell)^{-1}\}$ gives $\mathrm{SNR}_\ell \ge \widehat{V}_\ell \min\{e^{12}, 1+c_\star\widehat{R}_\ell\}$, so the clamp can cost the rule an inclusion only for a column with $\widehat{V}_\ell \le e^{-12}$, and it can never create one.

Lemma 5 (what the SNR denominator is). *For a feasible ϕ let $\widehat{F}_{i\ell} := \mathbb{E}_q\big[\sum_j \alpha_{\ell j}^2 (Y_{ij} + e^{\xi_j})\sigma_{ij}(1-\sigma_{ij})\big] \ge 0$, the q-averaged observed information for $W_{i\ell}$; $\widehat{R}_\ell := \sum_j (m_{\alpha_{\ell j}}^2 + s_{\alpha_{\ell j}}^2)$, the fitted mass of loading row ℓ; $S_\ell := \frac{1}{n}\sum_i s_{W_{i\ell}}^2$; $\rho_\ell := S_\ell(1+c_\star\widehat{R}_\ell)$ and $\rho_\star := \max_\ell \rho_\ell$.*

- *(a) (deterministic lower bound) There is $c_\star > 0$, depending only on the boxes, with $c_\star\widehat{R}_\ell \le \widehat{F}_{i\ell}$ for every feasible ϕ and every (i,ℓ).*
- *(b) (the identity) Under Assumption 4(ii), $s_{W_{i\ell}}^2 = (1+\widehat{F}_{i\ell})^{-1}$ for every (i,ℓ); hence $S_\ell \le (1+c_\star\widehat{R}_\ell)^{-1}$, that is $\rho_\ell \le 1$ for every ℓ, and $\mathrm{SNR}_\ell \ge \widehat{V}_\ell(1+c_\star\widehat{R}_\ell)$. Without (ii), $\mathrm{SNR}_\ell \ge \widehat{V}_\ell(1+c_\star\widehat{R}_\ell)/\rho_\star$, every $c_\star$ in the second part of Theorem 1 and in Theorem 2 is replaced by $c_\star/\rho_\star$ with no other change, and $\rho_\star$ is computable from any fit.*
- *(c) (upper half) On $\{\max_j \bar{Y}_j \le C\}$, with $\bar{Y}_j = \frac{1}{n}\sum_i Y_{ij}$, $\frac{1}{n}\sum_i \widehat{F}_{i\ell} \le C\widehat{R}_\ell$, and under (b) also $\mathrm{SNR}_\ell \le \widehat{V}_\ell(1+C\widehat{R}_\ell)$.*

Proof. (b) The identity is the stationarity of $\mathcal{L}$ in $s_{W_{i\ell}}^2$: by Price's theorem (Price, 1958) the likelihood contributes half its q-averaged second $W_{i\ell}$-derivative, which is $-\frac{1}{2}\widehat{F}_{i\ell}$ with $\widehat{F}_{i\ell}$ as displayed, the interchange of derivative and expectation being justified by the sub-exponential envelope of Lemma 2(iii); the prior KL contributes $-\frac{1}{2}(1 - 1/s^2)$, and the graph penalty is s-free, so the implementation's gradient in $\log s_W$ is an ordinary one, the stop-gradient of §E.1 acting on the embedding means alone and never reaching this block; solve (Remark 1). Given the identity and (a), $s_{W_{i\ell}}^2 \le (1+c_\star\widehat{R}_\ell)^{-1}$ termwise, so $S_\ell \le (1+c_\star\widehat{R}_\ell)^{-1}$ and $\mathrm{SNR}_\ell = \widehat{V}_\ell/S_\ell \ge \widehat{V}_\ell(1+c_\star\widehat{R}_\ell)$; without the identity divide by $\rho_\ell \le \rho_\star$ instead. (a) Dropping $Y \ge 0$, $\widehat{F}_{i\ell} \ge \mathbb{E}_q[\sum_j \alpha_{\ell j}^2 \varkappa(\eta_{ij}, e^{\xi_j})]$ with $\varkappa = \theta^2 e^\eta/(e^\eta+\theta)^2 > 0$. Write $\eta_{ij} = \alpha_{\ell j}W_{i\ell} + T$ with the blocks mutually independent under mean-field. The α-draw enters through its truncated second moment, $\mathbb{E}[\alpha^2 \mathbf{1}_{|\alpha-m|\le 2s}] \ge 0.73(m^2+s^2)_{\alpha_{\ell j}}$ (Gaussian computation; the cross term vanishes by symmetry); the three remaining blocks are restricted to bulk events of probabilities at least 0.95 (a two-standard-deviation window for the W-draw), $\frac{1}{2}$ (a Chebyshev window for T, whose variance is bounded by the feasible boxes), and $\frac{15}{16}$ (a unit window for ξ); on the intersection, (η,θ) lies in a fixed compact set on which $\varkappa \ge c_G > 0$. Multiplying the three probabilities and $c_G \cdot 0.73$ gives $c_\star$; nothing random enters. (c) Upper bound: $\widehat{F}_{i\ell} \le \frac{1}{4}\sum_j (m^2+s^2)_{\ell j}(Y_{ij} + \mathbb{E}_q e^{\xi_j})$; average over i and use $\max_j \bar{Y}_j \le C$ (Bernstein plus a union over genes). The sandwich divides $\widehat{V}_\ell$ by the two ends of $\frac{1}{n}\sum_i (1+\widehat{F}_{i\ell})^{-1} \in [(1+C\widehat{R}_\ell)^{-1}, (1+c_\star\widehat{R}_\ell)^{-1}]$ (Jensen on the left, termwise on the right). □

Lemma 6 (the selected part). *For the fitted configuration, or for any feasible ϕ with its own statistics $\widehat{\mathcal{S}}$, $\widehat{V}_\ell$, S_ℓ, with $\rho_\star$ as in Lemma 5 ($\rho_\star \le 1$ under Assumption 4(ii)):*

- *(a) for $\ell \notin \widehat{\mathcal{S}}$, $\widehat{V}_\ell\widehat{R}_\ell \le \rho_\star/c_\star$ and $\|\widetilde{W}_{\cdot\ell}\widehat{\alpha}_{\ell\cdot}\|_F^2 = n\widehat{V}_\ell\|\widehat{\alpha}_{\ell\cdot}\|^2 \le n\rho_\star/c_\star$;*
- *(b) $\|\widehat{M}^\perp - \widehat{M}_{\widehat{\mathcal{S}}}^\perp\|_F \le L\sqrt{n\rho_\star/c_\star}$, hence $\|\widehat{M}_{\widehat{\mathcal{S}}}^\perp - M^{0\perp}\|_F \le \|\widehat{M}^\perp - M^{0\perp}\|_F + L\sqrt{n\rho_\star/c_\star}$;*
- *(c) every ℓ with $\|\widetilde{W}_{\cdot\ell}\widehat{\alpha}_{\ell\cdot}\|_F^2 > n\rho_\star/c_\star$ lies in $\widehat{\mathcal{S}}$.*

Proof. (a) If $\ell \notin \widehat{\mathcal{S}}$ then $\mathrm{SNR}_\ell \le 1$, so $\widehat{V}_\ell \le S_\ell$, and $S_\ell\widehat{R}_\ell = \rho_\ell\widehat{R}_\ell/(1+c_\star\widehat{R}_\ell) \le \rho_\star/c_\star$ by the definition of ρ_ℓ; hence $\widehat{V}_\ell\widehat{R}_\ell \le \rho_\star/c_\star$, and $\|\widehat{\alpha}_{\ell\cdot}\|^2 \le \widehat{R}_\ell$ with $\|\widetilde{W}_{\cdot\ell}\|^2 = n\widehat{V}_\ell$ gives the display. (Centering earns its keep here: the uncentered product carries the score-mean energy $n\bar{m}_\ell^2\|\widehat{\alpha}_{\ell\cdot}\|^2$,

$\bar{m}_\ell = \frac{1}{n}\sum_i m_{\widehat{W}_{i\ell}}$, which no selection bound controls and which is invisible to both the SNR statistic and k-means.) (b) $\widehat{M}^\perp - \widehat{M}^\perp_{\widehat{\mathcal{S}}} = \sum_{\ell\notin\widehat{\mathcal{S}}}(P_X^\perp \widetilde{W}_{\cdot\ell})(\widehat{\alpha}_{\ell\cdot}P_V^\perp)$ by $P_X^\perp = P_X^\perp P_1^\perp$; each term has Frobenius norm at most $\|\widetilde{W}_{\cdot\ell}\|\|\widehat{\alpha}_{\ell\cdot}\| \le \sqrt{n\rho_\star/c_\star}$ by (a), projections being contractions, and there are at most L terms. (c) If $\|\widetilde{W}_{\cdot\ell}\widehat{\alpha}_{\ell\cdot}\|_F^2 > n\rho_\star/c_\star$ then $\widehat{V}_\ell\widehat{R}_\ell \ge \widehat{V}_\ell\|\widehat{\alpha}_{\ell\cdot}\|^2 > \rho_\star/c_\star$, and Lemma 5(b) gives $\mathrm{SNR}_\ell \ge c_\star\widehat{V}_\ell\widehat{R}_\ell/\rho_\star > 1$. □

Lemma 7 (identifiability decomposition)**.** *For any difference $\Delta\eta = X\Delta\beta + \Delta\gamma V^\top + \Delta M$,*

$$P_X^\perp\,\Delta\eta\,P_V^\perp = P_X^\perp\,\Delta M\,P_V^\perp, \qquad \textit{so} \qquad \|P_X^\perp\Delta M P_V^\perp\|_F^2 \le \|\Delta\eta\|_F^2$$

with constant 1 and no assumption, projections being contractions. Moreover $\|P_X^\perp\Delta M P_V^\perp\|_F = \min_{\beta,\gamma}\|\Delta M - X\beta - \gamma V^\top\|_F$: the left-hand side is exactly the distance from ΔM to the design gauge orbit, which is why no stronger functional of $W\alpha$ is identifiable and no weaker one is needed. The restriction to functionals of $W\alpha$ matters: η^0 itself is identifiable and strictly stronger.

Proof. Uniqueness of the canonical representative: if $X\beta' + \gamma' V^\top = X\beta'' + \gamma'' V^\top$ with both γ's X-orthogonal, multiplying by $X^\top$ kills the γ-terms and full column rank gives $\beta' = \beta''$, then $\gamma' = \gamma''$. For the inequality: $P_X^\perp\Delta\eta P_V^\perp = P_X^\perp\Delta M P_V^\perp$ since the projections annihilate the design terms, and projections are contractions, so $\|P_X^\perp\Delta M P_V^\perp\|_F \le \|\Delta\eta\|_F$ with constant 1 and no assumption. The Frisch–Waugh–Lovell characterisation is the standard least-squares identity $\min_{\beta,\gamma}\|\Delta M - X\beta - \gamma V^\top\|_F = \|P_X^\perp\Delta M P_V^\perp\|_F$. Sharpness: if M' is any feasible signal with $P_X^\perp M' P_V^\perp = P_X^\perp M^0 P_V^\perp$, then $\Delta = M' - M^0$ has $P_X^\perp\Delta P_V^\perp = 0$, so by the attainment clause of Lemma 1 $\Delta = X\beta + \gamma V^\top$ for some (β,γ). Replacing β^0 by $\beta^0 - \beta$ and γ^0 by $\gamma^0 - \gamma$ leaves η^0 unchanged, so every such M' is observationally equivalent to M^0 whenever the replacement stays in Φ: nothing finer than $M^{0\perp}$ is identifiable. □

Proposition 1 (identifiable target; restated). *Fix $C \in \mathbb{R}^{H\times L}$ and let θ^0 and its translate θ^C, given by $W^0 \mapsto W^0 + XC$ and $\beta^0 \mapsto \beta^0 - C\alpha^0$, both lie in Φ. Then the two induce the same law of Y, for every estimator T*

$$\max_{\theta\in\{\theta^0,\theta^C\}} \tfrac{1}{np}\,\mathbb{E}_\theta\|T - W(\theta)\alpha(\theta)\|_F^2 \;\ge\; \tfrac{1}{4np}\|XC\alpha^0\|_F^2,$$

and $P_X^\perp W(\theta^C)\alpha(\theta^C)P_V^\perp = M^{0\perp}$ for every C. The same holds on the gene side: for $D \in \mathbb{R}^{L\times G}$ the translate θ^D given by $\alpha^0 \mapsto \alpha^0 + DV^\top$ and $\gamma^0 \mapsto \gamma^0 - W^0 D$, both settings again in Φ, satisfies the display with $\|W^0 DV^\top\|_F^2$ in place of $\|XC\alpha^0\|_F^2$, and $P_X^\perp W(\theta^D)\alpha(\theta^D)P_V^\perp = M^{0\perp}$.

Proof. The predictor is unchanged, $\eta(\theta^C) = X(\beta^0 - C\alpha^0) + \gamma^0 V^\top + (W^0 + XC)\alpha^0 + O = \eta^0$, and the dispersions are untouched, so the two parameters induce the same law $\mathbb{P}$ of Y. Write $a = T - W(\theta^0)\alpha(\theta^0)$ and $b = T - W(\theta^C)\alpha(\theta^C)$, so that $a - b = XC\alpha^0$. Since $\|a\|_F^2 + \|b\|_F^2 \ge \frac{1}{2}\|a-b\|_F^2$,

$$2\max_\theta \mathbb{E}_\theta\|T - W(\theta)\alpha(\theta)\|_F^2 \;\ge\; \mathbb{E}_{\mathbb{P}}\big[\|a\|_F^2 + \|b\|_F^2\big] \;\ge\; \tfrac{1}{2}\|XC\alpha^0\|_F^2,$$

which is the display after dividing by $2np$. For the last clause, $P_X^\perp XC\alpha^0 P_V^\perp = 0$ because $P_X^\perp X = 0$. For the gene-side translate, $\eta(\theta^D) = X\beta^0 + (\gamma^0 - W^0 D)V^\top + W^0(\alpha^0 + DV^\top) + O = \eta^0$, the $\pm W^0 DV^\top$ terms cancelling; the two-point argument is unchanged with $a - b = W^0 DV^\top$; and $P_X^\perp W^0 DV^\top P_V^\perp = 0$ because $V^\top P_V^\perp = 0$. Its feasibility is a separate requirement from the first translate's, since it constrains α and γ rather than W and β. Under the first translate $P_X^\perp W\alpha$ is invariant and under the second $W\alpha P_V^\perp$ is, so only the two-sided projection survives both. □

E.4 Consistency of the factor signal

Conventions for this section. $C, c > 0$ denote constants depending only on the boxes defining Φ (through C_η; the box $[\xi_{\min}, \xi_{\max}]$ of the dispersion means m_{ξ_j}; the cap $s^2_{\xi,\max} \le 1/16$ on the dispersion mixing variances; and the entrywise mean boxes $C_W, C_\alpha, C_\beta, C_\gamma$), on the row bounds $C_X := \max_i\|X_{i\cdot}\|$ and $C_V := \max_j\|V_{j\cdot}\|$ of Assumption 3, and on (L, H, G, σ_β); they change from line to line; λ_g is displayed explicitly and never absorbed into C, and σ_ξ does not enter (the

dispersion block's divergence cancels in Step 1 of the proof of Theorem 1). As before, $\omega_{n,p} := r(n+p)\log(np)$ with $r := L + H + G$, and $\mu_{ij}^0 := e^{\eta_{ij}^0}$. For gene j the dispersion-mixture parameter is $\nu_j = (m_{\xi_j}, s_{\xi_j})$ with $m_{\xi_j} \in [\xi_{\min}, \xi_{\max}]$ and $0 < s_{\xi_j}^2 \le s_{\xi,\max}^2$; Ξ is the set of feasible $\nu = (\nu_j)_{j\le p}$, and $\bar{g}_{ij}(m;\nu_j)$ is the plug-in objective of Lemma 3, recalled in the display below. A *feasible structured mean predictor* is $m = Xm_\beta + \gamma V^\top + m_W m_\alpha + O\mathbf{1}_p^\top$ with all blocks in their boxes and $|m_{ij}| \le C_\eta$; it is a function of at most $Hp + nG + nL + Lp \le r(n+p)$ scalars. For such m and $\nu \in \Xi$ set

$$\bar{g}_{ij}(m;\nu_j) := \mathbb{E}_{\xi\sim\mathcal{N}(m_{\xi_j},s_{\xi_j}^2)}\big[\ell_{e^\xi}(Y_{ij};m)\big], \qquad \bar{w}_{ij}(\nu_j) := \mathbb{E}_{\xi\sim\mathcal{N}(m_{\xi_j},s_{\xi_j}^2)}\Big[\frac{e^\xi}{\mu_{ij}^0+e^\xi}\Big],$$

$$S_{ij}(\nu_j) := \bar{w}_{ij}(\nu_j)\,\big(Y_{ij}-\mu_{ij}^0\big), \qquad Z(m,\nu) := \sum_{i,j} S_{ij}(\nu_j)\,\big(m_{ij}-\eta_{ij}^0\big) = \langle S(\nu),\, m-\eta^0\rangle.$$

$S(\nu)$ is the score of the mixture objective at the true predictor (Lemma 8(a)). We write $\widehat{\nu} := \nu(\widehat{\phi})$ for the fitted dispersion block and $\mathrm{KL}_\xi(\widehat{q}) := \sum_j \mathrm{KL}\,\big(\widehat{q}(\xi_j)\,\|\,\mathcal{N}(\mu_\xi,\sigma_\xi^2)\big)$ for its Kullback–Leibler term.

Step 0: curvature in the mean holds for every realization of the counts.

Lemma 8 (sure strong concavity in the mean)**.** *Let* $\kappa(\eta,\theta) := \theta^2 e^\eta/(e^\eta+\theta)^2$,

$$\kappa_0 := \min\big\{\kappa(\eta,\theta):\ |\eta|\le C_\eta,\ \xi_{\min}-1 \le \log\theta \le \xi_{\max}+1\big\} \ \ge\ \frac{e^{2\xi_{\min}-2-C_\eta}}{(e^{C_\eta}+e^{\xi_{\max}+1})^2}\ >0,$$

$c_0 := \frac{15}{16}\kappa_0$ *and* $c_w := \frac{15}{16}\,e^{\xi_{\min}-1}/(e^{C_\eta}+e^{\xi_{\min}-1})$.

(a) For every $y \ge 0$, $\theta > 0$ *and* $\eta \in \mathbb{R}$, $-\partial_\eta^2\ell_\theta(y;\eta) = (y+\theta)\sigma_\theta(\eta)(1-\sigma_\theta(\eta)) \ge \kappa(\eta,\theta)$ *and* $\partial_\eta\ell_\theta(y;\eta_{ij}^0) = \frac{\theta}{\mu_{ij}^0+\theta}\,(y-\mu_{ij}^0)$*; consequently, for every feasible* ν_j, $\partial_m\bar{g}_{ij}(\eta_{ij}^0;\nu_j) = S_{ij}(\nu_j)$ *with* $\bar{w}_{ij}(\nu_j) \in [c_w, 1]$.

(b) For every feasible ν_j *and every* $m \in [-C_\eta, C_\eta]$, $-\partial_m^2\bar{g}_{ij}(m;\nu_j) \ge c_0$.

(c) For every feasible ν_j *and every* $m \in [-C_\eta, C_\eta]$,

$$\bar{g}_{ij}(m;\nu_j) - \bar{g}_{ij}(\eta_{ij}^0;\nu_j)\ \le\ S_{ij}(\nu_j)\,(m-\eta_{ij}^0) - \tfrac{c_0}{2}\,(m-\eta_{ij}^0)^2.$$

The constants κ_0, c_0, c_w *depend only on* $(C_\eta, \xi_{\min}, \xi_{\max})$*, and (a) to (c) hold for every realization of* Y*. The lower bound* c_w *is recorded and not used below; the proofs use only* $\bar{w}_{ij} \le 1$.

Proof. (a) From $\ell_\theta(y;\eta) = c(y,\theta) + y\eta + \theta\log\theta - (y+\theta)\log(e^\eta+\theta)$ (the display in the proof of Lemma 2(i)), $\partial_\eta\ell_\theta = y - (y+\theta)\sigma_\theta(\eta)$, and since $\partial_\eta\sigma_\theta = \sigma_\theta(1-\sigma_\theta)$, $\partial_\eta^2\ell_\theta = -(y+\theta)\sigma_\theta(1-\sigma_\theta)$. Because $y \ge 0$, $(y+\theta)\sigma_\theta(1-\sigma_\theta) \ge \theta\sigma_\theta(1-\sigma_\theta) = \theta\cdot\frac{e^\eta}{e^\eta+\theta}\cdot\frac{\theta}{e^\eta+\theta} = \kappa(\eta,\theta)$. At $\eta = \eta_{ij}^0$, $\sigma_\theta(\eta_{ij}^0) = \mu_{ij}^0/(\mu_{ij}^0+\theta)$ and $y-(y+\theta)\mu_{ij}^0/(\mu_{ij}^0+\theta) = \theta(y-\mu_{ij}^0)/(\mu_{ij}^0+\theta)$. The score is dominated uniformly in θ: $y-(y+\theta)\sigma_\theta(\eta) = y(1-\sigma_\theta)-\theta\sigma_\theta \in [-e^\eta,\, y]$ because $\theta\sigma_\theta(\eta) = \theta e^\eta/(e^\eta+\theta) \le e^\eta$; so $|\partial_\eta\ell_\theta(y;\eta)| \le y + e^{C_\eta+1}$ on $|\eta| \le C_\eta + 1$, a ξ-free envelope on a neighborhood of the box, and dominated convergence allows differentiation of $\bar{g}_{ij}(m;\nu_j) = \mathbb{E}_\xi[\ell_{e^\xi}(Y_{ij};m)]$ under the Gaussian expectation. Hence $\partial_m\bar{g}_{ij}(\eta_{ij}^0;\nu_j) = \mathbb{E}_\xi\big[\frac{e^\xi}{\mu_{ij}^0+e^\xi}\big](Y_{ij}-\mu_{ij}^0) = S_{ij}(\nu_j)$. The integrand of $\bar{w}_{ij}$ lies in $(0,1)$, so $\bar{w}_{ij} \le 1$; on the event $|\xi - m_{\xi_j}| \le 1$, whose probability is at least $1 - s_{\xi_j}^2 \ge 15/16$ by Chebyshev's inequality, $e^\xi \ge e^{\xi_{\min}-1}$ and the integrand is at least $e^{\xi_{\min}-1}/(e^{C_\eta}+e^{\xi_{\min}-1})$ because $\mu_{ij}^0 \le e^{C_\eta}$ and $t \mapsto t/(\mu_{ij}^0+t)$ is increasing; hence $\bar{w}_{ij} \ge c_w$.

(b) The second derivative is dominated by $(y+e^\xi)/4$, which is integrable against the Gaussian law of ξ (a log-normal has finite mean), so differentiating twice under the expectation,

$$\begin{aligned}-\partial_m^2\bar{g}_{ij}(m;\nu_j) &= \mathbb{E}_\xi\big[(Y_{ij}+e^\xi)\sigma_{e^\xi}(m)(1-\sigma_{e^\xi}(m))\big]\ \ge\ \mathbb{E}_\xi\big[\kappa(m,e^\xi)\big]\\ &\ge\ \mathbb{E}_\xi\big[\kappa(m,e^\xi)\mathbf{1}\{|\xi-m_{\xi_j}|\le 1\}\big]\ \ge\ \kappa_0\,\mathbb{P}\big(|\xi-m_{\xi_j}|\le 1\big)\ \ge\ \tfrac{15}{16}\kappa_0 = c_0,\end{aligned}$$

by (a), because on the event $|\xi - m_{\xi_j}| \le 1$ one has $\log\theta = \xi \in [\xi_{\min} - 1, \xi_{\max} + 1]$ while $|m| \le C_\eta$, and by Chebyshev's inequality as in (a). The bound on κ_0: $\theta^2 \ge e^{2\xi_{\min}-2}$, $e^\eta \ge e^{-C_\eta}$ and $(e^\eta + \theta)^2 \le (e^{C_\eta} + e^{\xi_{\max}+1})^2$ on its defining set.

(c) Taylor's formula with integral remainder, $\bar g_{ij}(m;\nu_j) - \bar g_{ij}(\eta^0_{ij};\nu_j) = \partial_m \bar g_{ij}(\eta^0_{ij};\nu_j)(m-\eta^0_{ij}) + \int_{\eta^0_{ij}}^{m} (m-t)\,\partial_m^2 \bar g_{ij}(t;\nu_j)\,dt$. Every t between η^0_{ij} and m lies in $[-C_\eta, C_\eta]$ (Assumption 2 for η^0_{ij}, the box for m), so by (b) the integrand is at most $-c_0(m-t)$ in the orientation of the integral, whichever of m, η^0_{ij} is larger, and the remainder is at most $-\frac{c_0}{2}(m-\eta^0_{ij})^2$. Insert (a) for the first-order term. □

Remark 2 (population curvature). Taking expectations over Y in Lemma 8(c), with $\mathbb{E}\, S_{ij}(\nu_j) = 0$, gives the population form $\mathbb{E}_Y[\bar g_{ij}(\eta^0_{ij};\nu)] - \mathbb{E}_Y[\bar g_{ij}(m;\nu)] \ge \frac{c_0}{2}(m-\eta^0_{ij})^2$ for every feasible ν and $m \in [-C_\eta, C_\eta]$; the proofs below use only the sure form (c), the count-data analogue of Chen and Li's deterministic curvature bound $\inf b'' > 0$ (Chen & Li, 2022).

Step 3: the empirical process, localized in the mean and anchored at the fitted dispersion.

Lemma 9 (localized empirical-process bound at the fitted dispersion)**.** *Let Θ be the set of pairs (m,ν) with m a feasible structured mean predictor and $\nu \in \Xi$, and for $t \ge 1$ let $\Theta_t := \{(m,\nu) \in \Theta : \|m - \eta^0\|_F \le t\}$. Let K_1 be the sub-exponential scale of Lemma 2(iii), so that $\max_{i,j}\|Y_{ij} - \mu^0_{ij}\|_{\psi_1} \le K_1$; K_1 depends only on $(C_\eta, \theta_{\min})$: the moment generating function of $\mathrm{NB}(\mu,\theta)$ is finite at λ exactly when $e^\lambda < 1 + \theta/\mu$, so on the boxes $\|Y_{ij}\|_{\psi_1}$ is bounded in terms of e^{C_η} and $\theta_{\min}$ alone, growing as $\theta_{\min} \to 0$, and centering costs an absolute factor (Vershynin, 2018). Let $\mathcal{E}_Y := \{\sum_{i,j} Y_{ij} \le C_Y\, np\}$ with $C_Y := e^{C_\eta} + K_1$. There are an absolute constant $c_Y > 0$, a constant C_1 depending only on the boxes, (C_X, C_V) and r, and n_0 depending only on the boxes, (C_X, C_V) and r, such that $\mathbb{P}(\mathcal{E}_Y) \ge 1 - e^{-c_Y np}$ for $np \ge n_0$ and, for all $np \ge n_0$, $t \ge 1$ and $x \ge 1$, with probability at least $1 - 2e^{-x} - e^{-c_Y np}$,*

$$\sup_{(m,\nu)\in\Theta_t} \big|Z(m,\nu)\big| \;\le\; C_1\Big[\, t\sqrt{\omega_{n,p} + x} + (\omega_{n,p} + x)\Big].$$

Proof. (i) Pointwise bound. Fix $(m,\nu) \in \Theta$. The summands $\zeta_{ij} := S_{ij}(\nu_j)(m_{ij} - \eta^0_{ij})$ are independent across (i,j), each being a function of Y_{ij} alone once (m,ν) is fixed; they are centered because $\mathbb{E}Y_{ij} = \mu^0_{ij}$; and $\|\zeta_{ij}\|_{\psi_1} \le K_1|m_{ij} - \eta^0_{ij}|$ because $|\bar w_{ij}(\nu_j)| \le 1$ (Lemma 8(a)). Bernstein's inequality for sums of independent centered sub-exponential variables (Vershynin, 2018, Theorem 2.8.1) gives an absolute constant $c_B > 0$ such that, for all $s \ge 0$,

$$\mathbb{P}\big(|Z(m,\nu)| \ge s\big) \le 2\exp\Big[-c_B \min\Big(\frac{s^2}{K_1^2\|m-\eta^0\|_F^2},\; \frac{s}{K_1 \max_{i,j}|m_{ij} - \eta^0_{ij}|}\Big)\Big].$$

If $\|m - \eta^0\|_F \le \tau$ and $\max_{i,j}|m_{ij} - \eta^0_{ij}| \le b$, then $s := K_1[\tau\sqrt{z/c_B} + bz/c_B]$ makes both ratios at least z/c_B, so

$$\mathbb{P}\Big(|Z(m,\nu)| \ge K_1\big[\tau\sqrt{z/c_B} + b\,z/c_B\big]\Big) \le 2e^{-z} \qquad \text{for every } z \ge 0. \tag{24}$$

(ii) The event $\mathcal{E}_Y$. $\sum_{i,j}(Y_{ij} - \mu^0_{ij})$ is a sum of np independent centered variables with ψ_1-norm at most K_1; the same inequality with all coefficients equal to one and $s = K_1 np$ gives $\mathbb{P}(\sum_{i,j}(Y_{ij} - \mu^0_{ij}) \ge K_1 np) \le 2e^{-c_B np} \le e^{-c_Y np}$ with $c_Y := c_B/2$ once $np \ge (2\log 2)/c_B$. Since $\sum_{i,j}\mu^0_{ij} \le e^{C_\eta} np$, the complement of that event is contained in $\mathcal{E}_Y$.

(iii) The parameter net. The pairs $(m,\nu) \in \Theta$ are images of the parameter vector $\vartheta := (m_\beta, \gamma, m_W, m_\alpha, (m_{\xi_j}, s_{\xi_j})_{j\le p})$, of dimension $N \le r(n+p) + 2p \le 2r(n+p)$ (as $r \ge 3$), each coordinate ranging in an interval of length at most D_0, where D_0 depends only on the boxes. Let $\mathcal{N}$ be an ϵ-net of the parameter box in ℓ_∞ with $\epsilon := (np)^{-3}$ and net points inside the box, so that $|\mathcal{N}| \le \lceil D_0/\epsilon\rceil^N$ and, for $np \ge \max\{3, D_0^{-1/3}\}$, so that $\lceil D_0/\epsilon\rceil \le 2D_0(np)^3$,

$$\log|\mathcal{N}| \;\le\; 2r(n+p)\log\big(2D_0(np)^3\big) \;\le\; C_2\,\omega_{n,p}, \qquad C_2 := 2\big(3 + \log^+(2D_0)\big).$$

A net point $\vartheta^\sharp$ defines $(m^\sharp, \nu^\sharp)$ by the same formulas; $m^\sharp$ need not satisfy $|m^\sharp_{ij}| \le C_\eta$, which is not used below.

(iv) Net error and shell preservation. For $(m,\nu) \in \Theta_t$ with parameter ϑ choose $\vartheta^\sharp \in \mathcal{N}$ with $\|\vartheta - \vartheta^\sharp\|_\infty \le \epsilon$. Three elementary bounds hold. First, writing $m_W m_\alpha - m_W^\sharp m_\alpha^\sharp = (m_W - m_W^\sharp)m_\alpha + m_W^\sharp(m_\alpha - m_\alpha^\sharp)$ and using the bounded rows of Assumption 3, $\sum_h |X_{ih}| \le \sqrt{H}\, C_X$ and $\sum_g |V_{jg}| \le \sqrt{G}\, C_V$,

$$|m_{ij} - m_{ij}^\sharp| \;\le\; \epsilon\Big[\sqrt{H}\, C_X + \sqrt{G}\, C_V + L(C_\alpha + C_W)\Big] =: C_3\,\epsilon.$$

Second, $|S_{ij}(\nu_j)| \le |Y_{ij} - \mu_{ij}^0| \le Y_{ij} + e^{C_\eta}$, because $0 < \bar w_{ij} \le 1$. Third, with $h(\xi) := e^\xi/(\mu_{ij}^0 + e^\xi)$, $0 \le h' = h(1-h) \le \frac14$, and $\zeta \sim \mathcal{N}(0,1)$,

$$\big|\bar w_{ij}(\nu_j) - \bar w_{ij}(\nu_j^\sharp)\big| \le \mathbb{E}\big|h(m_{\xi_j} + s_{\xi_j}\zeta) - h(m_{\xi_j}^\sharp + s_{\xi_j}^\sharp\zeta)\big| \le \tfrac14\big(\epsilon + \epsilon\,\mathbb{E}|\zeta|\big) \le \tfrac{\epsilon}{2},$$

so that $|S_{ij}(\nu_j) - S_{ij}(\nu_j^\sharp)| \le \frac{\epsilon}{2}(Y_{ij} + e^{C_\eta})$. Since $|m_{ij}^\sharp - \eta_{ij}^0| \le |m_{ij} - \eta_{ij}^0| + C_3\epsilon \le 2C_\eta + 1$ for $np \ge C_3^{1/3}$,

$$\begin{aligned}|Z(m,\nu) - Z(m^\sharp,\nu^\sharp)| &\le \sum_{i,j} |S_{ij}(\nu_j)|\,|m_{ij} - m_{ij}^\sharp| + \sum_{i,j} |S_{ij}(\nu_j) - S_{ij}(\nu_j^\sharp)|\,|m_{ij}^\sharp - \eta_{ij}^0| \\ &\le \big(C_3 + C_\eta + 1\big)\,\epsilon \sum_{i,j}\big(Y_{ij} + e^{C_\eta}\big),\end{aligned}$$

which on $\mathcal{E}_Y$ is at most $(C_3 + C_\eta + 1)(C_Y + e^{C_\eta})\,(np)^{-2} =: C_4(np)^{-2} \le 1$ for $np \ge \sqrt{C_4}$. Moreover $\|m - m^\sharp\|_F \le \sqrt{np}\,\max_{i,j}|m_{ij} - m_{ij}^\sharp| \le C_3(np)^{-5/2} \le 1 \le t$ for $np \ge C_3^{2/5}$, hence $\|m^\sharp - \eta^0\|_F \le 2t$: the net point of a member of Θ_t lies in the shell of radius $2t$, and $\max_{i,j}|m_{ij}^\sharp - \eta_{ij}^0| \le b := 2C_\eta + 1$.

(v) Union bound. Let $\mathcal{N}_t := \{\vartheta^\sharp \in \mathcal{N} : \|m^\sharp - \eta^0\|_F \le 2t\}$ and $z := x + C_2\omega_{n,p}$. By (24) with $\tau = 2t$ and $b = 2C_\eta + 1$, applied to every point of $\mathcal{N}_t$,

$$\mathbb{P}\Big(\exists\,\vartheta^\sharp \in \mathcal{N}_t :\; |Z(m^\sharp,\nu^\sharp)| \ge K_1\big[2t\sqrt{z/c_B} + b\,z/c_B\big]\Big) \le |\mathcal{N}|\,2e^{-z} \le 2e^{-x}.$$

On the complement of this event intersected with $\mathcal{E}_Y$, every $(m,\nu) \in \Theta_t$ satisfies, by (iv), $|Z(m,\nu)| \le K_1[2t\sqrt{z/c_B} + bz/c_B] + 1$. Since $z \le \max(C_2,1)(\omega_{n,p} + x)$ and $1 \le \omega_{n,p} + x$, the claim holds with $C_1 := K_1 \max(2,b)\max(C_2,1)\max(c_B^{-1}, c_B^{-1/2}) + 1$; the requirements on np collected above define n_0. □

Theorem 1 and its proof.

Theorem 1 (consistency of the factor signal; restated). *Let Assumptions 1 to 3 hold and let $\widehat\phi \in \Phi$ satisfy Assumption 4(i), the basic inequality against the comparison configuration $\widetilde\phi(\widehat\nu)$ of §E.1. There are constants $C > 0$ and n_0, depending only on the boxes, (C_X, C_V) and (L, H, G, σ_β) and not on λ_g, such that for $np \ge n_0$ the following hold on an event of probability at least $1 - C(np)^{-1}$.*

(a) $\|\widehat M^\perp - M^{0\perp}\|_F^2 \;\le\; \|\widehat\eta - \eta^0\|_F^2 \;\le\; C\big[r(n+p)\log(np) + \lambda_g\, n\big]$.

(b) If in addition Assumption 4(ii) holds, then, with $c_\star$ the constant of Lemma 5,

$$\|\widehat M_{\widehat{\mathcal{S}}}^\perp - M^{0\perp}\|_F \;\le\; \sqrt{C\big[r(n+p)\log(np) + \lambda_g\, n\big]} \;+\; L\sqrt{n/c_\star}.$$

In particular, under Assumption 2 and for every fixed λ_g, $\frac{1}{np}\|\widehat M^\perp - M^{0\perp}\|_F^2 \xrightarrow{p} 0$ and $\frac{1}{np}\|\widehat M_{\widehat{\mathcal{S}}}^\perp - M^{0\perp}\|_F^2 \xrightarrow{p} 0$, with no assumption on the strength of any factor.

Proof. Step 1 (the comparison configuration and its ledger). Let $\widetilde\phi := \widetilde\phi(\widehat\nu)$ be the comparison configuration of §E.1: its means are a parametrization $(\beta^0, \gamma^0, W^0, \alpha^0)$ of the truth inside the boxes, with live rows $\mathcal{A} := \{\ell : \alpha_{\ell\cdot}^0 \ne 0\}$, so that $\bar\eta(\widetilde\phi) = \eta^0$ and the mean-predictor box $|\bar\eta_{ij}(\widetilde\phi)| =$

$|\eta^0_{ij}| \le C_\eta$ holds by Assumption 2; every variance of the blocks β, $W_{\cdot\mathcal{A}}$ and $\alpha_{\mathcal{A}\cdot}$ equals $(n+p)^{-1}$; $\widetilde{\lambda}^2_\ell = \widetilde{R}_\ell/p$ with $\widetilde{R}_\ell = \|\alpha^0_{\ell\cdot}\|^2 + p(n+p)^{-1}$ for $\ell \in \mathcal{A}$; the rows $\ell \notin \mathcal{A}$ are collapsed, with $m_{W\cdot\ell} = 0$ and $s^2_{W\cdot\ell} = 1$; and $\nu(\widetilde{\phi}) = \widehat{\nu}$. It lies in Φ by the check of §E.1, which uses the mean boxes, the mean-predictor box, the λ box with $\lambda^2_{\max} \ge C^2_\alpha + 1$, the absence of positive lower bounds on the W, α and β variance ranges, the value 1 in the W-variance range, the closure at collapsed rows, and $\widehat{\nu} \in \Xi$. (With a positive floor $\lambda_{\min}$ on the λ box one takes $\widetilde{\lambda}_\ell := \max\{\lambda_{\min}, \sqrt{\widetilde{R}_\ell/p}\}$; then $\sum_j (s^2 + m^2)_{\alpha_{\ell j}}/\widetilde{\lambda}^2_\ell \le p$, $p\log\widetilde{\lambda}^2_\ell \le p\log\max\{\lambda^2_{\min}, C^2_\alpha + 1\}$, and (L3) holds with another constant.) The configuration depends on the fit only through $\widehat{\nu}$.

Its ledger, on the event $\mathcal{E}_Y$ of Lemma 9, is as follows. (L1) *Likelihood.* By the lower bound of Lemma 3 at $\widetilde{\phi}$, $\mathbb{E}_{\widetilde{q}}[\ell_{e^{\xi_j}}(Y_{ij}; \eta_{ij})] \ge \bar{g}_{ij}(\eta^0_{ij}; \widehat{\nu}_j) - \frac{1}{8}\mathbb{E}_{\widetilde{q}(\xi_j)}[Y_{ij} + e^{\xi_j}]\operatorname{Var}_{\widetilde{q}}(\eta_{ij})$. Under $\widetilde{q}$ the blocks are independent Gaussians, so $\operatorname{Var}_{\widetilde{q}}(\eta_{ij}) = \sum_h X^2_{ih} s^2_{\beta_{hj}} + \sum_{\ell\in\mathcal{A}} \big[(W^0_{i\ell})^2 s^2_{\alpha_{\ell j}} + (\alpha^0_{\ell j})^2 s^2_{W_{i\ell}} + s^2_{W_{i\ell}} s^2_{\alpha_{\ell j}}\big] \le (n+p)^{-1}\big[C^2_X + L(C^2_W + C^2_\alpha + 1)\big]$, where a collapsed row contributes nothing because under $\widetilde{q}$ its $\alpha_{\ell\cdot}$ is a point mass at zero, so $W_{i\ell}\alpha_{\ell j} = 0$ almost surely whatever $s^2_{W_{i\ell}} = 1$ does, and $\mathbb{E}_{\widehat{q}} e^{\xi_j} = \exp(m_{\widehat{\xi}_j} + s^2_{\widehat{\xi}_j}/2) \le e^{\xi_{\max}+1/32}$. Summing over (i,j) on $\mathcal{E}_Y$, the total Taylor loss is at most $C\,np/(n+p)$. (L2) $\mathrm{KL}_W(\widetilde{\phi}) = \frac{1}{2}\sum_i \sum_{\ell\in\mathcal{A}} \big[(n+p)^{-1} + (W^0_{i\ell})^2 - 1 + \log(n+p)\big] \le \frac{nL}{2}\big[C^2_W + \log(n+p)\big]$; a collapsed column contributes $\frac{1}{2}\sum_i (1 + 0 - 1 - \log 1) = 0$. (L3) $\mathrm{KL}_\alpha(\widetilde{\phi})$: for $\ell \in \mathcal{A}$, with $\widetilde{\lambda}^2_\ell = \widetilde{R}_\ell/p$, $\frac{1}{2}\sum_j \big[\frac{s^2_{\alpha_{\ell j}} + m^2_{\alpha_{\ell j}}}{\widetilde{\lambda}^2_\ell} - 1 + \log\widetilde{\lambda}^2_\ell - \log s^2_{\alpha_{\ell j}}\big] = \frac{1}{2}\big[p - p + p\log(\widetilde{R}_\ell/p) + p\log(n+p)\big] \le \frac{p}{2}\big[\log(C^2_\alpha + 1) + \log(n+p)\big]$, because $\widetilde{R}_\ell/p \le C^2_\alpha + 1$; a collapsed row contributes 0 by the closure convention. The total is at most $\frac{pL}{2}[\log(C^2_\alpha + 1) + \log(n+p)]$. (L4) $\mathrm{KL}_\beta(\widetilde{\phi}) = \frac{1}{2}\sum_{h,j}\big[\frac{(n+p)^{-1} + (\beta^0_{hj})^2}{\sigma^2_\beta} - 1 + \log\sigma^2_\beta + \log(n+p)\big] \le \frac{Hp}{2}\big[\frac{1+C^2_\beta}{\sigma^2_\beta} + |\log\sigma^2_\beta| + \log(n+p)\big]$. (L5) The half-Cauchy term: $\sum_{\ell\in\mathcal{A}} \log(1 + \widetilde{\lambda}^2_\ell) \le L\log(2 + C^2_\alpha)$. (L6) The graph term: A has nonnegative entries and unit row sums, so $|((I - A)W^0)_{i\ell}| \le |W^0_{i\ell}| + \sum_{i'} A_{ii'}|W^0_{i'\ell}| \le 2C_W$ and $\lambda_g\|(I - A)W^0\|^2_F \le 4\lambda_g\, nLC^2_W$; no other property of A is used. (L7) γ is a point estimate under a flat prior and contributes no divergence; the dispersion block contributes exactly $\mathrm{KL}_\xi(\widehat{q})$, which is finite because $\widehat{\nu} \in \Xi$ has $0 < s^2_{\widehat{\xi}_j} \le s^2_{\xi,\max}$ and $m_{\widehat{\xi}_j}$ in its box. Collecting (L1) to (L7), and using $\log(n+p) \le \log(np)$ for $n, p \ge 2$,

$$\mathcal{L}\big(\widetilde{\phi}(\widehat{\nu})\big) \;\ge\; \sum_{i,j} \bar{g}_{ij}(\eta^0_{ij}; \widehat{\nu}_j) - \mathrm{KL}_\xi(\widehat{q}) - \mathrm{Rem}, \qquad \mathrm{Rem} \le C_5\Big[\omega_{n,p} + \lambda_g n + \frac{np}{n+p}\Big]. \tag{25}$$

On the fitted side, Jensen's inequality (the upper bound of Lemma 3) and the nonnegativity of every dropped term (the divergences of the W, α and β blocks, collapsed rows contributing 0; the half-Cauchy term $\sum_\ell \log(1 + \widehat{\lambda}^2_\ell)$; and the penalty at its analyzed value $\lambda_g\|(I - A)\widehat{m}_W\|^2_F$, the stop-gradient changing no value) give

$$\mathcal{L}(\widehat{\phi}) \;\le\; \sum_{i,j} \bar{g}_{ij}(\widehat{\eta}_{ij}; \widehat{\nu}_j) - \mathrm{KL}_\xi(\widehat{q}), \tag{26}$$

where $\mathrm{KL}_\xi(\widehat{q})$ is kept. Assumption 4(i), (25) and (26) give, after cancelling the finite term $\mathrm{KL}_\xi(\widehat{q})$ on both sides,

$$\sum_{i,j}\Big[\bar{g}_{ij}(\widehat{\eta}_{ij}; \widehat{\nu}_j) - \bar{g}_{ij}(\eta^0_{ij}; \widehat{\nu}_j)\Big] \;\ge\; -\mathrm{Rem}. \tag{27}$$

The cancellation is the only place where the choice $\nu(\widetilde{\phi}) = \widehat{\nu}$ is used; it is what removes any reference dispersion, and with it any need to bound the fitted dispersions, from the argument.

Step 2 (curvature). Since $|\widehat{\eta}_{ij}| \le C_\eta$ by the box of Φ and $\widehat{\nu} \in \Xi$, Lemma 8(c) applied at every (i,j) with $m = \widehat{\eta}_{ij}$ and $\nu_j = \widehat{\nu}_j$, summed, bounds the left side of (27) by $Z(\widehat{\eta}, \widehat{\nu}) - \frac{c_0}{2}\|\widehat{\eta} - \eta^0\|^2_F$. Writing $\rho := \|\widehat{\eta} - \eta^0\|_F$,

$$\tfrac{c_0}{2}\,\rho^2 \;\le\; Z(\widehat{\eta}, \widehat{\nu}) + \mathrm{Rem}. \tag{28}$$

Step 3 (the event). Set $x := 2\log(np)$, $t_* := \sqrt{\omega_{n,p}}$ (≥ 1), $K := \lceil \log_2(2C_\eta\sqrt{np})\rceil$ and $t_k := 2^k t_*$ for $0 \le k \le K$. Let $\mathcal{E}$ be the intersection of $\mathcal{E}_Y$ with the $K + 1$ events on which the display of

Lemma 9 holds for $t = t_k$. Then $\mathbb{P}(\mathcal{E}^c) \le e^{-c_Y np} + (K+1)\big[2(np)^{-2} + e^{-c_Y np}\big] \le C(np)^{-1}$ for $np \ge n_0$, because $K + 1 \le \log_2(np) + C$.

Step 4 (peeling). Work on $\mathcal{E}$. Since both $\widehat{\eta}$ and η^0 have entries in $[-C_\eta, C_\eta]$, $\rho \le 2C_\eta\sqrt{np} \le t_K$. If $\rho \le t_*$ then $\rho^2 \le \omega_{n,p}$. Otherwise there is $k \in \{1, \dots, K\}$ with $t_{k-1} < \rho \le t_k$; then $(\widehat{\eta}, \widehat{\nu}) \in \Theta_{t_k}$, and Lemma 9 with $t_k = 2t_{k-1} < 2\rho$ and $x = 2\log(np) \le \omega_{n,p}$ gives $Z(\widehat{\eta}, \widehat{\nu}) \le C_1[t_k\sqrt{\omega_{n,p} + x} + \omega_{n,p} + x] \le C_1[2\sqrt{2}\,\rho\sqrt{\omega_{n,p}} + 2\omega_{n,p}]$. With (28), $\frac{c_0}{2}\rho^2 \le 2\sqrt{2}C_1\sqrt{\omega_{n,p}}\,\rho + 2C_1\omega_{n,p} + \mathrm{Rem}$. For $a, b, d > 0$ the inequality $a\rho^2 \le b\rho + d$ implies $\rho \le (b + \sqrt{b^2 + 4ad})/(2a)$, hence $\rho^2 \le (4b^2 + 8ad)/(4a^2) = b^2/a^2 + 2d/a$; with $a = c_0/2$, $b = 2\sqrt{2}C_1\sqrt{\omega_{n,p}}$ and $d = 2C_1\omega_{n,p} + \mathrm{Rem}$,

$$\rho^2 \;\le\; \frac{32C_1^2}{c_0^2}\,\omega_{n,p} + \frac{8C_1}{c_0}\,\omega_{n,p} + \frac{4}{c_0}\,\mathrm{Rem} \;\le\; C\Big[\omega_{n,p} + \lambda_g n + \frac{np}{n+p}\Big] \;\le\; C\big[\omega_{n,p} + \lambda_g n\big],$$

since $np/(n+p) \le n \wedge p \le \omega_{n,p}$. In both cases

$$\|\widehat{\eta} - \eta^0\|_F^2 \;\le\; C\big[r(n+p)\log(np) + \lambda_g n\big]. \tag{29}$$

Step 5 (transfer). $\widehat{\eta} - \eta^0 = X(\widehat{m}_\beta - \beta^0) + (\widehat{\gamma} - \gamma^0)V^\top + (\widehat{M} - M^0)$, the offset cancelling, so Lemma 7 gives $\|\widehat{M}^\perp - M^{0\perp}\|_F = \|P_X^\perp(\widehat{M} - M^0)P_V^\perp\|_F \le \|\widehat{\eta} - \eta^0\|_F$, with constant one and no condition on the designs beyond Assumption 3. With (29) this is (a).

Step 6 (the selected part). Under Assumption 4(ii) one has $\rho_\star \le 1$ by Lemma 5(b), and Lemma 6(b) gives

$$\|\widehat{M}^\perp_{\widehat{\mathcal{S}}} - M^{0\perp}\|_F \;\le\; \|\widehat{M}^\perp - M^{0\perp}\|_F + L\sqrt{n/c_\star};$$

with (a) this is (b). Without Assumption 4(ii) the same lemma holds with $c_\star/\rho_\star$ in place of $c_\star$, $\rho_\star$ being computable from the fit.

Asymptotic form. Divide by np: $r(n+p)\log(np)/(np) = r\log(np)(\frac{1}{n} + \frac{1}{p}) \to 0$ by Assumption 2, $\lambda_g n/(np) = \lambda_g/p \to 0$, and $L^2 n/(c_\star np) = L^2/(c_\star p) \to 0$. □

Remark 3 (what the proof uses, and what survives). (i) *The estimator.* Part (a) uses Assumption 4(i) alone, at one line, (27), against one configuration, $\widetilde{\phi}(\widehat{\nu})$, which depends on the fit only through its dispersion block and can be evaluated from a fit and the truth in a simulation; for a maximizer of $\mathcal{L}$ over Φ the inequality is automatic. This is why Assumption 4(i) names $\widetilde{\phi}(\widehat{\nu})$ rather than a fit-free configuration. Part (b) uses Assumption 4(ii) in addition, through the identity of Lemma 5. The deployed optimizer's stop-gradient acts on the field of m_W alone, so the dispersion block, like every block other than m_W, is optimized against $\mathcal{L}$ itself, and copying it adds no idealization beyond Assumption 4(i). (ii) *Dispersions.* The proof assumes nothing about $\widehat{\nu}$ beyond feasibility and concludes nothing about it: the mixture score $S(\nu)$ at η^0 has mean zero for every ν (Lemma 2(ii) in mixture form) and the curvature is bounded below for every count and every ν (Lemma 8). No reference dispersion, no curvature lower bound in ξ, and no term comparing two mixtures appears; a bound anchored at a fixed reference mixture would have to pay for such a term with variance proxy np in place of $\|m - \eta^0\|_F^2$. (iii) *The graph.* It enters part (a) twice, at (L6), where only the nonnegativity and unit row sums of A are used, and in (26), where the penalty is dropped by its sign; and part (b) once more, through the penalty's independence of the variance coordinates, which makes the identity of Lemma 5 exact. The $\lambda_g n$ term of Theorem 1 is its only trace; no property of the neighborhoods, no independence between graph and counts, and no smoothness of the truth along the graph is assumed. (iv) *No angle condition* between the factor signal and the design spaces appears: the transfer is Lemma 7, with constant one. Finally, the localization is in the mean block alone and the net is at the ℓ_∞ scale of the parameters, so no Frobenius-scale covering of the low-rank predictor class is needed; and the logarithm in Theorem 1 has two sources of the same order, the entropy ledger (L2) to (L4), which charges $\log(n+p)$ per live variational coordinate of the W, α and β blocks of the comparison configuration, and the union over the net.

Remark 4 (the duality view of the linear term). The only stochastic object in the proof is $\langle S(\widehat{\nu}), \Delta\rangle$ with $\Delta := \widehat{\eta} - \eta^0$ of rank at most $r' := H + G + 2L$. For fixed ν, nuclear-norm duality, $\langle S(\nu), \Delta\rangle \le \sqrt{r'}\,\|S(\nu)\|_{\mathrm{op}}\|\Delta\|_F$, and an operator-norm bound $\|S(\nu)\|_{\mathrm{op}} \le C(\sqrt{n} + \sqrt{p})$ up to logarithmic lower-order terms (Bandeira & van Handel, 2016) give Chen and Li's log-free linear term (Chen & Li, 2022, Lemmas 1 and 3), transplanted to the log link by Lemma 8(c). Uniformity over the $2p$

fitted dispersion parameters then costs a union bound of order $\sqrt{p\log(np)}$ in the operator norm, up to polylogarithmic factors from the truncation in their Lemma 3, and the entropy ledger costs $r(n+p)\log(n+p)$ regardless, so this route gives the order of Theorem 1 and locates its logarithm in the variational family's entropy ledger and the dispersion block, not in the noise. No log-free rate is claimed.

Remark 5 (three floors). Let $d^2 = d_k^2$ be the strength of one true factor, with canonical scores $\sqrt{n}\,u_k$ and loadings $(d_k/\sqrt{n})v_k^\top$. Three bands calibrate the assumptions, on the Gaussian rank-one spiked model $d\,uv^\top + E$ with E an $n \times p$ matrix of independent unit-variance entries, the local picture of a count model after standardization by the curvature weights. Everything quoted about that model is cited, not proved; the incidental-parameter count below is a calibration, not a theorem; and no converse is proved for the count model. *Detection,* $d^2 \lesssim \sqrt{np}$. Below the Baik–Ben Arous–Péché threshold $d^2 = \sqrt{np}$ the spike leaves the spectrum unchanged (Baik et al., 2005; Paul, 2007); in the spiked covariance model, whose $n \times p$ sample matrix is the additive model $d\,uv^\top + E$ with Gaussian scores u, the likelihood ratio is contiguous below that threshold for a spherical prior on v, so no test is consistent against it (Onatski et al., 2013; Perry et al., 2018; El Alaoui & Jordan, 2018). *Partial recovery,* $\sqrt{np} \ll d^2 \lesssim n \vee p$. Above the threshold the spike is detected, but the top singular vectors of the data have limiting squared overlaps

$$|\langle \widehat{u}, u\rangle|^2 \;\to\; \frac{d^4 - np}{d^4 + n\,d^2}, \qquad |\langle \widehat{v}, v\rangle|^2 \;\to\; \frac{d^4 - np}{d^4 + p\,d^2}$$

in the proportional regime $n/p \to c \in (0,\infty)$ (Benaych-Georges & Nadakuditi, 2012); they tend to one only when $d^2 \gg n \vee \sqrt{np}$ and $d^2 \gg p \vee \sqrt{np}$ respectively. Outside the proportional regime an incidental-parameter count (Neyman & Scott, 1948) indicates the same thresholds: each of the n canonical scores is estimated from p counts that carry information of order $\|(d/\sqrt{n})v\|^2 = d^2/n$, so the score vector's total squared error is of order n^2/d^2 against $\|\sqrt{n}\,u\|^2 = n$, and it vanishes relative to the signal only if $d^2 \gg n$, even with the loadings known; symmetrically, the loading vector needs $d^2 \gg p$. The count thus indicates that no method can recover the score direction u_k to a vanishing angle when $d_k^2 \lesssim n \vee \sqrt{np}$, nor the loading direction v_k when $d_k^2 \lesssim p \vee \sqrt{np}$, and that recovering the canonical pair needs $d_k^2 \gg n \vee p$; the proved, matrix-level form of the same count is the minimax lower bound $\|\widehat{M} - M^*\|_F^2 \gtrsim K(N \vee J)$ of Chen and Li (Chen & Li, 2022, Proposition 1), in their notation, that is $L^0(n \vee p)$ in ours. *Pervasive,* $d^2 \asymp np$. A non-vanishing fraction of cells participate, with loadings touching a non-vanishing fraction of genes at non-vanishing log fold change: the regime of the major axes of variation and the hypothesis of Theorem 2(d). *Where the results sit.* Theorem 1 bounds the error of the fitted signal by $r(n+p)\log(np)$ with no strength assumption; against Chen and Li's Proposition 1 its rate is minimax up to the logarithm. Assumption 5 asks the weakest true factor to clear that error, $d_{L^0}^2 \gg \omega_{n,p}$, and Theorem 2 then places every true direction in the span of the selected columns. In the calibrated sense above, Assumption 5 is sharp up to the logarithm for the recovery of the canonical pairs (Proposition 2) and, when $p \lesssim n$, for the span statement of Theorem 2, where the score floor $n \vee \sqrt{np}$ coincides with $n \vee p$. When $p \gg n$ the score floor is $\sqrt{np}$ and Assumption 5, at $\omega_{n,p} \asymp p\log(np)$, exceeds it by the factor $\sqrt{p/n}\,\log(np)$: the price of controlling the scores through the Frobenius error of the whole signal, which the loading side sets. Remark 11 compares the strength conditions with those of Chen and Li. Factors in the second band are captured only in aggregate, the bound of Theorem 2 degrading as $\|\widehat{M}^{\perp}_{\widehat{S}} - M^{0\perp}\|_F/d_k$ grows, while Theorem 1 holds regardless. Whether the selected set coincides with the set of true factors is a separate question (Remark 10) and is not claimed.

E.5 Selection: no true factor is lost

Throughout this section E denotes any number with $\|\widehat{M}^{\perp} - M^{0\perp}\|_F \le E$. Theorem 1 supplies, on an event of probability at least $1 - C(np)^{-1}$, $E^2 \le C[\omega_{n,p} + \lambda_g n] \le C'\omega_{n,p}$ for fixed λ_g, since $n \le \omega_{n,p}$, with C' depending on λ_g as the constants of this section may; we write $E^2 \le C\omega_{n,p}$ for this bound. Write $E_S := \|\widehat{M}^{\perp}_{\widehat{S}} - M^{0\perp}\|_F$, so that $E_S \le E + L\sqrt{n\rho_\star/c_\star}$ by Lemma 6(b), with $\rho_\star \le 1$ under Assumption 4. Only two properties of the floor are used: $E^2 \le C\omega_{n,p}$, and $n = o(\omega_{n,p})$ together with $\omega_{n,p} = o(np)$; both hold for $\omega_{n,p} = r(n+p)\log(np)$ and for the alternative floor $\sqrt{np\,\omega_{n,p}}$.

Theorem 2 (no true factor is lost; restated). *Let Assumptions 1, 2, 3 and 4 hold, and let E, E_S be as above. Then, deterministically:*

(i) (span) $\max_{k\le L^0}\operatorname{dist}(u_k,\widehat{\mathcal{U}}) \le \|(I-\widetilde{P})P_{\mathcal{U}^0}\|_{\mathrm{op}} \le \|(I-\widetilde{P})P_{\mathcal{U}^0}\|_F \le E_S/d_{L^0}$;

(ii) (count) if $E_S < d_{L^0}$, then $\widetilde{P}$ is injective on $\mathcal{U}^0$ and $L_{\mathrm{active}} \ge \dim\widehat{\mathcal{U}} \ge L^0$;

(iii) (rank concentration) $\sum_{k>L^0}\sigma_k^2(\widehat{M}^\perp) \le E^2$, *and if* $E < d_1$ *then* $\sum_{k>L^0}\sigma_k^2(\widehat{M}^\perp)/\|\widehat{M}^\perp\|_F^2 \le E^2/(d_1-E)^2$.

Consequently, under Assumptions 1 to 5, as $n,p\to\infty$:

(a) $\mathbb{P}(L_{\mathrm{active}} \ge L^0) \to 1$;

(b) $\max_{k\le L^0}\operatorname{dist}(u_k,\widehat{\mathcal{U}}) \xrightarrow{p} 0$;

(c) $\sum_{k>L^0}\sigma_k^2(\widehat{M}^\perp)/\|\widehat{M}^\perp\|_F^2 \xrightarrow{p} 0$, *where (c) needs only $d_1^2/\omega_{n,p}\to\infty$ in place of Assumption 5.*

Part (d) is restated and proved after the proof of (a) to (c).

Proof of Theorem 2. (i) For $\ell\in\widehat{\mathcal{S}}$, $P_X^\perp m_{\widehat{W}_{\cdot\ell}}\in\widehat{\mathcal{U}}$, so $(I-\widetilde{P})P_X^\perp m_{\widehat{W}_{\cdot\ell}}=0$ and, summing the rank-one terms, $(I-\widetilde{P})\widehat{M}^\perp_{\widehat{\mathcal{S}}}=0$. Hence $(I-\widetilde{P})M^{0\perp}=(I-\widetilde{P})(M^{0\perp}-\widehat{M}^\perp_{\widehat{\mathcal{S}}})$ and $\|(I-\widetilde{P})M^{0\perp}\|_F\le E_S$. Right multiplication by a matrix with orthonormal columns preserves the Frobenius norm ($\|A(V^0)^\top\|_F^2=\operatorname{Tr}(A(V^0)^\top V^0A^\top)=\|A\|_F^2$), so

$$\begin{aligned}\|(I-\widetilde{P})M^{0\perp}\|_F^2 = \|(I-\widetilde{P})U^0D^0\|_F^2 &= \sum_{k\le L^0} d_k^2\,\|(I-\widetilde{P})u_k\|^2\\ &\ge\ d_{L^0}^2\sum_{k\le L^0}\|(I-\widetilde{P})u_k\|^2 = d_{L^0}^2\,\|(I-\widetilde{P})P_{\mathcal{U}^0}\|_F^2,\end{aligned}$$

the last equality by the same identity applied to $P_{\mathcal{U}^0}=U^0(U^0)^\top$. Thus $\|(I-\widetilde{P})P_{\mathcal{U}^0}\|_F\le E_S/d_{L^0}$, and $\operatorname{dist}(u_k,\widehat{\mathcal{U}})=\|(I-\widetilde{P})P_{\mathcal{U}^0}u_k\|\le\|(I-\widetilde{P})P_{\mathcal{U}^0}\|_{\mathrm{op}}\le\|(I-\widetilde{P})P_{\mathcal{U}^0}\|_F$.

(ii) For a unit $u\in\mathcal{U}^0$, $\|(I-\widetilde{P})u\|\le\|(I-\widetilde{P})P_{\mathcal{U}^0}\|_{\mathrm{op}}\le E_S/d_{L^0}<1$, so $\|\widetilde{P}u\|^2=1-\|(I-\widetilde{P})u\|^2>0$: $\widetilde{P}$ restricted to $\mathcal{U}^0$ has trivial kernel, $\dim\widetilde{P}\mathcal{U}^0=L^0$, and $\widetilde{P}\mathcal{U}^0\subseteq\widehat{\mathcal{U}}$ gives $\dim\widehat{\mathcal{U}}\ge L^0$; $\widehat{\mathcal{U}}$ is spanned by $|\widehat{\mathcal{S}}|=L_{\mathrm{active}}$ vectors, so $L_{\mathrm{active}}\ge\dim\widehat{\mathcal{U}}\ge L^0$.

(iii) $M^{0\perp}$ has rank L^0, so by the Eckart–Young theorem (Eckart & Young, 1936)

$$\sum_{k>L^0}\sigma_k^2(\widehat{M}^\perp)=\min_{\operatorname{rank}R\le L^0}\|\widehat{M}^\perp-R\|_F^2\le\|\widehat{M}^\perp-M^{0\perp}\|_F^2\le E^2,$$

and $\|\widehat{M}^\perp\|_F\ge\|M^{0\perp}\|_F-\|\widehat{M}^\perp-M^{0\perp}\|_F\ge d_1-E$ because $\|M^{0\perp}\|_F^2=\sum_k d_k^2\ge d_1^2$; divide. (In the raw frame, $\widehat{M}-\widehat{M}^\perp=P_X\widehat{M}+P_X^\perp\widehat{M}P_V$ has rank at most $\min(H,L)+G$, so $\widehat{M}$ lies within E of a matrix of rank at most $L^0+\min(H,L)+G$ and $\sum_{k>L^0+\min(H,L)+G}\sigma_k^2(\widehat{M})\le E^2$.)

(a)–(c) On the event of Theorem 1, under Assumption 4, $E_S\le\sqrt{C\omega_{n,p}}+L\sqrt{n/c_\star}$, so

$$\frac{E_S}{d_{L^0}}\ \le\ \sqrt{C\,\frac{\omega_{n,p}}{d_{L^0}^2}}+\frac{L}{\sqrt{c_\star}}\sqrt{\frac{n}{\omega_{n,p}}\cdot\frac{\omega_{n,p}}{d_{L^0}^2}}\ \longrightarrow\ 0$$

under Assumption 5, since $n=o(\omega_{n,p})$: (b) follows from (i), and (a) from (ii) once the right-hand side is below 1. For (c), $E/d_1\le E/d_{L^0}\le\sqrt{C\omega_{n,p}/d_{L^0}^2}\to0$ and (iii) applies; only $E/d_1\to0$ was used. □

Reading Theorem 2. Clause (i) says every direction of the true factor space lies within angle $\arcsin(E_S/d_{L^0})$ of the span of the selected, design-residualized score columns; (ii) that at least

L^0 columns are selected; (iii) that the fitted design-residual signal is of rank L^0 up to a vanishing share of its energy. Not claimed: that each selected column is a true factor or that $L_{\text{active}} = L^0$ (the count is one-sided; no over-selection bound is claimed, see the remark on exactness); that any single column is close to any single factor (Theorem 2(d) is the statement of that kind that holds, and Remark 7 shows where it stops); or anything about the score embedding: Theorem 1 controls $\widehat{\eta}$ and hence the denoised means, (iii) says that the fitted design-residual signal carries a vanishing share of its energy beyond rank L^0, and a selected column enters the embedding at its own variance whatever its product energy. The rule that keeps every column also satisfies (i) and (ii); what the deployed rule drops is controlled by Lemma 6(a), whose per-column budget $n/c_\star$ is the content specific to the unit threshold and is what makes $\widehat{M}^{\perp}_{\widehat{\mathcal{S}}}$, not only $\widehat{M}^{\perp}$, consistent.

Theorem 2(d) (a distinct selected factor for every true factor; restated). *Let Assumptions 1, 2, 3 and 4 hold, let E be as above, and suppose*

$$E + L\sqrt{n/c_\star} \;\le\; \tfrac{1}{4}\, d_{L^0} \qquad \text{and} \qquad \tau \in (0,1] \;\text{ satisfies }\; L^{3/2} C_W C_\alpha\, \tau \sqrt{np} \;\le\; \tfrac{1}{4}\, d_{L^0}.$$

(a) Then there is an injective map $\pi : \{1, \dots, L^0\} \to \widehat{\mathcal{S}}$ with $|\operatorname{corr}(u_k, \widehat{W}_{\cdot\pi(k)})| \ge \tau$ for every $k \le L^0$.

(b) Under Assumption 5$^+$ the second condition holds for all large n, p with the constant $\tau := \sqrt{c_h}/(4L^{3/2} C_W C_\alpha)$ and the first holds on the event of Theorem 1 for all large n, p; hence under Assumptions 1, 2, 3, 5$^+$, 4, with probability tending to one such a π exists with this τ.

(c) Under Assumption 5 alone the same holds with $\tau_{n,p} := d_{L^0}/(4L^{3/2} C_W C_\alpha \sqrt{np})$ in place of τ, which lies in $(0, 1]$ because $d_{L^0} \le \|M^{0\perp}\|_{\text{op}} \le \|W^0\|_F \|\alpha^0\|_F \le L C_W C_\alpha \sqrt{np}$ gives $\tau_{n,p} \le 1/(4\sqrt{L})$, and $\tau_{n,p} \to 0$ when $d^2_{L^0} = o(np)$.

Proof. The graph. Let $\mathcal{G}_\tau$ be the bipartite graph with left vertex set $\{1, \dots, L^0\}$, right vertex set $\widehat{\mathcal{S}}$, and an edge (k, ℓ) iff $|u_k^\top \widetilde{W}_{\cdot\ell}| \ge \tau \|\widetilde{W}_{\cdot\ell}\|$. Every $\ell \in \widehat{\mathcal{S}}$ has $\widehat{V}_\ell > 0$, hence $\widetilde{W}_{\cdot\ell} \ne 0$, so an edge means exactly $|\operatorname{corr}(u_k, \widehat{W}_{\cdot\ell})| \ge \tau$ (§E.1.2). By Hall's marriage theorem (Hall, 1935), $\mathcal{G}_\tau$ has a matching saturating the left side, which is the π claimed, provided every nonempty $K' \subseteq \{1, \dots, L^0\}$ satisfies $|N(K')| \ge |K'|$, where $N(K') := \{\ell \in \widehat{\mathcal{S}} : (k, \ell) \text{ is an edge for some } k \in K'\}$.

Decomposition. Fix K' and $P_{K'} := \sum_{k \in K'} u_k u_k^\top$. Since $P_X^\perp u_k = u_k = P_1^\perp u_k$, $P_{K'} P_X^\perp = P_{K'} = P_{K'} P_1^\perp$, so

$$P_{K'} \widehat{M}^{\perp} = \sum_{\ell=1}^{L} A_\ell, \qquad A_\ell := \big(P_{K'} \widetilde{W}_{\cdot\ell}\big)\big(\widehat{\alpha}_{\ell\cdot} P_V^\perp\big), \qquad P_{K'} M^{0\perp} = \sum_{k \in K'} d_k u_k v_k^\top,$$

each A_ℓ of rank at most one, and $P_{K'} M^{0\perp}$ with exactly $|K'|$ nonzero singular values, the smallest at least d_{L^0}.

First application of Weyl's inequality. For any two matrices and every m, $|\sigma_m(B + \Delta) - \sigma_m(B)| \le \|\Delta\|_{\text{op}}$ (Horn & Johnson, 2012, Sec. 7.3). With $B = P_{K'} M^{0\perp}$ and $\Delta = P_{K'}(\widehat{M}^{\perp} - M^{0\perp})$, $\|\Delta\|_{\text{op}} \le E$, so $\sigma_{|K'|}(P_{K'} \widehat{M}^{\perp}) \ge d_{L^0} - E$.

The columns outside $N(K')$. Split $\{1, \dots, L\} = N(K') \cup (\widehat{\mathcal{S}} \setminus N(K')) \cup \widehat{\mathcal{S}}^c$. For $\ell \in \widehat{\mathcal{S}} \setminus N(K')$ no $k \in K'$ is joined to ℓ, so $\|P_{K'} \widetilde{W}_{\cdot\ell}\|^2 = \sum_{k \in K'} (u_k^\top \widetilde{W}_{\cdot\ell})^2 < |K'| \tau^2 \|\widetilde{W}_{\cdot\ell}\|^2 \le L\tau^2 C_W^2 n$, using $\|\widetilde{W}_{\cdot\ell}\| \le \|m_{\widehat{W}_{\cdot\ell}}\| \le C_W \sqrt{n}$ (entrywise box), and with $\|\widehat{\alpha}_{\ell\cdot} P_V^\perp\| \le \|\widehat{\alpha}_{\ell\cdot}\| \le C_\alpha \sqrt{p}$ (entrywise box on the loadings), $\|A_\ell\|_F = \|P_{K'} \widetilde{W}_{\cdot\ell}\| \|\widehat{\alpha}_{\ell\cdot} P_V^\perp\| \le \sqrt{L}\, \tau C_W C_\alpha \sqrt{np}$. For $\ell \in \widehat{\mathcal{S}}^c$, Lemma 6(a) and $\|P_{K'} \widetilde{W}_{\cdot\ell}\| \le \|\widetilde{W}_{\cdot\ell}\|$ give $\|A_\ell\|_F \le \|\widetilde{W}_{\cdot\ell}\| \|\widehat{\alpha}_{\ell\cdot}\| \le \sqrt{n\rho_\star/c_\star} \le \sqrt{n/c_\star}$, the last step by Assumption 4(ii). At most L columns lie outside $N(K')$, so $\|\sum_{\ell \notin N(K')} A_\ell\|_{\text{op}} \le \sum_{\ell \notin N(K')} \|A_\ell\|_F \le L^{3/2} \tau C_W C_\alpha \sqrt{np} + L\sqrt{n/c_\star}$.

Second application of Weyl's inequality, and the rank count. With $B = P_{K'}\widehat{M}^{\perp}$ and $\Delta = -\sum_{\ell \notin N(K')} A_\ell$,

$$\sigma_{|K'|}\Big(\sum_{\ell \in N(K')} A_\ell\Big) \;\geq\; d_{L^0} - E - L\sqrt{n/c_\star} - L^{3/2}\tau C_W C_\alpha \sqrt{np} \;\geq\; d_{L^0} - \tfrac{1}{4}d_{L^0} - \tfrac{1}{4}d_{L^0} \;=\; \tfrac{1}{2}d_{L^0} \;>\; 0$$

under the two hypotheses. A matrix with $|K'|$ nonzero singular values has rank at least $|K'|$; a sum of $|N(K')|$ matrices of rank at most one has rank at most $|N(K')|$. Hence $|N(K')| \geq |K'|$, Hall's condition holds, and π exists.

The constants. Under Assumption 5$^+$, $d_{L^0} \geq \sqrt{c_h np}$ for all large n, p, so $\tau = \sqrt{c_h}/(4L^{3/2}C_W C_\alpha)$ gives $L^{3/2}C_W C_\alpha \tau\sqrt{np} = \frac{1}{4}\sqrt{c_h np} \leq \frac{1}{4}d_{L^0}$; on the event of Theorem 1, $E + L\sqrt{n/c_\star} \leq \sqrt{C\omega_{n,p}} + L\sqrt{n/c_\star} = o(d_{L^0})$ by Assumption 5 (implied by Assumption 5$^+$) and $n = o(\omega_{n,p})$. Under Assumption 5 alone, $\tau_{n,p}$ makes the second hypothesis an equality and the first holds as before. □

What the proof used, and what it did not. It used Assumption 4(ii), through Lemma 6(a), for the unselected columns; the entrywise boxes C_W, C_α for the selected non-neighbors; and Theorem 1 for E. It used no gap and no ordering of the d_k: the only property of $\{u_k\}$ that entered is that they form an orthonormal basis of $\mathcal{U}^0$, since $\sigma_{|K'|}(P_{K'}M^{0\perp}) \geq d_{L^0}$ holds for $P_{K'}$ the projector onto any $|K'|$-dimensional subspace of $\mathcal{U}^0$ (for a unit y there, $\|(M^{0\perp})^\top y\| = \|D^0(U^0)^\top y\| \geq d_{L^0}$). The statement therefore holds for any orthonormal basis of $\mathcal{U}^0$, in particular for any orthonormal basis of a block of tied strengths. Not asserted: that $\pi(k)$ carries factor k alone or predominantly, that the matching is unique, or that the angle vanishes; the last cannot be asserted because the fitted basis is not the canonical one (Remark 9).

Remark 6 (the matched correlation under an exact rotation). Suppose the selected columns are an orthogonal rotation of the canonical factors, $\widehat{W}_{\cdot\widehat{\mathcal{S}}} = \sqrt{n}\,U^0 R$ with $R \in \mathbb{R}^{L^0 \times L^0}$ orthogonal, so that $\mathrm{corr}(u_k, \widehat{W}_{\cdot\ell}) = R_{k\ell}$ and $D := R \circ R$ is doubly stochastic. Then some permutation π has $\min_k D_{k\pi(k)} \geq 4/(L^0+1)^2$, so the bottleneck matching has $|\,\mathrm{corr}(u_k, \widehat{W}_{\cdot\pi(k)})| \geq 2/(L^0+1)$ for every k. Indeed, if the graph $\{(k,\ell) : D_{k\ell} \geq t\}$ has no perfect matching, Hall's theorem gives q rows whose neighborhood has at most $q-1$ columns; with $q' \geq L^0 - q + 1$ the number of columns outside it, the $q \times q'$ block has all entries below t, and comparing the rows' mass q with the mass $< qq't$ inside the block and $\leq L^0 - q'$ in the remaining columns (column sums are 1) gives $t > (q + q' - L^0)/(qq')$. Writing $q' = L^0 + 1 - q + j$, $j \geq 0$, the right-hand side is $(1+j)/(q(L^0 + 1 - q + j))$, nondecreasing in j (its derivative has the sign of $q(L^0+1-q) - q \geq 0$) and at $j = 0$ at least $4/(L^0+1)^2$, since $q(L^0+1-q) \leq (L^0+1)^2/4$. So a level $t \leq 4/(L^0+1)^2$ always admits a perfect matching (the Frobenius–König theorem on the level set; Lovász & Plummer, 2009). The bound is attained: for $L^0 = 3$, the orthogonal R with rows $(1/\sqrt{2}, 1/2, 1/2)$, $(1/\sqrt{2}, -1/2, -1/2)$, $(0, 1/\sqrt{2}, -1/\sqrt{2})$ has D with rows $(.5, .25, .25)$, $(.5, .25, .25)$, $(0, .5, .5)$, and the bottleneck value of D, $\max_\pi \min_k D_{k\pi(k)}$, is exactly $1/4$. In this idealization the worst-case τ of Theorem 2(d) is replaced by $2/(L^0+1)$.

Remark 7 (obstructions: what fails, and why the constant is worst-case). Three configurations mark the boundaries of Theorem 2(d). Each reproduces the truth exactly ($\widehat{M}^{\perp} = M^{0\perp}$, $E = 0$), lies in Φ when the singular vectors are delocalized ($\max_i |u_{k,i}| \leq C_W/(2\sqrt{n})$, $d_k \max_j |v_{k,j}| \leq C_\alpha\sqrt{n}/2$), and has every column selected under Assumption 4(ii), by $\mathrm{SNR}_\ell \geq \widehat{V}_\ell(1 + c_\star \widehat{R}_\ell)$: in O1 and O3 every column has $\widehat{V}_\ell \geq 1$ and $\widehat{R}_\ell \geq \|\widehat{\alpha}_{\ell\cdot}\|^2 > 0$; in O2, with a, b as defined there, both columns have $\widehat{R}_\ell \geq d_1^2/(a^2 n) \geq d_1^2/n$, which exceeds $1/c_\star$ once $d_1^2 > n/c_\star$, with $\widehat{V}_1 = 1$ and $\widehat{V}_2 = b^2$. Take intercept-only designs, so $u_k \perp \mathbf{1}_n$, $v_k \perp \mathbf{1}_p$.

O1 (pure-cross columns; every bilinear matching fails). $L^0 = 2$, $d_1 = d_2 = d$, $\widehat{W} = \sqrt{n}\,[u_1 + u_2,\ u_1,\ u_2]$, $\widehat{\alpha} = (d/\sqrt{n})\,[(v_1 + v_2)^\top;\ -v_2^\top;\ -v_1^\top]$. Then $\widehat{W}\widehat{\alpha} = d(u_1 v_1^\top + u_2 v_2^\top)$, the cross terms cancelling, while the bilinear overlaps $B_{k\ell} := |u_k^\top \widetilde{W}_{\cdot\ell}|\,|\widehat{\alpha}_{\ell\cdot} v_k|$ form the matrix with rows $(d, 0, 0)$, $(d, 0, 0)$: both factors have column 1 as their only bilinear witness, and no injective assignment with nonzero bilinear overlaps exists. The score correlations are $(1/\sqrt{2}, 1, 0)$ for $k = 1$ and $(1/\sqrt{2}, 0, 1)$ for $k = 2$, so the matching $1 \mapsto 2$, $2 \mapsto 3$ exists with correlation one. A small score angle bounds

the whole k-row of a column's rank-one term, cross terms included; a small bilinear overlap does not, which is why part (d) is stated for correlations.

O2 (a hidden factor; no angle bound below pervasiveness). $L^0 = 1$, $d_1^2 = o(np)$, a unit $z \perp \{u_1, \mathbf{1}_n\}$, $a^2+b^2 = 1$, $\widehat{W} = \sqrt{n}\,[au_1+bz,\ bz]$, $\widehat{\alpha} = (d_1/(a\sqrt{n}))\,[v_1^\top;\ -v_1^\top]$. Then $\widehat{W}\widehat{\alpha} = d_1 u_1 v_1^\top$; the loading box requires only $a \geq d_1 \max_j |v_{1,j}|/(C_\alpha\sqrt{n})$, which permits $a \to 0$ exactly when $d_1^2 = o(np)$; $\mathrm{corr}(u_1, \widehat{W}_{\cdot 1}) = a$ and $\mathrm{corr}(u_1, \widehat{W}_{\cdot 2}) = 0$; yet $u_1 = (\widehat{W}_{\cdot 1} - \widehat{W}_{\cdot 2})/(a\sqrt{n}) \in \widehat{\mathcal{U}}$, so the span clause of Theorem 2 holds with distance zero. At the pervasive scale, $d_1 = \sqrt{c_h np}$, the loading box forces $a \geq \sqrt{c_h p}\,\max_j |v_{1,j}|/C_\alpha$, which is of order $\sqrt{c_h}/C_\alpha$ for delocalized v_1 ($\max_j |v_{1,j}| \asymp p^{-1/2}$): this is what Assumption 5$^+$ adds, and $\tau_{n,p}$ matches a at $d_1^2 = o(np)$.

O3 (an argmax-defeating rotation). $L^0 = 3$, $\widehat{W} = \sqrt{n}\,U^0 R$, $\widehat{\alpha} = R^\top D^0 (V^0)^\top/\sqrt{n}$ with the R of Remark 6: $\mathrm{corr}(u_k, \widehat{W}_{\cdot \ell}) = R_{k\ell}$ and $B_{k\ell} = d_k R_{k\ell}^2$, so factors 1 and 2 both choose column 1 under the maximal-correlation and the maximal-overlap rule, whatever the strengths, while $1 \mapsto 1$, $2 \mapsto 2$, $3 \mapsto 3$ has correlations $1/\sqrt{2}$, $1/2$, $1/\sqrt{2}$. Part (d) is an existence statement for this reason.

What the analysis resolves. Under Assumption 4(ii) a column with $\widehat{R}_\ell \asymp p$ has $s^2_{W_{i\ell}} \asymp 1/p$, so its score-block entropy alone is $\frac{1}{2}\sum_i \log(1+\widehat{F}_{i\ell}) \asymp \frac{n}{2}\log p$ and its loading row costs of order $p \log n$: O1 has objective value lower by about $\frac{1}{2}(n \log p + p \log n)$ than the two-column configuration with the same product, and O2 lower by an amount of order $n \log(1/a) + n \log(d_1^2/n) + p \log n$ than the one-column one. These are of order $(n+p)\log(np) = O(\omega_{n,p})$, the resolution of the analysis: Theorem 1 bounds $\mathcal{L}(\widetilde{\phi}(\widehat{\nu}))$ from below with a loss of order $\omega_{n,p}$ and uses nothing finer, so the analysis does not resolve the objective below $O(\omega_{n,p})$ and cannot exclude such configurations. Maximality itself indicates more: comparing $\widehat{\phi}$ with the configuration that merges the two columns of O2 into one (same product; a score-block saving of order $(n+p)\log(np)$ against a change of at most $C\lambda_g n$ in the graph term and of order n in the mean-field second-order terms) indicates that O2 with $a \to 0$ is excluded at the level of orders, constants untracked. The obstructions are obstructions to the method of proof, which uses Assumption 4 through (i) and (ii) only; this is why τ is a worst-case constant, and Remark 6 shows how far a finer argument could move it.

Remark 8 (a per-factor witness). For each $k \leq L^0$, $u_k^\top \widehat{M}^\perp v_k \geq d_k - E$ and the left side is $\sum_{\ell \leq L}(u_k^\top \widehat{W}_{\cdot \ell})(\widehat{\alpha}_{\ell \cdot} v_k)$, so some column $\ell(k)$ has bilinear overlap at least $(d_k - E)/L$, hence centered product energy at least $(d_k - E)^2/L^2$, and Lemma 6(c) puts it in $\widehat{\mathcal{S}}$ once $d_k > E + L\sqrt{n\rho_\star/c_\star}$. Witnesses need not be distinct (O1), so this adds nothing to the count clause of Theorem 2; Theorem 2(d) is the distinct-column statement that does hold.

E.6 Factor-level recovery of the canonical basis

Definition 1 (factor read-off)**.** Let $\widehat{M}^\perp = \sum_k \widehat{d}_k \widehat{u}_k \widehat{v}_k^\top$ be the singular value decomposition of the design-residual fitted signal. The *read-off* of factor k is $\widehat{W}^{\mathrm{ro}}_{\cdot k} = \sqrt{n}\,\widehat{u}_k$, $\widehat{\alpha}^{\mathrm{ro}}_{k \cdot} = (\widehat{d}_k/\sqrt{n})\widehat{v}_k^\top$, the principal factors of the quantity the model fits, computable from any fitted model.

Lemma 10 (Wedin's $\sin\Theta$ theorem)**.** *Let $B, \widehat{B} \in \mathbb{R}^{n \times p}$ have singular value decompositions partitioned as $B = U_1\Sigma_1 V_1^\top + U_2 \Sigma_2 V_2^\top$ and $\widehat{B} = \widehat{U}_1\widehat{\Sigma}_1\widehat{V}_1^\top + \widehat{U}_2\widehat{\Sigma}_2\widehat{V}_2^\top$, with U_1 and $\widehat{U}_1$ of the same width. If $\sigma_{\min}(\widehat{\Sigma}_1) \geq a + \bar{\delta}$ and $\sigma_{\max}(\Sigma_2) \leq a$ for some $a \geq 0$, $\bar{\delta} > 0$, then, with $\Delta = \widehat{B} - B$, $\max\{\|(I - \widehat{U}_1\widehat{U}_1^\top)U_1\|_{\mathrm{op}}, \|(I - \widehat{V}_1\widehat{V}_1^\top)V_1\|_{\mathrm{op}}\} \leq \|\Delta\|_{\mathrm{op}}/\bar{\delta}$.*

Proof. This is the generalized $\sin\Theta$ theorem of Wedin (1972), stated there for unitarily invariant norms: with the residuals $B\widehat{V}_1 - \widehat{U}_1\widehat{\Sigma}_1 = -\Delta\widehat{V}_1$ and $B^\top\widehat{U}_1 - \widehat{V}_1\widehat{\Sigma}_1 = -\Delta^\top\widehat{U}_1$, whose operator norms are at most $\|\Delta\|_{\mathrm{op}}$, it bounds the sines of the largest principal angles between $\mathrm{col}(U_1)$ and $\mathrm{col}(\widehat{U}_1)$, and between $\mathrm{col}(V_1)$ and $\mathrm{col}(\widehat{V}_1)$, by $\|\Delta\|_{\mathrm{op}}/\bar{\delta}$; for subspaces of equal dimension $\|(I - \widehat{U}_1\widehat{U}_1^\top)U_1\|_{\mathrm{op}}$ is that sine, which is symmetric in the two subspaces. □

Proposition 2 (canonical factors by the singular value decomposition)**.** *Let Assumptions 1 to 6 hold and let E be as in §E.5; the two displays below are deterministic given E, and the vanishing clause*

holds on the event of Theorem 1. Fix $k \le L^0$ with $E \le \delta d_k/8$ and set $s_k := \operatorname{sign}(\widehat{u}_k^\top u_k)$. Then

$$\max\left\{\|\widehat{u}_k - s_k u_k\|,\ \|\widehat{v}_k - s_k v_k\|\right\} \le \frac{4\sqrt{2}\,E}{\delta\, d_k},$$

$$\frac{1}{n}\big\|\widehat{W}^{\mathrm{ro}}_{\cdot k} - s_k\sqrt{n}\,u_k\big\|^2 + \frac{1}{p}\big\|\widehat{\alpha}^{\mathrm{ro}}_{k\cdot} - s_k\tfrac{d_k}{\sqrt{n}}v_k^\top\big\|^2 \le \frac{32E^2}{\delta^2 d_k^2} + \frac{2E^2}{np}\Big(1+\frac{32}{\delta^2}\Big).$$

Both bounds tend to zero in probability for every $k \le L^0$: the read-off recovers the canonical factors up to sign, matched in order of strength. If Assumption 6 holds only between blocks of consecutive strengths, the same argument bounds the principal angles between the corresponding block subspaces, and nothing is claimed within a block.

Proof. Let $\Delta := \widehat{M}^\perp - M^{0\perp}$, so $\|\Delta\|_{\mathrm{op}} \le E$; for $m \le L^0$ let P^0_m, $\widehat{P}_m$ be the orthogonal projectors onto the spans of $u_1,\dots,u_m$ and of $\widehat{u}_1,\dots,\widehat{u}_m$ ($P^0_0 = \widehat{P}_0 = 0$), and Q^0_m, $\widehat{Q}_m$ their right-side analogues. Weyl's inequality gives $\widehat{d}_m \ge d_m - E$ for every m.

Top-m subspaces, $m \in \{k-1, k\}$, $m \ge 1$. Assumption 6 and $d_m \ge d_k$ give $d_m - d_{m+1} \ge \delta d_m \ge \delta d_k$. Apply Lemma 10 with $B = M^{0\perp}$, $\widehat{B} = \widehat{M}^\perp$, block 1 the top m singular triples, $a := d_{m+1}$, $\bar\delta := \delta d_k - E > 0$: $\sigma_{\min}(\widehat{\Sigma}_1) = \widehat{d}_m \ge d_m - E \ge d_{m+1} + \delta d_k - E = a + \bar\delta$ and $\sigma_{\max}(\Sigma_2) = d_{m+1} = a$ (with $d_{L^0+1} = 0$; the trailing zero singular values of $M^{0\perp}$ belong to block 2). Hence, using $E \le \delta d_k/2$,

$$\|(I - \widehat{P}_m)P^0_m\|_{\mathrm{op}} \ \vee\ \|(I - \widehat{Q}_m)Q^0_m\|_{\mathrm{op}} \ \le\ \frac{E}{\delta d_k - E} \ \le\ \frac{2E}{\delta d_k}.$$

From the top-m subspaces to the k-th direction. Since $I - \widehat{u}_k\widehat{u}_k^\top = (I - \widehat{P}_k) + \widehat{P}_{k-1}$,

$$\sin\angle(\widehat{u}_k, u_k) = \|(I - \widehat{u}_k\widehat{u}_k^\top)u_k\| \le \|(I - \widehat{P}_k)u_k\| + \|\widehat{P}_{k-1}u_k\|.$$

The first term is at most $\|(I - \widehat{P}_k)P^0_k\|_{\mathrm{op}}$ since $u_k = P^0_k u_k$. For the second, $u_k \perp \mathrm{range}(P^0_{k-1})$ gives $\widehat{P}_{k-1}u_k = \widehat{P}_{k-1}(I - P^0_{k-1})u_k$, and $\|\widehat{P}_{k-1}(I - P^0_{k-1})\|_{\mathrm{op}} = \|(I - P^0_{k-1})\widehat{P}_{k-1}\|_{\mathrm{op}} = \|(I - \widehat{P}_{k-1})P^0_{k-1}\|_{\mathrm{op}}$, by adjoints and because both are the sine of the largest principal angle between two subspaces of dimension $k-1$ (the term is absent for $k = 1$). Hence $\sin\angle(\widehat{u}_k, u_k) \le 4E/(\delta d_k) \le 1/2$, and likewise $\sin\angle(\widehat{v}_k, v_k) \le 4E/(\delta d_k)$.

Signs. $\|\widehat{u}_k - s_k u_k\|^2 = 2(1 - |\widehat{u}_k^\top u_k|) \le 2(1 - |\widehat{u}_k^\top u_k|^2) = 2\sin^2\angle(\widehat{u}_k, u_k) \le 32E^2/(\delta d_k)^2$. The same sign works for $\widehat{v}_k$: $\widehat{d}_k\, v_k^\top\widehat{v}_k = v_k^\top(\widehat{M}^\perp)^\top\widehat{u}_k = d_k\, u_k^\top\widehat{u}_k + v_k^\top\Delta^\top\widehat{u}_k$, where $|u_k^\top\widehat{u}_k| \ge \cos\angle(\widehat{u}_k, u_k) \ge \sqrt{3}/2$ and $|v_k^\top\Delta^\top\widehat{u}_k| \le E \le d_k/8$, so $v_k^\top\widehat{v}_k$ has the sign of $u_k^\top\widehat{u}_k$ ($\widehat{d}_k > 0$), and $\|\widehat{v}_k - s_k v_k\|^2 \le 2\sin^2\angle(\widehat{v}_k, v_k) \le 32E^2/(\delta d_k)^2$.

Scores, loadings, asymptotics. $\frac{1}{n}\|\widehat{W}^{\mathrm{ro}}_{\cdot k} - s_k\sqrt{n}\,u_k\|^2 = \|\widehat{u}_k - s_k u_k\|^2$, and

$$\frac{1}{p}\Big\|\widehat{\alpha}^{\mathrm{ro}}_{k\cdot} - s_k\tfrac{d_k}{\sqrt{n}}v_k^\top\Big\|^2 = \frac{\|\widehat{d}_k\widehat{v}_k - s_k d_k v_k\|^2}{np} \le \frac{2}{np}\Big[(\widehat{d}_k - d_k)^2 + d_k^2\|\widehat{v}_k - s_k v_k\|^2\Big] \le \frac{2E^2}{np}\Big(1+\frac{32}{\delta^2}\Big),$$

by Weyl's inequality for the first term. On the event of Theorem 1, $E^2 \le C\omega_{n,p}$, and Assumption 5 gives $\omega_{n,p}/d_k^2 \le \omega_{n,p}/d_{L^0}^2 \to 0$, so $E \le \delta d_k/8$ holds eventually and the first bound vanishes; the second vanishes because $\omega_{n,p} = o(np)$ under Assumption 2. For a block of consecutive strengths separated from its neighbors by the gap of Assumption 6, apply the first step to the two top-m subspaces bounding the block and subtract projectors as above. □

Reading Proposition 2. The proposition identifies the canonical factors of the signal at the price one should expect: the factor's strength must clear the floor of Assumption 5, and it must be separated from its neighbors' (tied factors are identifiable only as a block, as in any spectral method). The same holds for the read-off of the selected product $\widehat{M}^\perp_{\widehat{S}}$, with E_S in place of E (Lemma 6(b)). What the proposition does not do is make the canonical basis the right target for biological reading: a canonical direction is defined by variance geometry, and nothing in the model makes generative

programs orthogonal or energy-ranked, so where strengths are close a single canonical direction can combine several programs at once. The paper's biological analyses therefore read the fitted columns, the basis the sparsity prior scales and selects factor by factor, and audit that reading empirically (cross-seed reproducibility, marker-removal robustness, technical screening; Section 4). Theorem 2(d) is the bridge between the two bases: each canonical direction has a distinct selected fitted column at bounded correlation, bounded rather than vanishing because the two bases differ, as the next remark measures.

Remark 9 (which basis, empirically). On the atlas reference fit (Section 4), the fitted and read-off bases differ substantially: the Hungarian-matched |cosine| between the fitted score columns and the read-off of Definition 1 averages 0.48 over the 24 active factors, with minimum 0.13 and fourteen of the 24 below 0.5. The design-residual frame costs little of what matters here: it retains 63% of the row-centered signal's energy and removes no canonical direction, all 24 retaining design-residual strength $\widehat{d}_k^2$ above $n + p$ (the smallest $\widehat{d}_k^2/(n + p)$ is 3.3). Both bases reproduce across ten seeds, at a mean matched loading cosine over the 45 seed pairs of 0.999 for the fitted basis and 0.998 for the read-off, with no single matched factor below 0.99 and 0.98 respectively (the matching of Appendix C.2), and every analysis that depends only on the fitted signal (denoised means, decoder differential expression) is identical in the two by construction. The disagreement is larger where the theorem tolerates it: agreement tracks the relative singular-value gap at rank correlation 0.51, so canonical directions with nearly tied strengths — which merge distinct rare-type programs that the fitted basis keeps apart — are where the two bases part company. This is the empirical form of the discussion above: the canonical basis is identifiable, and it is the wrong target.

E.7 Exactness at the unit threshold

Remark 10 (exactness at the unit threshold: an account, not a theorem). Theorem 2 bounds under-selection only. The reason no over-selection bound is claimed is one of orders, recorded at the level of a surrogate objective that no proof uses.

The implicit per-factor cost. Substituting the variance fixed points $s^2_{W_{i\ell}} = (1 + \widehat{F}_{i\ell})^{-1}$ of Lemma 5 and $s^2_{\alpha_{\ell j}} = (\lambda_\ell^{-2} + F_{\alpha,\ell j})^{-1}$, where $F_{\alpha,\ell j}$ is the q-averaged observed information for $\alpha_{\ell j}$, into the KL terms and the variance leakage of the likelihood (Lemma 3) turns the two blocks of an alive column ℓ into Gaussian-evidence forms:

$$\tfrac{1}{2}\|m_{W\cdot\ell}\|^2 + \tfrac{1}{2}\sum_i \log(1 + \widehat{F}_{i\ell}), \qquad \frac{\|m_{\alpha_{\ell\cdot}}\|^2}{2\lambda_\ell^2} + \tfrac{1}{2}\sum_j \log\left(1 + \lambda_\ell^2 F_{\alpha,\ell j}\right) + \log(1 + \lambda_\ell^2).$$

Profiling the scale at its stationary value $\lambda_\ell^2 \approx \|m_{\alpha_{\ell\cdot}}\|^2/p$ and the split of scale between scores and loadings, under a homogeneous curvature κ per entry, leaves $\frac{n+p}{2} + \frac{n}{2}\log\left(1+\kappa e_\ell(1+2\lambda_g r_\ell)/n\right) + \frac{p}{2}\log\left(1 + \kappa e_\ell/p\right)$, where $e_\ell := \|m_{W\cdot\ell}\|^2\|m_{\alpha_{\ell\cdot}}\|^2$ is the column's product energy and $r_\ell = \|(I - A)m_{W\cdot\ell}\|^2/\|m_{W\cdot\ell}\|^2$ its roughness on the graph; the same profile fixes $\widehat{V}_\ell = 1/(1 + 2\lambda_g r_\ell)$ and $\mathrm{SNR}_\ell \approx \kappa e_\ell/n$, so $\widehat{V}_\ell$ alone is not a signal variance. Any alive column therefore costs $\Theta(n + p)$ up to logarithms, which a true pervasive factor pays out of a likelihood gain of order $d_k^2 \asymp np$, while a surplus column on the top singular pair of the standardized residual gains at most $\kappa\sigma_1^2/2$ with $\sigma_1 \approx \sqrt{n} + \sqrt{p}$ (Bai & Yin, 1993), again $\Theta(n+p)$. Penalty and spurious gain share the order $N \vee J$ of Remark 5 in Chen & Li (2022), whose consistency condition requires the penalty increment to dominate it; no such domination holds here, and exactness at the unit threshold is not proved.

The corrected ledger. In the exact-KL Gaussian surrogate (independent standardized entries, unit curvature, one surplus column with mean-field Gaussian factors and the priors of Eq. (2), means $m = xu_1$, $\mu = yv_1$ on the top singular pair), the objective change relative to collapse is

$$\Delta\mathcal{L} = \sigma_1 xy - \tfrac{1}{2}QR - \mathrm{KL}_W - \mathrm{KL}_\alpha - \log(1 + \lambda^2),$$

where $Q = x^2 + ns^2$ and $R = y^2 + p\tau^2$ with (s^2, τ^2) the common variances, $\mathrm{KL}_W = \frac{1}{2}[x^2 + ns^2 - n - n\log s^2]$ and $\mathrm{KL}_\alpha = \frac{1}{2}[R/\lambda^2 - p + p\log\lambda^2 - p\log\tau^2]$. The term $\frac{1}{2}x^2 = n\widehat{V}/2$ is the prior-mean part of KL_W; costing the score block by its information terms alone omits it and reverses the sign. Maximizing at $n = 2000$ and $\sigma_1 = \sqrt{n} + \sqrt{p}$ (logarithmic grid in (x^2, y^2), with (s^2, τ^2, λ^2) at their stationary values), the alive branch's supremum is negative at every aspect ratio, $\Delta\mathcal{L}/n = -0.026, -0.058, -0.082, -0.114, -0.156$ at $p/n = 0.1, 0.5, 1, 2, 4$: the edge chaser

collapses. The alive branch first wins at σ_1 between 1.09 and 1.11 times the edge (an energy margin of 19 to 23%), appearing through a first-order transition with $\mathrm{SNR} = 0.94, 1.67, 2.56, 3.83, 5.21$, between 0.54 and 0.66 times $(1+\sqrt{p/n})^2$, of the same order as the largest inactive and smallest active SNR across the benchmark, 0.92 and 2.33.

Why this is not a theorem. The surplus column of the deployed estimator sees noise plus the estimation error of the fitted true block, which Theorem 1 controls only at the order $r(n+p)\log(np)$ of the gain itself; the negative-binomial residual is heteroskedastic, with an edge the analysis does not locate; and every available constant is a worst case over the boxes, while the margin is of order one. A theorem would need a null-SNR upper bound of order $(1+\sqrt{p/n})^2$ for the global maximizer, uniform over the law of the residual.

Splitting. In the same account the per-column penalty is concave in e_ℓ and vanishes at 0, hence subadditive: more alive columns for a fixed product cost more, and rotating a pervasive column into a surplus column of energy e' by an angle φ costs $\frac{n+p}{2}\log(1+\sin^2\varphi\, e_\ell/e')$, which dominates every competing $O(n+p)$ term once $\sin^2\varphi \gg (n+p)/(np)$. At this level no true factor is split across alive columns, and exactness reduces to the survival of a chaser.

E.8 Relation to Chen and Li (2022)

Remark 11 (Relation to Chen and Li (2022)). Theorems 1 and 2 follow the pattern of Chen and Li's results for the constrained joint maximum likelihood estimator of a generalized latent factor model (Chen & Li, 2022). The dictionary is as follows.

Model and nuisance. Their natural parameter is $m_{ij} = d_j + A_j^\top F_i$, with d_j their intercept (unrelated to our strengths d_k), so their signal M^* has rank K^*+1: K^* factors plus a column intercept that every candidate model contains. Our design terms $X\beta + \gamma V^\top$ are that intercept generalized to a nuisance of rank at most $H+G$ that every candidate model also contains; their complete-data model is the case $X = \mathbf{1}_n$, $G = 0$, no offset, canonical link and known dispersion. Their constraint set $(\|F_i\|^2+1)^{1/2} \le C$, $(d_j^2 + \|A_j\|^2)^{1/2} \le C$ is our Φ, their Assumption 2 is Assumption 2, their $K_{\max}$ is L and their $K^* \le K_{\max}$ is $L^0 \le L$.

The weakest factor. Their Theorem 3 is stated through $\sigma_{K^*+1}(M^*)$, the least nonzero singular value of the full signal, intercept included. In their Lemma 9 it enters as a proxy: a fit with fewer than K^* factors has rank at most K^* and so lies at Frobenius distance at least $\sigma_{K^*+1}(M^*)$ from M^* by the Eckart–Young theorem (Eckart & Young, 1936), whereas the distance that matters is the smaller one, to the fits that also carry the intercept. Our $d_{L^0} = \sigma_{L^0}(P_X^\perp \eta^0 P_V^\perp)$ is that gap exactly for the design at hand, since $d_k = \min_{\beta,\gamma}\sigma_k(\eta^0 - X\beta - \gamma V^\top)$ by Lemma 1; it is also the only version that is well defined, $\sigma_k(W^0\alpha^0)$ moving along the gauge orbit of Proposition 1 while d_k does not.

Estimator and rule; one-sidedness. They refit for every $K \le K_{\max}$ and minimize $-2\hat{l}_K + v_{\mathrm{JIC}}(K)$; the increment $u_{\mathrm{JIC}} = v_{\mathrm{JIC}}(K) - v_{\mathrm{JIC}}(K-1)$ (their v and u) is a design choice, and their Theorem 3 needs $N \vee J = o(u_{\mathrm{JIC}})$ against over-selection and $u_{\mathrm{JIC}} = o(\sigma^2_{K^*+1}(M^*))$ against under-selection. We fit once, at the ceiling L, with no penalty on the count; a column is kept by $\mathrm{SNR}_\ell > 1$ and switched off by the sparsity prior's collapse. The place of u_{JIC} is taken by an implicit per-factor cost, the Kullback–Leibler terms of a live score column and loading row, which the exactness remark (Remark 10) computes, at the surrogate level, to be of order $n+p$: the order of the likelihood gain a surplus column extracts from noise, their Remark 5 ($2(\hat{l}_K - \hat{l}_{K-1}) = O_p(N \vee J)$ for $K > K^*$). Their over-selection lemma (their Lemma 4) needs the increment to grow faster than $N \vee J$; ours is not dialed above $n \vee p$, so that lemma has no analogue here and Theorem 2 is one-sided, $\mathbb{P}(L_{\mathrm{active}} \ge L^0) \to 1$ with span inclusion in place of $\mathbb{P}(\widehat{K} = K^*) \to 1$. Their under-selection lemma has an analogue: where a fit of too low a rank trails the truth by a multiple of $\sigma^2_{K^*+1}(M^*)$ once the error bound holds, here a true direction left outside the selected span would have to be carried by unselected columns, whose design-residual energy is at most $L\sqrt{n/c_\star}$ in total, and the proof divides by d_{L^0}.

Error bounds and the logarithm. Their Theorem 1 bounds the squared Frobenius error of every fit with $K^* \le K \le K_{\max}$ by a constant times $K_{\max}(N \vee J)$, with no logarithm, and their Proposition 1 shows this order to be minimax. The proof is not a covering argument: for a canonical link the likelihood difference is $\langle Z, M - M^*\rangle$ minus a Bregman remainder with deterministic curvature, the

linear term obeys $\langle Z, \Delta \rangle \le \|Z\|_{\mathrm{op}} \|\Delta\|_* \le \sqrt{\operatorname{rank} \Delta}\, \|Z\|_{\mathrm{op}} \|\Delta\|_F$ with $\operatorname{rank} \Delta \le K + K^* + 2$, and the only probabilistic input is $\|Z\|_{\mathrm{op}} = O_p(\sqrt{N} + \sqrt{J})$ (their Lemma 3). Theorem 1 has the same skeleton: strong concavity of the smoothed negative binomial log-likelihood in the mean, deterministic for every dispersion and every $Y \ge 0$; a score at the truth equal to the centred count times a weight $\bar{w}_{ij}(\nu_j) \in (0, 1]$; and $\operatorname{rank}(\widehat{\eta} - \eta^0) \le H + G + 2L$. Its floor $\omega_{n,p} = r(n+p)\log(np)$ carries a logarithm for two reasons, neither removed by the duality view. First, the comparison configuration of the basic inequality belongs to the variational family: each of its order $r(n+p)$ live variances equals $(n+p)^{-1}$ and costs $\frac{1}{2}\log(n+p)$ of entropy, so the ledger carries $\omega_{n,p}$ before any probability enters, and a point estimate has no such ledger. Second, the score weights depend on the p fitted dispersions, which they take as known (their Remark 6 plugs in one estimate from the largest model); the supremum over them is paid by the parameter net of Lemma 9, of log-cardinality of order $\omega_{n,p}$, and duality bounds the linear term only at a fixed dispersion. We do not claim their log-free rate: their logarithm is a penalty they choose, ours sits in the estimation rate.

Selection conditions. Their Corollary 1 takes $v_{\mathrm{JIC}}(K) = K(N \vee J)\log(N \wedge J)$ under $(N \vee J)\log(N \wedge J) = o(\sigma^2_{K^*+1}(M^*))$; Assumption 5 is $\omega_{n,p} = o(d^2_{L^0})$. The two conditions have the same order whenever $\log(n \wedge p) \asymp \log(np)$, so for every polynomial aspect ratio $p \asymp n^a$, and ours is heavier only when one dimension is sub-polynomial in the other. They obtain the upper bound $\widehat{K} \le K^*$, which we do not. We state the span clause of Theorem 2(b) and the reconstruction from the selected part in Theorem 1, which needs no strength condition; they state neither, and both concern the fitted columns that downstream analysis consumes.

Per-factor statements. Their Remark 1 makes no per-factor claim: the factors are not identifiable under rotation, so every bound is for M and for K. Theorem 2(d) is consistent with that position. It is an existence statement about the fitted columns, that the selected set contains distinct columns correlated with distinct true score directions above a fixed level, and it identifies nothing: no column is asserted to estimate a factor, the level τ is worst case over the boxes, and what is identifiable per factor is the canonical basis of $M^{0\perp}$, recovered under a gap condition by a spectral read-off rather than by the fitted columns.

E.9 Scope: what is deliberately not claimed

For calibration, the claims this appendix does *not* make, and why.

(i) *Exact recovery of the selected set.* Theorem 2 bounds under-selection and Theorem 2(d) assigns a distinct selected column to every true factor; neither bounds over-selection. The implicit per-factor cost of the sparsity prior under mean-field inference is of the same order as a noise factor's likelihood gain (Remark 10), so exactness at the unit threshold is not proved; the remark records the surrogate account, its margin, and what a proof would need.

(ii) *Factor naming below the floor of Assumption 5*: such factors are captured in aggregate (the subspace statement degrades continuously) but are not individually licensed for interpretation; programs of tens of genes are below any method's resolution at these scales.

(iii) *A minimax converse* at the detection floor: the three-band picture of Remark 5 is stated as calibration, not as a proved lower bound.

(iv) *Constants.* Theorem 1 is quantitative, but every constant is a worst case over the feasible boxes. In particular $c_\star$ of Lemma 5 is a minimum of the negative-binomial curvature over the boxes, and the budget $n/c_\star$ it yields for an unselected column is not claimed to be numerically informative, which is why the computable $\rho_\star$ accompanies the selection diagnostics. Assumption 5 is a sufficient condition proved under these constants, not a threshold the reported factors are claimed to clear.

(v) *The canonical basis as the biological target*: Proposition 2 identifies the canonical factors, and that license is deliberately not transferred to per-factor biological reading, since a canonical direction can combine several generative programs (the discussion and Remark 9 there); no analogous convergence is asserted for the fitted columns, whose reading Section 4 audits empirically.

(vi) *The mathematical estimator.* The guarantees concern a maximizer of $\mathcal{L}$ over Φ under exact variational expectations, through the two consequences of maximality named in Assumption 4; the deployed optimizer's stop-gradient, and why Assumption 4(i) is therefore named

rather than derived, are discussed in §E.1. Optimization itself is not analyzed, and the variational family is analyzed as deployed: no claim that the fitted variances approximate posterior variances is made, Lemma 5 characterizes the SNR denominator directly, and the family's implicit penalty enters only in Remark 10, as an account.

(vii) *Recovery of the unprojected signal $W^0\alpha^0$*: no estimator attains it, by Proposition 1. Every statement here is for the design-residual signal $P_X^{\perp}(\cdot)P_V^{\perp}$ of §E.1.2; the deployed embedding may still carry $\mathrm{col}(X)$ directions, which the depth screening of Section 4 is what detects. Nor is $\mathrm{col}(X)$ purely technical: a cell covariate such as the detection rate carries biological variation — on the benchmark data it tracks cell-cycle state and cell size — so the projection removes some signal along with the nuisance, and the estimand is narrower than "$W^0\alpha^0$ with the technical part removed".

(viii) *Covariates computed from the counts*: the designs are fixed (Assumption 3), so a covariate computed from Y, such as the deployed detection rate, is outside the theorems as stated (see the discussion of Assumption 3 in §E.2).