LEARNING MULTIOME REPRESENTATIONS ACROSS INDIVIDUALS

Anonymous authors
Paper under double-blind review

ABSTRACT

Paired single-cell multiome assays measure gene expression and chromatin accessibility in the same cell. Contrastive learning on the two measurements has therefore become a standard way to learn cell representations. In population studies, however, such representations are useful only if they transfer across donors. We show that the contrastive loss alone does not ensure this. When negatives are sampled from all donors, the contrastive loss rewards features that identify the donor. When negatives are drawn within donors, each donor's representation can be rotated separately without changing the loss. Removing donor information altogether resolves neither issue without also discarding phenotype. We propose MINUET, which combines within-donor contrastive learning with a prediction task shared across donors and a donor penalty applied within groups of donors that share biological covariates. We show that these objectives determine the representation up to one shift for each group of donors connected by shared covariates. Applied to six paired cohorts, MINUET achieves the best average rank among nine methods in integration, cross-donor and cross-dataset transfer, and phenotype prediction.

1 INTRODUCTION

Paired single-cell multiome assays measure gene expression and chromatin accessibility in the same cell (Chen et al., 2019; Ma et al., 2020). The pairing provides supervision without labels: a model can learn by identifying which RNA and ATAC profiles come from the same cell (van den Oord et al., 2018; Chen et al., 2020). Yet the studies that produce these data are rarely about individual cells. They sample donors and ask which cell states and phenotypes are shared with donors not yet seen. This paper asks when contrastive learning on paired cells yields representations that transfer to such donors.

Matching paired cells well does not ensure transfer. When negatives are sampled from all donors, any feature that identifies the donor lowers the contrastive loss while carrying no information about cell state, a shortcut solution (Robinson et al., 2021). Sampling negatives within each donor removes this shortcut, but the loss then never compares two donors. Rotating the representations of one donor, both assays together, leaves every similarity unchanged, and each donor can be rotated differently. A model can therefore reach a low contrastive loss while representing the same cell state differently in different donors. Removing all donor information goes too far. Donors differ in disease, age and genotype, and a representation that is fully invariant to the donor loses the phenotypes a population study aims to find (Zhao et al., 2019).

MINUET addresses each difficulty in turn. It draws negatives only from the anchor's donor. It trains one prediction head, shared by all donors and both assays, to predict a common target from the representation. It penalizes the dependence between the representation and donor identity only among donors with the same biological covariates, such as disease status (Aliee et al., 2023). After training, the frozen encoders map each query cell on its own, without fine-tuning. We prove that, even for nonlinear encoders, these objectives determine the representation up to one shift for each connected component of a graph that links donors sharing a covariate group (Theorem 3.1). The shared head removes the rotations, the donor penalties remove the shifts between donors of the same group, and how well the graph is connected determines how stable the result is (Singer, 2011; Fiedler, 1973). The training losses are small exactly when the representation is close to this solution, and the same check applies to a new donor using only its unlabeled cells.

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.

Synthetic experiments test each failure mode separately. On 636,509 cells from 110 donors in six cohorts, we compare MINUET with seven published methods and with its own prediction target used as the representation, under four protocols: observed-cell integration, transfer to held-out donors, transfer to a held-out dataset, and phenotype prediction for held-out donors. Every protocol treats donors rather than cells as the unit of evaluation (Zimmerman et al., 2021; Squair et al., 2021). MINUET achieves the best average rank under all four. In population studies, what a representation learns depends on how the self-supervised task is constructed, not only on the encoder.

2 RELATED WORK

Learning from paired views. Paired views have been used for learning since canonical correlation analysis (Gao et al., 2015), and they are the foundation of modern contrastive learning (van den Oord et al., 2018; Chen et al., 2020). Same-cell multiome data have made paired views central to single-cell genomics (Argelaguet et al., 2021). Existing methods use the pairing in three ways. Some learn a joint latent space (Ashuach et al., 2023). Some separate shared from assay-specific variation (Zhang et al., 2026; Gui et al., 2025). Others predict or contrastively align one assay from the other (Hu & Quon, 2024; Xiong et al., 2023a; Niu et al., 2025). Regulatory graphs add feature-level links (Cao & Gao, 2022). When cells are unpaired, correspondence has to be inferred from shared or weakly linked features (Zhu et al., 2023; Chen et al., 2024). In all of these methods, pairing aligns the two assays. Which cells are compared, and across which donors, is left as a detail. We treat that choice as the central question.

Shortcuts, symmetries and invariance. A contrastive loss can be lowered by features that are easy to discriminate and useless downstream (Robinson et al., 2021). Conditioning negatives on a nuisance variable is one remedy (Tsai et al., 2022). Identifiability results show that a learned representation is determined only up to a group of transformations (Zimmermann et al., 2021; Roeder et al., 2021; Daunhawer et al., 2023). In our setting, each donor carries its own symmetry. Rather than resolving each symmetry separately, we tie donors together through a shared predictor. As in synchronization problems (Singer, 2011), how well this works depends on how well the donors are connected. Removing the donor altogether (Ganin & Lempitsky, 2015) is costly when labels and donors are associated (Zhao et al., 2019). We therefore require invariance only within biological groups (Aliee et al., 2023).

Transfer to new individuals. Mapping new samples onto a reference involves a choice between keeping the reference fixed and adapting it to the query (Kang et al., 2021; Lotfollahi et al., 2022; De Donno et al., 2023). Population-scale models describe how samples and conditions vary (Boyeau et al., 2025; Madrigal et al., 2024). Conditional batch correction links batches through shared annotations before fitting pairwise adapters (McCarter, 2026). MINUET adapts nothing at query time. Evaluation raises a separate problem. Multiome benchmarks score integration on the cells used for fitting (Lee et al., 2023; Liu et al., 2025; Wang et al., 2026). However, cells from one donor are not independent replicates (Zimmerman et al., 2021; Squair et al., 2021). We therefore treat the donor, not the cell, as the unit of evaluation.

3 FROM SAME-CELL MATCHING TO CROSS-DONOR REPRESENTATIONS

For cell i, let $(X_i^{\mathrm{R}}, X_i^{\mathrm{A}})$ be its paired RNA and ATAC profiles, D_i its donor, and C_i its biological covariates, such as disease status or brain region. The two encoders map the profiles to $Z_i^{\mathrm{R}}, Z_i^{\mathrm{A}} \in \mathbb{R}^k$. We ask what lowers the contrastive loss, what the loss leaves undetermined, and how to fix this without removing biology. Figure 1 links each answer to an objective of MINUET, Figures 2 and 3 test the answers on synthetic data, and Section A gives the proofs.

Negatives from other donors reward learning the donor. Suppose the representation contains a feature that identifies the donor and raises the similarity of every same-donor pair by a. For one anchor, let $\ell(a)$ be its InfoNCE loss at temperature τ, and let ω_a be the total softmax probability of negatives from other donors. Then (Proposition A.1)

$$\ell'(a) = -\frac{\omega_a}{\tau}, \qquad \sup_{a \geq 0} \big[\ell(0) - \ell(a)\big] = -\log(1 - \omega_0). \tag{1}$$

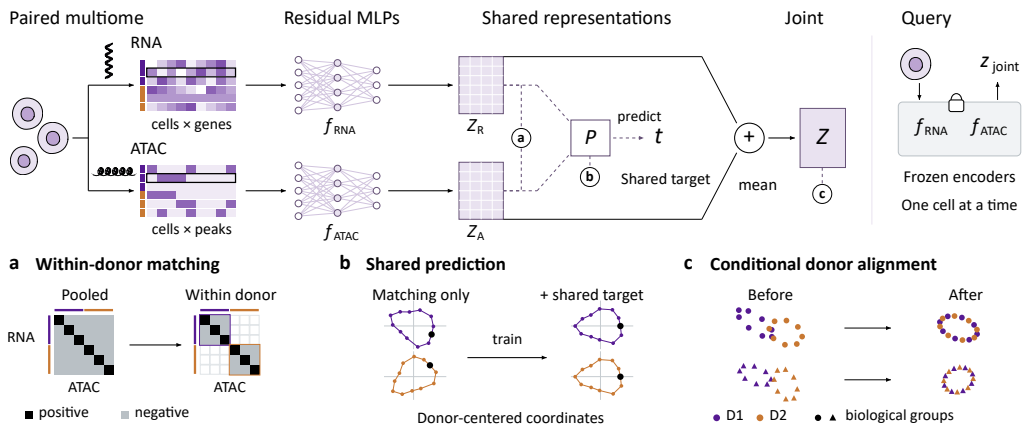

Figure 1: **Minuet.** Two encoders map the RNA and ATAC profiles of a cell to a shared representation, and frozen encoders embed query cells one at a time. **a** Each cell is matched against cells of its own donor rather than of all donors. **b** Matching alone leaves each donor free to rotate; one head, shared by all donors, predicts a common target from either assay and fixes the orientation. **c** Donor penalties align donors within biological groups. Icons adapted from Servier Medical Art (CC BY 3.0) and Bioicons (CC0).

As long as some negatives come from other donors, the loss decreases as a grows, by up to $\log |\mathcal{D}|$ when negatives are drawn evenly from $|\mathcal{D}|$ donors. When all negatives come from the anchor's donor, the feature raises every similarity equally and $\ell'(a) = 0$ (Figure 1a). Sampling negatives within donors, as in conditional contrastive learning (Tsai et al., 2022), therefore removes the *incentive* to learn this shortcut (Robinson et al., 2021). Donor information that is already in the input is handled by the donor penalties below.

Within-donor contrastive learning does not compare donors. Each cell is compared only with cells of its own donor, after that donor's mean is subtracted, so no term of the loss compares two donors.

Proposition 3.1 (Invariance to donor-specific orthogonal transformations)**.** *If a paired loss depends only on inner products between donor-centered, normalized representations within each donor, then rotating and translating each donor separately,*

$$Z_i^m \mapsto Q_d Z_i^m + a_d^m, \qquad D_i = d, \quad Q_d \in \mathsf{O}(k), \tag{2}$$

leaves the loss unchanged (Proposition A.2).

The two assays of a donor must be rotated together, but different donors can be rotated differently. Reflecting one donor, $Q_d = -I$, leaves the loss unchanged, yet a linear classifier of cell state trained on the other donors then reverses its predictions for this donor. An encoder can learn such a reflection whenever the donor can be identified from the input, as it can under batch effects.

A shared prediction head and covariate groups identify the representation. The shared head removes the rotations. One linear head P must predict the same target S from the cells of every donor, so a rotated donor would be predicted incorrectly. Each donor's representation is therefore fixed at AS, where A is the pseudoinverse of P, up to a shift of the donor mean. The donor penalty removes these shifts by requiring donors of the same covariate group to have the same mean. If donors 1 and 2 share a group, and donors 2 and 3 share another, all three are aligned. Alignment thus propagates through a graph that connects donors sharing a group, and each connected component of this graph keeps one common shift. Let $\mathcal{Z}$ be the set of these representations, AS plus one shift per component. Write ε^2 for the prediction loss, B_Z for the variance of donor means within covariate groups, which the penalty drives to zero, and $\mathrm{dist}^2(Z, \mathcal{Z})$ for the mean squared distance from Z to $\mathcal{Z}$.

Theorem 3.1 (Identification, informal)**.** *If the donor means of the target agree within covariate groups, then, up to constants that depend only on P,*

$$\varepsilon^2 + B_Z \;\lesssim\; \mathrm{dist}^2(Z, \mathcal{Z}) \;\lesssim\; \frac{\varepsilon^2 + B_Z}{\lambda_\pi}, \tag{3}$$

where $\lambda_\pi \in (0, 1]$ is the spectral gap of the donor graph, which measures how well connected the graph is. In particular, $Z \in \mathcal{Z}$ if and only if both losses are zero.

The losses are therefore small exactly when the representation is close to $\mathcal{Z}$, and Theorem A.1 gives the constants. In MINUET the target is centered within each donor, so the condition holds for covariates that are constant within a donor, such as disease status. The donors of each group are then all connected to each other, $\lambda_\pi = 1$, and at $\varepsilon = 0$ the bound holds with equality: the penalty B_Z, computed from donor means alone, equals the distance of the whole representation from $\mathcal{Z}$. Covariates that vary within a donor, such as brain region, connect different groups, and λ_π measures how strongly.

The guarantee extends to new donors. The same argument applies to a donor not used in fitting, because its cells are embedded by the same frozen encoders. For a query donor q in group c, let ε_q be the prediction loss on its cells, μ_q its mean representation, and μ_c the mean over reference donors in group c.

Corollary 3.1 (Query donors, informal)**.** *For a donor-centered target and a covariate constant within donor, up to constants that depend only on P,*

$$\mathbb{E}\big[\|Z - AS - \mu_c\|_2^2 \;\big|\; D = q\big] \;\asymp\; \varepsilon_q^2 + \|\mu_q - \mu_c\|_2^2.$$

A new donor is therefore mapped correctly exactly when the head still predicts its target and its mean representation matches the reference mean of its group. Both quantities are computed from unlabeled cells (Corollary A.1).

Removing donor information can remove phenotype. Donors differ in biology, so a representation that is fully invariant to the donor can lose the phenotypes a study aims to find (Zhao et al., 2019). For a phenotype Y with finitely many values, the chain rule of mutual information (Cover & Thomas, 2006, Thm. 2.5.2) gives (Proposition A.3)

$$H(Y \mid Z, C) \;\geq\; I(Y; D \mid C) - I(Z; D \mid C). \tag{4}$$

If the representation keeps less donor information than the phenotype carries, Y cannot be predicted exactly from Z and C. For a phenotype that is constant within each donor, $I(Z; D) = I(Z; Y) + I(Z; D \mid Y)$: removing all donor information, $I(Z; D) = 0$, also removes Y, whereas removing it only within phenotype groups, $I(Z; D \mid Y) = 0$, keeps Y. The donor penalties therefore act within covariate groups (Figure 1c). The representations in Theorem 3.1 remove even less. On $\mathcal{Z}$, Z equals AS plus a shift determined by C, so for every Y (Proposition A.4)

$$I(Z; D \mid C) = I(AS; D \mid C), \qquad I(Z; Y \mid C) = I(AS; Y \mid C).$$

Only differences between donor means are removed, and all information that the target carries about the phenotype is kept. In the phenotype-prediction protocol, Y is not given as a covariate, so all donor means are aligned, yet phenotype information in the cells of each donor, such as their cell-state composition, is kept.

4 MINUET

MINUET is fitted once on paired reference data. After fitting, it maps query cells through frozen encoders. Each objective below corresponds to one result of Section 3, and the top band of Figure 1 summarizes the objectives and the inference path.

4.1 ENCODERS AND JOINT REPRESENTATION

Let $\mathcal{D}_{\mathrm{ref}} = \{(X_i^{\mathrm{R}}, X_i^{\mathrm{A}}, D_i, C_i)\}_{i=1}^n$, where D_i is the donor and C_i collects prespecified biological covariates. For each assay $m \in \{\mathrm{R}, \mathrm{A}\}$, an encoder returns a shared representation $z_i^m \in \mathbb{R}^k$ and a

private variable v_i^m. Decoders reconstruct each assay from v_i^m alone, so that assay-specific variation has a place to go. The shared representation is shaped only by the cross-assay objectives below. For a paired cell, the joint representation is

$$z_i^{\text{joint}} = \tfrac{1}{2}\big(z_i^{\text{R}} + z_i^{\text{A}}\big). \tag{5}$$

Encoders are residual MLPs with BatchNorm. For leave-one-dataset-out (LODO) transfer, we use LayerNorm instead, which normalizes each cell without reference batch statistics. Section B gives architectural and likelihood details.

4.2 Within-donor contrastive learning

Negatives come only from the anchor's donor, whatever their covariates. For cells $\mathcal{I}_d$ sampled from donor d, each assay's representation is centered and normalized within the donor:

$$\bar{z}_d^m = \frac{1}{|\mathcal{I}_d|}\sum_{j\in\mathcal{I}_d} z_j^m, \qquad u_i^m = \frac{z_i^m - \bar{z}_d^m}{\max\{\|z_i^m - \bar{z}_d^m\|_2, \epsilon_{\text{norm}}\}}. \tag{6}$$

For an RNA anchor, its paired ATAC profile is the positive, and the other ATAC profiles of the same donor are the negatives. The reverse direction is symmetric. Let $\ell_{\text{R}\to\text{A}}(i;d)$ denote the temperature-scaled InfoNCE loss on the scores $\langle u_i^{\text{R}}, u_j^{\text{A}}\rangle$. Weighting directions, cells and donors equally gives

$$\mathcal{L}_{\text{pair}} = \frac{1}{|\mathcal{D}|}\sum_{d\in\mathcal{D}}\frac{1}{2|\mathcal{I}_d|}\sum_{i\in\mathcal{I}_d}\big[\ell_{\text{R}\to\text{A}}(i;d) + \ell_{\text{A}\to\text{R}}(i;d)\big]. \tag{7}$$

Because every candidate comes from the anchor's donor, a donor signature raises all scores equally and cancels out (Proposition A.1), and centering removes shifts that a donor adds to all of its cells (Proposition A.2).

4.3 Shared prediction across donors

The within-donor contrastive loss alone does not determine how representations align across donors (Proposition 3.1). To provide a common reference, MINUET constructs a prediction target. For each assay, we center the inputs within donor, fit one PCA basis on all reference cells, and standardize the resulting scores. The target is $t_i = [s_i^{\text{R}}; s_i^{\text{A}}] \in \mathbb{R}^q$. Because the basis is common to all donors, the target expresses every donor in the same coordinates. A single bias-free head $P \in \mathbb{R}^{q\times k}$, shared by both assays and all donors, predicts this target from either representation, $\widehat{t}_i^m = Pz_i^m$. Predictions and targets are centered within each sampled donor, and the loss is

$$\mathcal{L}_{\text{pred}} = \frac{1}{|\mathcal{D}|}\sum_{d\in\mathcal{D}}\frac{1}{2}\sum_{m\in\{\text{R},\text{A}\}}\text{SL1}\big(\text{ctr}_d(\widehat{T}^m), \text{ctr}_d(T)\big), \tag{8}$$

where T and $\widehat{T}^m$ stack the targets and predictions of the sampled cells, and SL1 is the SmoothL1 loss. Because of this centering, the loss ignores donor means, which the next objective handles.

4.4 Donor penalties within covariate groups

The centered prediction loss leaves differences between donor means unconstrained. MINUET therefore penalizes the statistical dependence between the representation and donor identity within each covariate group. For covariate group c, let Z_c stack the joint representations, D_c the one-hot donor indicators, and $H_c = I - \mathbf{1}\mathbf{1}^\top/n_c$. We penalize

$$\mathcal{L}_{\text{donor}} = \frac{1}{|\mathcal{C}|}\sum_{c\in\mathcal{C}}\frac{\|Z_c^\top H_c D_c\|_F^2}{\max\big\{\text{sg}\big(\|Z_c^\top H_c Z_c\|_F\,\|D_c^\top H_c D_c\|_F\big), \epsilon\big\}}, \tag{9}$$

where sg stops gradients through the normalizer, so that the penalty can be lowered only by moving donor means together, not by rescaling the representation. This is linear centered kernel alignment (Kornblith et al., 2019), a normalized linear HSIC (Gretton et al., 2005), between Z_c and donor identity. Its numerator vanishes exactly when all donors in group c share one mean representation, which is $B_Z = 0$ in Theorem 3.1. A second penalty, $\mathcal{L}_{\text{cov}}$, aligns donor covariances, scaled to unit

trace, within the same groups. It is a conditional form of correlation alignment (Sun & Saenko, 2016) and starts after a warm-up period, set by the schedule $\eta(u)$ over updates u (Section B). When no covariate is supplied, $|\mathcal{C}| = 1$, and both terms reduce to their unconditional forms.

The reference-fitting objective is

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \lambda_{\text{pair}}\mathcal{L}_{\text{pair}} + \lambda_{\text{pred}}\mathcal{L}_{\text{pred}} + \lambda_{\text{donor}}\mathcal{L}_{\text{donor}} + \eta(u)\,\lambda_{\text{cov}}\mathcal{L}_{\text{cov}}. \tag{10}$$

Donor labels, covariates, the target and the head are used only in this objective. At inference, MINUET keeps the two encoders, embeds each query cell independently, and averages paired representations by Equation (5). Query mapping involves no parameter updates, no donor-specific parameters, and no centering across query cells.

5 EXPERIMENTS

We test the failure modes of Section 3 first where the ground truth is known, and then in real population data.

5.1 SETUP

Datasets. We analyze 636,509 cells from 110 donors in six published paired multiome cohorts from brain, blood, and synovium (Wang et al., 2025; Mannens et al., 2025; Xiong et al., 2023b; Brown et al., 2024; Gay-Mimbrera et al., 2025; Gupta et al., 2023; Weinand et al., 2024); Table 3 lists their covariates and phenotypes.

Baselines. The comparators are Apollo, scPair, scPairing, scCLIP, IndiSeek, MultiVI, and scGLUE (Zhang et al., 2026; Hu & Quon, 2024; Niu et al., 2025; Xiong et al., 2023a; Gui et al., 2025; Ashuach et al., 2023; Cao & Gao, 2022), and DC-PCA, which uses MINUET's donor-centered PCA target (Timmerman, 2006) as the representation. All methods share inputs, feature sets, splits, seeds, and metrics but keep their own preprocessing, architectures, and training schedules. Two in-house controls, Pooled contrastive and Within-donor contrastive, use MINUET's encoders with a matching loss only. MINUET uses residual MLP encoders with a single configuration shared by all cohorts and protocols (Section B).

Evaluation protocols. Observed-cell integration asks whether a representation organizes the cells it was trained on, scored with the metrics of the scMultiBench benchmark (Liu et al., 2025; Luecken et al., 2022; Liu, 2025). Held-out-donor transfer asks whether it carries over to new individuals; query donors are excluded from preprocessing, feature selection, and training (Ambroise & McLachlan, 2002; Varma & Simon, 2006). Leave-one-dataset-out (LODO) transfer asks whether it carries over to a new study, training on two of Wang, Mannens, and ROSMAP and evaluating the third. Phenotype prediction asks whether it retains donor-level biology it was never shown: ROSMAP Alzheimer disease status, Crohn disease status, and alopecia areata severity are predicted in held-out donors, with the predicted phenotype withheld from MINUET's covariates. In the last three protocols, query cells are excluded from training.

Metrics. ARI and NMI measure how well cell-type labels are preserved, and CiLISI measures donor mixing within each cell type. FOSCTTM, the fraction of cells closer than the true match, measures the alignment error between the two assays at the single-cell level (Liu et al., 2019; Klein et al., 2025); lower is better. It is the only metric that uses the known pairing, so it also penalizes representations that mix donors by collapsing cells together. Cell-type transfer and phenotype prediction are scored by macro-F1, and phenotype prediction also by macro-AUROC. Methods are ranked within each cohort and metric, and ranks are averaged over cohorts with equal weights (Section D).

5.2 SYNTHETIC EXPERIMENTS

Three synthetic generators each create one failure mode: a donor signature shared by both assays, a separate rotation for each donor, and a phenotype whose association with the donor is set by ρ (Section C). The true cell states and phenotypes are used only for evaluation, except that the conditional experiment gives phenotype to the methods as a biological covariate.

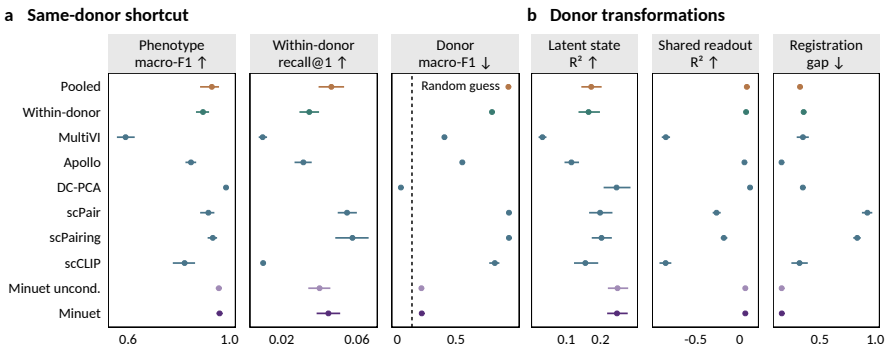

Figure 2: **Synthetic experiments.** Means and 95% intervals over 20 seeds; the dashed line marks random guessing. **a** Same-donor shortcut: phenotype prediction, within-donor matching, and donor prediction. **b** Donor-specific transformations: recovery of the latent state, a shared cross-donor readout, and the registration gap. Minuet uncond. applies the donor penalties without covariates.

Pooled negatives teach the model to recognize donors. Equation (1) predicts that when negatives come from other donors, a model can lower the matching loss simply by recognizing the donor, and the synthetic data confirm this: Pooled contrastive, scPair, and scPairing, which draw negatives from all donors, match cells best, yet the donor can be predicted from their representations almost perfectly (Figure 2a). Good matching is therefore no evidence that a representation has learned biology. Drawing negatives within donors removes the reason to encode the donor, but not the donor signal already present in the data, which Within-donor contrastive still carries. MINUET removes both: it matches cells at least as well as Within-donor contrastive, its donor prediction falls far lower, and it predicts phenotype better.

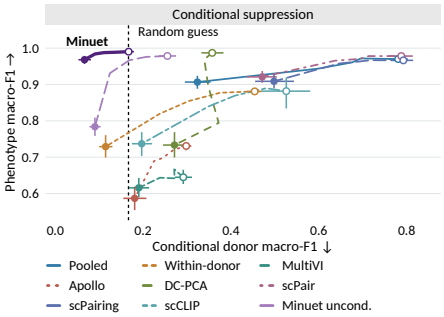

Figure 3: **Conditioning keeps phenotype.** Each curve follows one method as donor–phenotype association ρ rises from 0 (open) to 0.9 (filled). Unconditional penalties lose phenotype as ρ grows, whereas conditional penalties keep it. Means over 20 seeds; the dashed line marks random guessing.

Shared prediction registers donors. Within-donor matching cannot tell a donor's representation from a rotated copy (Proposition 3.1), so a classifier trained on some donors can fail on another. The registration gap measures this failure: how much better a linear map fitted to each donor separately performs than one map shared by all donors. Within-donor contrastive has a large gap, and methods that only align the two assays, such as scPair and scPairing, have the largest (Figure 2b). A prediction head shared by all donors takes away the freedom to rotate each donor separately (Theorem 3.1), and MINUET cuts the gap of Within-donor contrastive by more than half while recovering more of the true cell state.

Conditioning keeps phenotype while removing excess donor information. When donors differ in phenotype, Equation (4) implies that removing all donor information must also remove phenotype. The experiment shows this: as the association ρ between donor and phenotype grows, penalties applied to all donors at once lose more and more phenotype (Figure 3). Penalizing donor differences only within each phenotype group removes just the donor signal that phenotype does not explain, so MINUET leaves the least donor information within groups while keeping the phenotype.

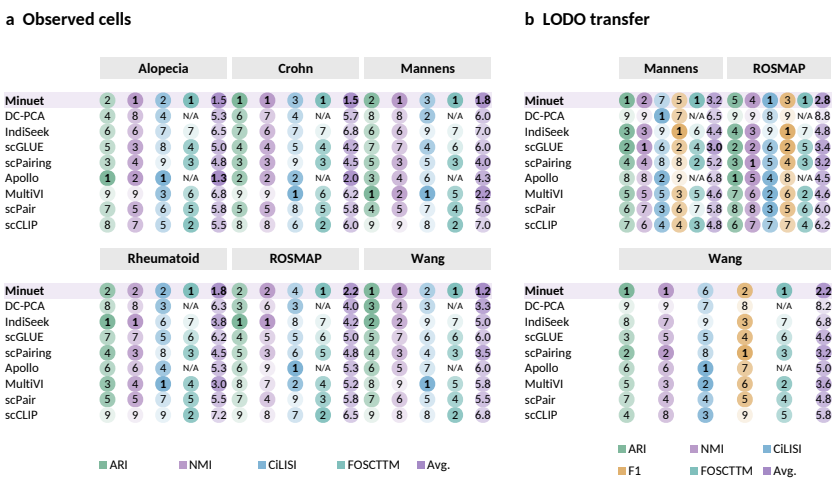

Figure 4: **Main results.** Ranks within each cohort and metric for **a** observed cells and **b** LODO transfer. Lower is better, darker fill marks a better score, and Avg. is the mean rank; N/A marks metrics a method does not define (Section B).

5.3 Main results

Donors are aligned within cell types, and the alignment extends to new donors. In the cells on which it was trained, Minuet has the best average rank and is the first in five of the six cohorts (Figure 4a). It keeps cell types together and matches the two assays of each cell better than the other methods, which is what its objectives are designed to do: pair matching aligns the two assays of each cell, and shared prediction with conditional penalties aligns donors of the same biological group. The advantage holds for query donors excluded from every training step, where Minuet ranks first overall and first or joint first in four of six cohorts (Figure 5). The encoders are frozen and map each query cell without fine-tuning on the new donor, so this alignment is learned by the encoders rather than fitted to the reference donors.

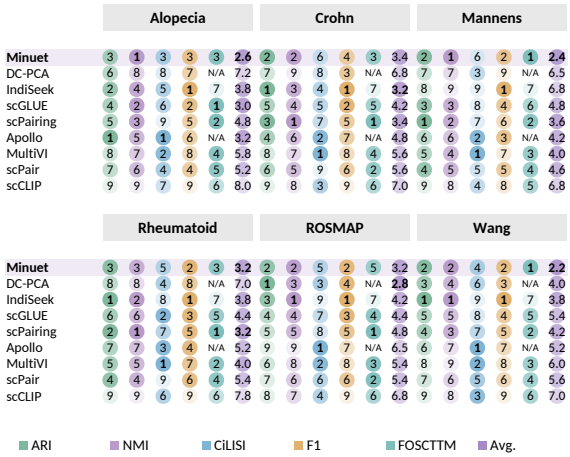

Figure 5: **Held-out donors.** Ranks by cohort for ARI, NMI, CiLISI, cell-type transfer macro-F1 (F1), and FOSCTTM, shown as in Figure 4.

Cell-type structure survives a change of study. When a whole study is left out, Minuet again has the best average rank and ranks first on ROSMAP and Wang (Figure 4b). It matches paired cells best in every held-out study and keeps cell types together best on average, so what it learns from two studies is carried over to a third. For this protocol Minuet uses LayerNorm, which normalizes each cell on its own, because BatchNorm would apply statistics of the training studies to a study that may differ from them (Section 4).

Representations learned without phenotype labels retain phenotype. Although phenotypes are never shown during training, Minuet predicts them best on average: it ranks first on Crohn and alopecia and has the best macro-F1 for Alzheimer disease in ROSMAP (Figure 6). The donor penalties weaken the dependence on the donor rather than remove it, so biology that differs between donors remains available to a classifier. Minuet maintains the state of Alzheimer's disease, whereas methods that remove all donor information predict it at the random-guess level (Figure 16).

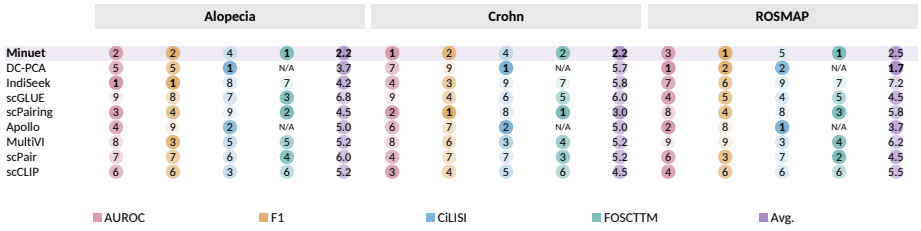

Figure 6: **Phenotype prediction.** Ranks by cohort for phenotype macro-AUROC (AUROC) and macro-F1 (F1), CiLISI, and FOSCTTM, shown as in Figure 4.

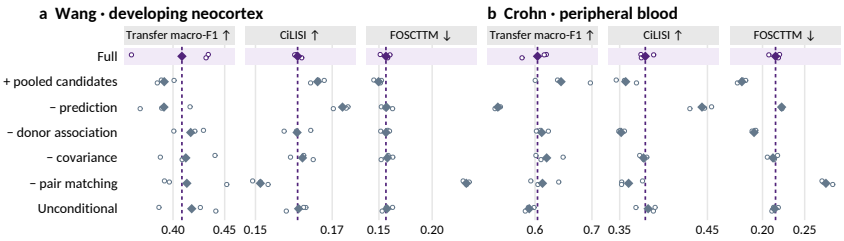

Figure 7: **Ablation study on held-out donors.** Each variant changes one element of Full.

5.4 Ablation

DC-PCA uses the target S itself as the representation and ranks below MINUET in every protocol, because it must see all cells of a query donor to center them and has no space shared by the two assays, whereas MINUET's encoders embed each cell on its own and place both assays in one space. Removing one objective at a time on held-out donors shows what each contributes (Figure 7). Without pair matching, the only term that aligns the two assays of a cell, paired-cell matching becomes worse on Wang and Crohn. Without shared prediction, cell-type transfer on Crohn falls, as expected when each donor's representation is free to rotate. The donor penalties act on differences between donors, which cell-type metrics do not measure. In the synthetic experiments, removing the donor-association penalty leaves more donor information within phenotype groups (Figure 14), and on Crohn a donor-level variable is kept better when it is included among the covariates (Section E).

6 Discussion

A low contrastive loss between the two assays of the same cell is a weak guide to transfer across donors. Negatives from other donors reward recognizing the donor, negatives from the same donor leave each donor free to rotate, and removing all donor information also removes phenotype. MINUET addresses each failure with a separate objective and ranks above its own prediction target in every protocol. Across cohorts, better paired-cell matching alone does not bring better transfer.

The remaining difficulty is deciding what counts as biology. A donor label stands for a bundle of causes: genotype, sex and age, environment, disease, and sample handling (Section E). The conditional penalties keep differences between covariate groups and reduce those within them, so the recorded covariates define the biology that survives. The covariates also carry the confounding of the study: if recruitment links sex to disease, the difference kept between disease groups includes sex, and a biological variable left out of the covariates is treated as a donor effect and removed (Section E).

A cleaner definition would rest on a donor variable that the study design cannot confound. Germline genotype is fixed at conception, the property that Mendelian randomization exploits (Davey Smith & Ebrahim, 2003). Describing donors by genotype, and replacing discrete groups with a similarity graph over continuous descriptors such as genotype and age, can be our next steps for future work. The same graph view implies that transfer depends on which donor groups a reference covers, not only on how many cells it contains.

AI USE STATEMENT

In this work, we have not used generative AI tools for tasks with required disclosure: they were not used to generate data, develop theoretical models or conceptual frameworks, formulate mathematical claims, implement methods, clean or reformat datasets, or interpret results, and translation is not applicable. We used generative AI tools to polish the writing, to draft parts of the text, including sections of the appendix, and to retrieve related work and check bibliographic entries against publisher and preprint records. We have reviewed all AI-assisted work: all drafted text was revised by the authors, and every citation was verified against its source. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REFERENCES

Hananeh Aliee, Ferdinand Kapl, Soroor Hediyeh-Zadeh, and Fabian J. Theis. Conditionally invariant representation learning for disentangling cellular heterogeneity. *arXiv preprint arXiv:2307.00558*, 2023. URL https://arxiv.org/abs/2307.00558.

Christophe Ambroise and Geoffrey J. McLachlan. Selection bias in gene extraction on the basis of microarray gene-expression data. *Proceedings of the National Academy of Sciences*, 99(10): 6562–6566, 2002. doi: 10.1073/pnas.102102699.

Joshua D. Angrist, Guido W. Imbens, and Donald B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996. doi: 10.1080/01621459.1996.10476902.

Ricard Argelaguet, Anna S. E. Cuomo, Oliver Stegle, and John C. Marioni. Computational principles and challenges in single-cell data integration. *Nature Biotechnology*, 39(10):1202–1215, 2021. doi: 10.1038/s41587-021-00895-7. URL https://doi.org/10.1038/s41587-021-00895-7.

Tal Ashuach, Mariano I. Gabitto, Rohan V. Koodli, Giuseppe-Antonio Saldi, Michael I. Jordan, and Nir Yosef. MultiVI: deep generative model for the integration of multimodal data. *Nature Methods*, 20(8):1222–1231, 2023. doi: 10.1038/s41592-023-01909-9. URL https://www.nature.com/articles/s41592-023-01909-9.

Pierre Boyeau, Justin Hong, Adam Gayoso, Martin Kim, José L. McFaline-Figueroa, Michael I. Jordan, Elham Azizi, Can Ergen, and Nir Yosef. Deep generative modeling of sample-level heterogeneity in single-cell genomics. *Nature Methods*, 22(11):2264–2274, 2025. doi: 10.1038/s41592-025-02808-x. URL https://www.nature.com/articles/s41592-025-02808-x.

Margaret Brown, Anne Dodd, Fang Shi, Emily Greenwood, Sini Nagpal, Vasantha L. Kolachala, Subra Kugathasan, and Greg Gibson. Concordant B and T cell heterogeneity inferred from the multi-omic landscape of peripheral blood mononuclear cells in a Crohn's disease cohort. *Journal of Crohn's and Colitis*, 18(12):1939–1956, 2024. doi: 10.1093/ecco-jcc/jjae055. URL https://doi.org/10.1093/ecco-jcc/jjae055.

Zhi-Jie Cao and Ge Gao. Multi-omics single-cell data integration and regulatory inference with graph-linked embedding. *Nature Biotechnology*, 40(10):1458–1466, 2022. doi: 10.1038/s41587-022-01284-4. URL https://www.nature.com/articles/s41587-022-01284-4.

Shuxiao Chen, Bokai Zhu, Sijia Huang, John W. Hickey, Kevin Z. Lin, Michael Snyder, William J. Greenleaf, Garry P. Nolan, Nancy R. Zhang, and Zongming Ma. Integration of spatial and single-cell data across modalities with weakly linked features. *Nature Biotechnology*, 42(7):1096–1106, 2024. doi: 10.1038/s41587-023-01935-0. URL https://www.nature.com/articles/s41587-023-01935-0.

Song Chen, Blue B. Lake, and Kun Zhang. High-throughput sequencing of the transcriptome and chromatin accessibility in the same cell. *Nature Biotechnology*, 37(12):1452–1457, 2019. doi: 10.1038/s41587-019-0290-0. URL https://pubmed.ncbi.nlm.nih.gov/31611697/.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 1597–1607, 2020. URL https://proceedings.mlr.press/v119/chen20j.html.

Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, RI, 1997. ISBN 978-0-8218-0315-8.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006. ISBN 978-0-471-24195-9.

Imant Daunhawer, Alice Bizeul, Emanuele Palumbo, Alexander Marx, and Julia E. Vogt. Identifiability results for multimodal contrastive learning. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=U_2kuqoTcB.

George Davey Smith and Shah Ebrahim. 'Mendelian randomization': can genetic epidemiology contribute to understanding environmental determinants of disease? *International Journal of Epidemiology*, 32(1):1–22, 2003. doi: 10.1093/ije/dyg070.

Carlo De Donno, Soroor Hediyeh-Zadeh, Amir Ali Moinfar, Marco Wagenstetter, Luke Zappia, Mohammad Lotfollahi, and Fabian J. Theis. Population-level integration of single-cell datasets enables multi-scale analysis across samples. *Nature Methods*, 20(11):1683–1692, 2023. doi: 10.1038/s41592-023-02035-2. URL https://www.nature.com/articles/s41592-023-02035-2.

Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2): 298–305, 1973. doi: 10.21136/CMJ.1973.101168. URL https://dml.cz/handle/10338.dmlcz/101168.

Kenji Fukumizu, Arthur Gretton, Xiaohai Sun, and Bernhard Schölkopf. Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems*, volume 20, 2007. URL https://proceedings.neurips.cc/paper_files/paper/2007/hash/3a0772443a0739141292a5429b952fe6-Abstract.html.

Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning (ICML)*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1180–1189, 2015. URL https://proceedings.mlr.press/v37/ganin15.html.

Chao Gao, Zongming Ma, Zhao Ren, and Harrison H. Zhou. Minimax estimation in sparse canonical correlation analysis. *The Annals of Statistics*, 43(5):2168–2197, 2015. doi: 10.1214/15-AOS1332. URL https://arxiv.org/abs/1405.1595.

Jesus Gay-Mimbrera, Pedro Jesús Gómez-Arias, Pablo Álvarez-Heredia, Alexander Batista-Duharte, Irene Rivera-Ruiz, Macarena Aguilar-Luque, Miguel Juan-Cencerrado, Carmen Mochón-Jiménez, Álvaro Cebrián-García, Eloísa Andújar-Pulido, Mónica Pérez-Alegre, Alejandra Pera, and Juan Ruano. Integrated single-cell chromatin and transcriptomic analyses of peripheral immune cells in patients with alopecia areata. *Frontiers in Immunology*, 16:1565241, 2025. doi: 10.3389/fimmu.2025.1565241. URL https://doi.org/10.3389/fimmu.2025.1565241.

Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with Hilbert–Schmidt norms. In *Proceedings of the 16th International Conference on Algorithmic Learning Theory (ALT)*, pp. 63–77, 2005. doi: 10.1007/11564089_7. URL https://doi.org/10.1007/11564089_7.

Yu Gui, Cong Ma, and Zongming Ma. IndiSeek learns information-guided disentangled representations. *arXiv preprint arXiv:2509.21584*, 2025. URL https://arxiv.org/abs/2509.21584.

Anika Gupta, Kathryn Weinand, Aparna Nathan, Saori Sakaue, Martin Jinye Zhang, Accelerating Medicines Partnership RA/SLE Program and Network, Laura Donlin, Kevin Wei, Alkes L.

Price, Tiffany Amariuta, and Soumya Raychaudhuri. Dynamic regulatory elements in single-cell multimodal data implicate key immune cell states enriched for autoimmune disease heritability. *Nature Genetics*, 55(12):2200–2210, 2023. doi: 10.1038/s41588-023-01577-7. URL https://doi.org/10.1038/s41588-023-01577-7.

Hongru Hu and Gerald Quon. scPair: Boosting single cell multimodal analysis by leveraging implicit feature selection and single cell atlases. *Nature Communications*, 15:9932, 2024. doi: 10.1038/s41467-024-53971-2. URL https://www.nature.com/articles/s41467-024-53971-2.

Joyce B. Kang, Aparna Nathan, Kathryn Weinand, Fan Zhang, Nghia Millard, Laurie Rumker, D. Branch Moody, Ilya Korsunsky, and Soumya Raychaudhuri. Efficient and precise single-cell reference atlas mapping with Symphony. *Nature Communications*, 12:5890, 2021. doi: 10.1038/s41467-021-25957-x. URL https://www.nature.com/articles/s41467-021-25957-x.

Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, Aimée Bastidas-Ponce, Perla Cota, Marta Tarquis-Medina, Shrey Parikh, Ilan Gold, Heiko Lickert, Mostafa Bakhti, Mor Nitzan, Marco Cuturi, and Fabian J. Theis. Mapping cells through time and space with moscot. *Nature*, 638(8052):1065–1075, 2025. doi: 10.1038/s41586-024-08453-2.

Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 3519–3529, 2019. URL https://proceedings.mlr.press/v97/kornblith19a.html.

Debbie A. Lawlor, Roger M. Harbord, Jonathan A. C. Sterne, Nic Timpson, and George Davey Smith. Mendelian randomization: using genes as instruments for making causal inferences in epidemiology. *Statistics in Medicine*, 27(8):1133–1163, 2008. doi: 10.1002/sim.3034.

Michelle Y. Y. Lee, Klaus H. Kaestner, and Mingyao Li. Benchmarking algorithms for joint integration of unpaired and paired single-cell RNA-seq and ATAC-seq data. *Genome Biology*, 24(1):244, 2023. doi: 10.1186/s13059-023-03073-x. URL https://doi.org/10.1186/s13059-023-03073-x.

Chunlei Liu. Multi-task benchmarking of single-cell multimodal omics integration methods. Zenodo, 2025. URL https://doi.org/10.5281/zenodo.15385334. Version 1.

Chunlei Liu, Sichang Ding, Hani Jieun Kim, Siqu Long, Di Xiao, Shila Ghazanfar, and Pengyi Yang. Multitask benchmarking of single-cell multimodal omics integration methods. *Nature Methods*, 22(11):2449–2460, 2025. doi: 10.1038/s41592-025-02856-3. URL https://www.nature.com/articles/s41592-025-02856-3.

Jie Liu, Yuanhao Huang, Ritambhara Singh, Jean-Philippe Vert, and William Stafford Noble. Jointly embedding multiple single-cell omics measurements. In *19th International Workshop on Algorithms in Bioinformatics (WABI 2019)*, volume 143 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 10:1–10:13, 2019. doi: 10.4230/LIPIcs.WABI.2019.10. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.WABI.2019.10.

Mohammad Lotfollahi, Mohsen Naghipourfar, Malte D. Luecken, Matin Khajavi, Maren Büttner, Marco Wagenstetter, Žiga Avsec, Adam Gayoso, Nir Yosef, Marta Interlandi, Sergei Rybakov, Alexander V. Misharin, and Fabian J. Theis. Mapping single-cell data to reference atlases by transfer learning. *Nature Biotechnology*, 40:121–130, 2022. doi: 10.1038/s41587-021-01001-7. URL https://www.nature.com/articles/s41587-021-01001-7.

Malte D. Luecken, Maren Büttner, Kridsadakorn Chaichoompu, Anna Danese, Marta Interlandi, Michaela F. Mueller, Daniel C. Strobl, Luke Zappia, Martin Dugas, Maria 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. URL https://www.nature.com/articles/s41592-021-01336-8.

Sai Ma, Bing Zhang, Lindsay M. LaFave, Andrew S. Earl, Zachary Chiang, Yan Hu, Jiarui Ding, Alison Brack, Vinay K. Kartha, Tristan Tay, Travis Law, Caleb Lareau, Ya-Chieh Hsu, Aviv Regev, and Jason D. Buenrostro. Chromatin potential identified by shared single-cell profiling of RNA and chromatin. *Cell*, 183(4):1103–1116.e20, 2020. doi: 10.1016/j.cell.2020.09.056. URL https://pubmed.ncbi.nlm.nih.gov/33098772/.

Ariel Madrigal, Tianyuan Lu, Larisa M. Soto, and Hamed S. Najafabadi. A unified model for interpretable latent embedding of multi-sample, multi-condition single-cell data. *Nature Communications*, 15(1):6573, 2024. doi: 10.1038/s41467-024-50963-0. URL https://www.nature.com/articles/s41467-024-50963-0.

Camiel C. A. Mannens, Lijuan Hu, Peter Lönnerberg, Marijn Schipper, Caleb C. Reagor, Xiaofei Li, Xiaoling He, Roger A. Barker, Erik Sundström, Danielle Posthuma, and Sten Linnarsson. Chromatin accessibility during human first-trimester neurodevelopment. *Nature*, 647(8088): 179–186, 2025. doi: 10.1038/s41586-024-07234-1. URL https://www.nature.com/articles/s41586-024-07234-1.

Calvin McCarter. Confounder-aware feature correction for single-cell batch integration, 2026. URL https://arxiv.org/abs/2608.28849.

Jeffrey Niu, Carlos Vasquez-Rios, and Jiarui Ding. Single-cell multiomics data integration and generation with scPairing. *Cell Reports Methods*, 5(11):101211, 2025. doi: 10.1016/j.crmeth.2025.101211. URL https://doi.org/10.1016/j.crmeth.2025.101211.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009. doi: 10.1017/CBO9780511803161.

Alkes L. Price, Nick J. Patterson, Robert M. Plenge, Michael E. Weinblatt, Nancy A. Shadick, and David Reich. Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics*, 38(8):904–909, 2006. doi: 10.1038/ng1847.

Pia Rautenstrauch and Uwe Ohler. Shortcomings of silhouette in single-cell integration benchmarking. *Nature Biotechnology*, 44(6):954–958, 2026. doi: 10.1038/s41587-025-02743-4. URL https://www.nature.com/articles/s41587-025-02743-4.

Joshua W. Robinson, Li Sun, Ke Yu, Kayhan Batmanghelich, Stefanie Jegelka, and Suvrit Sra. Can contrastive learning avoid shortcut solutions? In *Advances in Neural Information Processing Systems*, volume 34, pp. 4974–4986, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/hash/27934a1f19d678a1377c257b9a780e80-Abstract.html.

Geoffrey Roeder, Luke Metz, and Durk Kingma. On linear identifiability of learned representations. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9030–9039, 2021. URL https://proceedings.mlr.press/v139/roeder21a.html.

Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011. doi: 10.1016/j.acha.2010.02.001. URL https://doi.org/10.1016/j.acha.2010.02.001.

Jordan W. Squair, Matthieu Gautier, Claudia Kathe, Mark A. Anderson, Nicholas D. James, Thomas H. Hutson, Rémi Hudelle, Taha Qaiser, Kaya J. E. Matson, Quentin Barraud, Ariel J. Levine, Gioele La Manno, Michael A. Skinnider, and Grégoire Courtine. Confronting false discoveries in single-cell differential expression. *Nature Communications*, 12(1):5692, 2021. doi: 10.1038/s41467-021-25960-2. URL https://www.nature.com/articles/s41467-021-25960-2.

Baochen Sun and Kate Saenko. Deep CORAL: Correlation alignment for deep domain adaptation. In *Computer Vision – ECCV 2016 Workshops*, pp. 443–450, 2016. doi: 10.1007/978-3-319-49409-8_35. URL https://doi.org/10.1007/978-3-319-49409-8_35.

Marieke E. Timmerman. Multilevel component analysis. *British Journal of Mathematical and Statistical Psychology*, 59(2):301–320, 2006. doi: 10.1348/000711005X67599.

Yao-Hung Hubert Tsai, Tianqin Li, Martin Q. Ma, Han Zhao, Kun Zhang, Louis-Philippe Morency, and Ruslan Salakhutdinov. Conditional contrastive learning with kernel. In *International Conference on Learning Representations*, 2022. URL https://arxiv.org/abs/2202.05458.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. URL https://arxiv.org/abs/1807.03748.

Sudhir Varma and Richard Simon. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 7:91, 2006. doi: 10.1186/1471-2105-7-91.

Li Wang, Cheng Wang, Juan A. Moriano, Songcang Chen, Guolong Zuo, Arantxa Cebrián-Silla, Shaobo Zhang, Tanzila Mukhtar, Shaohui Wang, Mengyi Song, Lilian Gomes de Oliveira, Qiuli Bi, Jonathan J. Augustin, Xinxin Ge, Mercedes F. Paredes, Eric J. Huang, Arturo Alvarez-Buylla, Xin Duan, Jingjing Li, and Arnold R. Kriegstein. Molecular and cellular dynamics of the developing human neocortex. *Nature*, 647(8088):169–178, 2025. doi: 10.1038/s41586-024-08351-7. URL https://www.nature.com/articles/s41586-024-08351-7.

Yixuan Wang, Yimin Fan, Xuesong Wang, Tingyang Yu, Yongshuo Zong, Xinyuan Liu, Gaoyang Zhong, Meitong Liu, Qing Li, Kin Hei Lee, Khachatur Dallakyan, Zhichao Hu, Yaqian Qi, Junjie Huang, Gengjie Jia, Jiao Yuan, Ting-Fung Chan, Xin Gao, Irwin King, and Yu Li. SCMBench: benchmarking domain-specific and foundation models for single-cell multi-omics data integration. *Nature Communications*, 17:5967, 2026. doi: 10.1038/s41467-026-72570-x. URL https://www.nature.com/articles/s41467-026-72570-x.

Kathryn Weinand, Saori Sakaue, Aparna Nathan, Anna Helena Jonsson, Fan Zhang, Gerald F. M. Watts, Majd Al Suqri, Zhu Zhu, Accelerating Medicines Partnership Program: Rheumatoid Arthritis and Systemic Lupus Erythematosus (AMP RA/SLE) Network, Deepak A. Rao, Jennifer H. Anolik, Michael B. Brenner, Laura T. Donlin, Kevin Wei, and Soumya Raychaudhuri. The chromatin landscape of pathogenic transcriptional cell states in rheumatoid arthritis. *Nature Communications*, 15(1):4650, 2024. doi: 10.1038/s41467-024-48620-7. URL https://doi.org/10.1038/s41467-024-48620-7.

Lei Xiong, Tianlong Chen, and Manolis Kellis. scCLIP: Multi-modal single-cell contrastive learning integration pre-training. In *NeurIPS 2023 AI for Science Workshop*, 2023a. URL https://openreview.net/forum?id=KMtM5ZHxct.

Xushen Xiong, Benjamin T. James, Carles A. Boix, Yongjin P. Park, Kyriaki Galani, Matheus B. Victor, Na Sun, Lei Hou, Li-Lun Ho, Julio Mantero, Aine Ni Scannail, Vishnu Dileep, Weixiu Dong, Hansruedi Mathys, David A. Bennett, Li-Huei Tsai, and Manolis Kellis. Epigenomic dissection of Alzheimer's disease pinpoints causal variants and reveals epigenome erosion. *Cell*, 186(20):4422–4437.e21, 2023b. doi: 10.1016/j.cell.2023.08.040. URL https://pubmed.ncbi.nlm.nih.gov/37774680/.

Seyhan Yazar, Jose Alquicira-Hernandez, Kristof Wing, Anne Senabouth, M. Grace Gordon, Stacey Andersen, Qinyi Lu, Antonia Rowson, Thomas R. P. Taylor, Linda Clarke, Katia Maccora, Christine Chen, Anthony L. Cook, Chun Jimmie Ye, Kirsten A. Fairfax, Alex W. Hewitt, and Joseph E. Powell. Single-cell eQTL mapping identifies cell type–specific genetic control of autoimmune disease. *Science*, 376(6589):eabf3041, 2022. doi: 10.1126/science.abf3041.

Xinyi Zhang, G. V. Shivashankar, and Caroline Uhler. Partially shared multi-modal embedding learns holistic representation of cell state. *Nature Computational Science*, 6(3):285–300, 2026. doi: 10.1038/s43588-025-00948-w. URL https://www.nature.com/articles/s43588-025-00948-w.

Han Zhao, Remi Tachet des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7523–7532, 2019. URL https://proceedings.mlr.press/v97/zhao19a.html.

Bokai Zhu, Shuxiao Chen, Yunhao Bai, Han Chen, Guanrui Liao, Nilanjan Mukherjee, Gustavo Vazquez, David R. McIlwain, Alexandar Tzankov, Ivan T. Lee, Matthias S. Matter, Yury Goltsev,

Zongming Ma, Garry P. Nolan, and Sizun Jiang. Robust single-cell matching and multimodal analysis using shared and distinct features. *Nature Methods*, 20(2):304–315, 2023. doi: 10.1038/s41592-022-01709-7. URL https://www.nature.com/articles/s41592-022-01709-7.

Kip D. Zimmerman, Mark A. Espeland, and Carl D. Langefeld. A practical solution to pseudoreplication bias in single-cell studies. *Nature Communications*, 12(1):738, 2021. doi: 10.1038/s41467-021-21038-1. URL https://www.nature.com/articles/s41467-021-21038-1.

Roland S. Zimmermann, Yash Sharma, Steffen Schneider, Matthias Bethge, and Wieland Brendel. Contrastive learning inverts the data generating process. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12979–12990, 2021. URL https://proceedings.mlr.press/v139/zimmermann21a.html.

A PROOFS

We prove the results of Section 3 in the order in which they appear. Cell i has paired measurements $(X_i^{\mathrm{R}}, X_i^{\mathrm{A}})$, donor D_i and covariates C_i, and the encoders map it to $Z_i^{\mathrm{R}}, Z_i^{\mathrm{A}} \in \mathbb{R}^k$. Donor d has the cells $\mathcal{I}_d = \{i : D_i = d\}$, and $\overline{Z}_d^m$ is the mean of Z_j^m over $j \in \mathcal{I}_d$. Two cells $i, j \in \mathcal{I}_d$ receive the within-donor score

$$\big\langle \nu(\mathrm{ctr}_d^{\mathrm{R}}(Z)_i),\, \nu(\mathrm{ctr}_d^{\mathrm{A}}(Z)_j) \big\rangle, \qquad \mathrm{ctr}_d^m(Z)_i = Z_i^m - \overline{Z}_d^m, \qquad \nu(x) = \frac{x}{\max\{\|x\|_2, \epsilon\}},$$

with a floor $\epsilon > 0$. We write $\|\cdot\|_2$ for the Euclidean norm of a vector and the spectral norm of a matrix, $\|\cdot\|_F$ for the Frobenius norm, and M^+ for the Moore–Penrose pseudoinverse. For a $q \times k$ matrix M, $\sigma_{\max}(M) = \|M\|_2$ and $\sigma_{\min}(M) = \min_{\|x\|_2=1} \|Mx\|_2$. Logarithms are natural.

A.1 LOSS REDUCTION FROM A SAME-DONOR CUE

Fix one anchor with positive score r^+, negative scores r_j^- and temperature $\tau > 0$. Let d be the donor of the anchor and D_j the donor of negative j. A same-donor cue adds a to the positive score and to every negative score with $D_j = d$. At $a = 0$, the softmax denominator splits into same-donor and cross-donor mass,

$$M_= = e^{r^+/\tau} + \sum_{j:D_j=d} e^{r_j^-/\tau}, \qquad M_{\neq} = \sum_{j:D_j\neq d} e^{r_j^-/\tau}.$$

The cue multiplies every same-donor term by $e^{a/\tau}$. The denominator $\mathcal{N}(a)$, the InfoNCE loss $\ell(a)$ of the anchor and the softmax mass ω_a on negatives from other donors are then

$$\mathcal{N}(a) = M_= e^{a/\tau} + M_{\neq}, \qquad \ell(a) = \log \mathcal{N}(a) - \frac{r^+ + a}{\tau}, \qquad \omega_a = \frac{M_{\neq}}{\mathcal{N}(a)}.$$

Proposition A.1 (Loss reduction from a same-donor cue)**.** *For any finite set of negatives and $a \in \mathbb{R}$, $\ell'(a) = -\omega_a/\tau \leq 0$, with strict inequality exactly when a negative from another donor is present; otherwise ℓ is constant. For $a \geq 0$,*

$$\frac{a}{\tau}\,\omega_a \;\leq\; \ell(0) - \ell(a) \;=\; \log \frac{M_= + M_{\neq}}{M_= + M_{\neq} e^{-a/\tau}} \;\leq\; -\log(1 - \omega_0),$$

and the loss reduction increases to the upper bound as $a \to \infty$.

Proposition A.1 gives Equation (1).

Proof. Only the same-donor mass depends on a, so $\mathcal{N}'(a) = M_= e^{a/\tau}/\tau$ and

$$\begin{aligned} \ell'(a) &= \frac{\mathcal{N}'(a)}{\mathcal{N}(a)} - \frac{1}{\tau} \\ &= \frac{M_= e^{a/\tau} - \mathcal{N}(a)}{\tau\, \mathcal{N}(a)} \\ &= -\frac{M_{\neq}}{\tau\, \mathcal{N}(a)} = -\frac{\omega_a}{\tau}. \end{aligned}$$

The mass $M_{\neq}$ is positive exactly when a negative from another donor is present, and then $\ell' < 0$ everywhere. Otherwise $\ell' = 0$. For the loss reduction,

$$\begin{aligned} \ell(0) - \ell(a) &= \log \mathcal{N}(0) - \log \mathcal{N}(a) + \frac{a}{\tau} \\ &= \log \frac{\mathcal{N}(0)}{e^{-a/\tau}\mathcal{N}(a)} \\ &= \log \frac{M_= + M_{\neq}}{M_= + M_{\neq} e^{-a/\tau}}. \end{aligned}$$

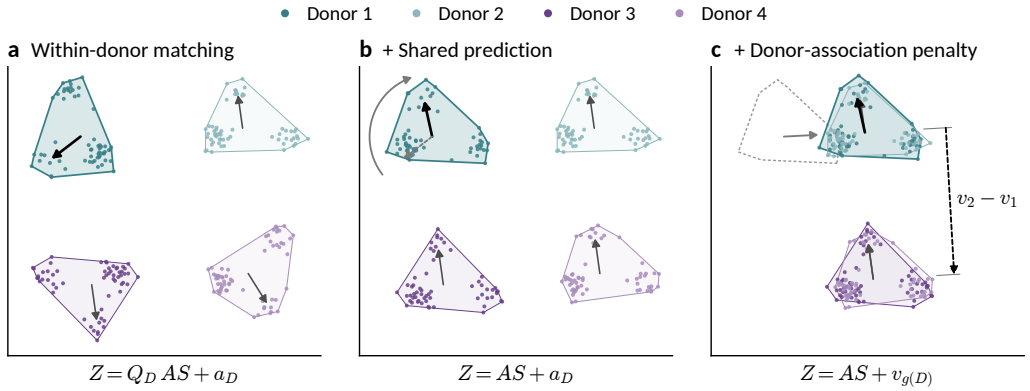

Figure 8: **What each objective leaves free.** Cells of four donors; donors 1 and 2 form one covariate group and donors 3 and 4 the other. Arrows show each donor's orientation, and gray marks show donor 1 in the previous panel. Panels **a**–**c** add the within-donor contrastive loss, shared prediction and the donor-association penalty in turn; the formula under each panel gives the representations that the objectives cannot yet tell apart from AS.

As a grows, $M_{\neq}e^{-a/\tau}$ decreases to zero, so the loss reduction increases to

$$\log \frac{M_{=} + M_{\neq}}{M_{=}} = -\log(1 - \omega_0).$$

For the lower bound, ω_s does not increase in s, so

$$\ell(0) - \ell(a) = \frac{1}{\tau}\int_0^a \omega_s\, ds \;\geq\; \frac{a}{\tau}\,\omega_a. \qquad \square$$

Balanced donors. If all scores are equal and the negatives are drawn uniformly from $|\mathcal{D}|$ equally sized donors, $1 - \omega_0$ is the fraction of same-donor candidates. It tends to $1/|\mathcal{D}|$ as the number of negatives grows, so the loss reduction $-\log(1 - \omega_0)$ approaches $\log|\mathcal{D}|$.

A.2 Invariances of the within-donor contrastive loss

Proposition A.2 (Invariance to donor-specific orthogonal transformations)**.** *For any $Q_d \in \mathsf{O}(k)$ and translations $a_d^{\mathrm{R}}, a_d^{\mathrm{A}} \in \mathbb{R}^k$, the transformations $Z_i^m \mapsto Q_d Z_i^m + a_d^m$ for $i \in \mathcal{I}_d$ and $m \in \{\mathrm{R}, \mathrm{A}\}$ leave every normalized within-donor score unchanged. The orthogonal transformation is shared by the two assays within a donor but may differ across donors.*

A paired loss of the form in Proposition 3.1 is a function of these scores, so Proposition A.2 proves Proposition 3.1.

Proof. Fix donor d and assay m, and let Z' be the transformed representation. Its donor mean is $Q_d\overline{Z}_d^m + a_d^m$. For all vectors x, u, v,

$$\begin{aligned}
\mathrm{ctr}_d^m(Z')_i &= Q_d Z_i^m + a_d^m - (Q_d\overline{Z}_d^m + a_d^m) = Q_d\,\mathrm{ctr}_d^m(Z)_i, \\
\nu(Q_d x) &= \frac{Q_d x}{\max\{\|Q_d x\|_2, \epsilon\}} = \frac{Q_d x}{\max\{\|x\|_2, \epsilon\}} = Q_d \nu(x), \\
\langle Q_d u, Q_d v\rangle &= u^\top Q_d^\top Q_d v = \langle u, v\rangle.
\end{aligned}$$

Applying the three lines in turn to the score of two cells of donor d leaves it unchanged. The translation cancels within each assay, so the two assays may use different translations. $\square$

A.3 Identification by shared prediction and covariate overlap

We prove Theorem 3.1 in the precise form of Theorem A.1, which Figure 8 illustrates on four donors. All expectations, including those used for centering, are taken under one distribution over reference

cells, such as their empirical distribution, with

$$q_c = \Pr(C = c), \qquad p_{d|c} = \Pr(D = d \mid C = c), \qquad \pi_d = \Pr(D = d) = \sum_c q_c p_{d|c}.$$

Sums over the donors of a group run over the donors with $p_{d|c} > 0$.

Weighted means. For vectors x_d with weights $p_d \geq 0$ that sum to one, $m = \sum_d p_d x_d$ and every vector u,

$$\sum_d p_d \|x_d - u\|_2^2 = \sum_d p_d \|x_d - m\|_2^2 + \|m - u\|_2^2, \tag{M1}$$

$$\sum_{d<d'} p_d p_{d'} \|x_d - x_{d'}\|_2^2 = \sum_d p_d \|x_d - m\|_2^2. \tag{M2}$$

In (M1) the cross term vanishes because $\sum_d p_d (x_d - m) = 0$. For (M2), set $u = x_{d'}$ in (M1) and average over d' with weights $p_{d'}$. The right side becomes $2 \sum_d p_d \|x_d - m\|_2^2$, and the left side counts each pair $d < d'$ twice.

Prediction loss. Let $S \in \mathbb{R}^q$ be the target of Section 4 before donor centering, $T = S - \mathbb{E}[S \mid D]$ the centered target and $Z = \frac{1}{2}(Z^{\mathrm{R}} + Z^{\mathrm{A}})$ the joint representation. The shared head $P \in \mathbb{R}^{q \times k}$ has full column rank, which requires $q \geq k$; MINUET uses $q = k = 64$. We write $\sigma_{\min} = \sigma_{\min}(P) > 0$, $\sigma_{\max} = \sigma_{\max}(P)$, $A = P^+$ and $\kappa = \|A\|_2 = 1/\sigma_{\min}$, so that $AP = I_k$ and PA is the orthogonal projection onto the range of P. The head leaves the residual $R = P(Z - \mathbb{E}[Z \mid D]) - T$, the average of the assay residuals $R^m = P(Z^m - \mathbb{E}[Z^m \mid D]) - T$. No representation fits the part $(I - PA)T$ of the target outside the range of P, so the prediction loss subtracts it:

$$\varepsilon^2 = \mathbb{E}\|R\|_2^2 - \mathbb{E}\|(I - PA)T\|_2^2, \qquad \mathbb{E}\|R\|_2^2 \leq \tfrac{1}{2}\big(\mathbb{E}\|R^{\mathrm{R}}\|_2^2 + \mathbb{E}\|R^{\mathrm{A}}\|_2^2\big).$$

Within-group variance of donor means. For a random vector U, let $\mu_{dc}(U) = \mathbb{E}[U \mid D = d, C = c]$ and $\mu_c(U) = \mathbb{E}[U \mid C = c] = \sum_d p_{d|c} \mu_{dc}(U)$. Let B_U be the variance of the donor means within a group, averaged over groups. By (M2), it is also a sum over pairs of donors:

$$B_U = \sum_c q_c \sum_d p_{d|c} \|\mu_{dc}(U) - \mu_c(U)\|_2^2 = \sum_c q_c \sum_{d<d'} p_{d|c} p_{d'|c} \|\mu_{dc}(U) - \mu_{d'c}(U)\|_2^2.$$

Donor graph. Two donors $d \neq d'$ are joined by an edge of weight

$$w_{dd'} = \sum_c q_c p_{d|c} p_{d'|c},$$

which is positive exactly when some group contains both. The Laplacian L is the matrix with

$$x^\top L x = \sum_{d<d'} w_{dd'} (x_d - x_{d'})^2.$$

With $\Pi = \mathrm{diag}(\pi_d)$, the spectral gap λ_π is the smallest positive eigenvalue of $\Pi^{-1/2} L \Pi^{-1/2}$, the normalized Laplacian (Chung, 1997, Ch. 1) of the donor graph with a loop of weight $\sum_c q_c p_{d|c}^2$ at each donor d, whose degrees are then π_d. If the graph has no edges, the terms carrying $1/\lambda_\pi$ are omitted below. Figure 9 shows two donor graphs.

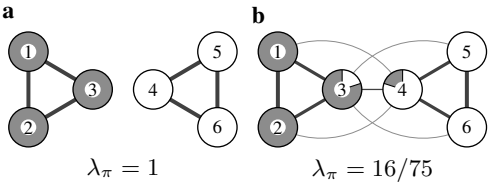

Figure 9: **Donor graphs.** Six donors of 100 cells; gray is the share of a donor's cells in the first covariate group, and line width is proportional to $w_{dd'}$. **a** A covariate constant within donor gives two cliques. **b** Donors 3 and 4 place 20 cells in the other group, which joins the cliques.

All donors of a group c are joined pairwise, so they lie in one connected component $j(c)$. We write $g(d)$ for the component of donor d and

$$\bar{b}_j = \frac{\sum_{g(d)=j} \pi_d b_d}{\sum_{g(d)=j} \pi_d}$$

for the π-weighted mean of vectors b_d over component j.

Identified set.

$$\mathcal{Z} = \{AS + v_{g(D)} : v_j \in \mathbb{R}^k \text{ for each component } j\}, \qquad \operatorname{dist}^2(Z, \mathcal{Z}) = \min_v \mathbb{E}\|Z - AS - v_{g(D)}\|_2^2.$$

Lemma A.1 (Poincaré inequality on the donor graph)**.** *For all vectors (v_d) and all component vectors (u_j),*

$$\lambda_\pi \sum_d \pi_d \|v_d - \bar{v}_{g(d)}\|_2^2 \;\le\; \sum_{d<d'} w_{dd'} \|v_d - v_{d'}\|_2^2 \;\le\; \sum_d \pi_d \|v_d - u_{g(d)}\|_2^2. \tag{11}$$

Every eigenvalue of $\Pi^{-1/2} L \Pi^{-1/2}$ lies in $[0, 1]$, and $\lambda_\pi \ge \lambda_* / \max_d \pi_d$, where λ_* is the smallest positive eigenvalue of L. If every donor belongs to a single covariate group, then $\lambda_\pi = 1$ and both inequalities in Equation (11) are equalities for $u_j = \bar{v}_j$.

The middle term of Equation (11) is the Dirichlet energy of (v_d) on the donor graph. When the donor graph is connected, λ_* is its algebraic connectivity (Fiedler, 1973).

Proof. Right inequality. Fix a group c and let $m_c = \sum_d p_{d|c} v_d$. Then

$$\begin{aligned} \sum_{d<d'} p_{d|c} p_{d'|c} \|v_d - v_{d'}\|_2^2 &= \sum_d p_{d|c} \|v_d - m_c\|_2^2 && \text{by (M2)} \\ &\le \sum_d p_{d|c} \|v_d - u_{j(c)}\|_2^2 && \text{by (M1).} \end{aligned}$$

Multiply by q_c and sum over c. The weights $q_c p_{d|c} p_{d'|c}$ sum to $w_{dd'}$, the weights $q_c p_{d|c}$ sum to π_d, and $u_{j(c)} = u_{g(d)}$ for every donor d of c.

Spectrum. Let $W_{dd'} = \sum_c q_c p_{d|c} p_{d'|c}$ for all d, d', including $d = d'$. Row d of W sums to π_d, and L has entries $-W_{dd'}$ off the diagonal and $\pi_d - W_{dd}$ on it, so

$$L = \Pi - W, \qquad N = \Pi^{-1/2} L \Pi^{-1/2} = I - \Pi^{-1/2} W \Pi^{-1/2}.$$

For $x \in \mathbb{R}^{\mathcal{D}}$ and $y = \Pi^{1/2} x$,

$$y^\top N y = x^\top L x \;\ge\; 0,$$

$$y^\top N y = \|y\|_2^2 - x^\top W x = \|y\|_2^2 - \sum_c q_c \Big(\sum_d p_{d|c} x_d \Big)^2 \;\le\; \|y\|_2^2,$$

so every eigenvalue of N lies in $[0, 1]$. Since $x^\top L x = 0$ exactly when x is constant on each component, the kernel of N is spanned by the vectors $\Pi^{1/2} \mathbf{1}_j$, where $\mathbf{1}_j$ indicates component j.

Left inequality. Fix one coordinate and let x_d be that coordinate of $v_d - \bar{v}_{g(d)}$. By the definition of $\bar{v}_j$, $\langle y, \Pi^{1/2} \mathbf{1}_j \rangle = \sum_{g(d)=j} \pi_d x_d = 0$, so y is orthogonal to the kernel and

$$\sum_{d<d'} w_{dd'} (x_d - x_{d'})^2 = y^\top N y \;\ge\; \lambda_\pi \|y\|_2^2 = \lambda_\pi \sum_d \pi_d x_d^2.$$

Across an edge both donors lie in one component, so $x_d - x_{d'}$ is the chosen coordinate of $v_d - v_{d'}$. Summing over coordinates gives the left inequality.

Comparison with λ_*. Let y be an eigenvector of N for λ_π and $x = \Pi^{-1/2} y$. Then y is orthogonal to the kernel, so $x \ne 0$ has π-weighted component means zero. Let z be x minus its unweighted mean on each component, so $z \perp \mathbf{1}_j$ for every j, and the $\mathbf{1}_j$ span the kernel of L. Then

$$\begin{aligned} \lambda_\pi \sum_d \pi_d x_d^2 = x^\top L x = z^\top L z && \text{since } x - z \text{ is constant on components} \\ \ge \lambda_* \sum_d z_d^2 \\ \ge \frac{\lambda_*}{\max_d \pi_d} \sum_d \pi_d z_d^2 \\ \ge \frac{\lambda_*}{\max_d \pi_d} \sum_d \pi_d x_d^2 && \text{by (M1) on each component.} \end{aligned}$$

Dividing by $\sum_d \pi_d x_d^2 > 0$ gives the claim.

One group per donor. If each donor d belongs only to group $c(d)$, the components are the groups and $\pi_d = q_{c(d)} p_{d|c(d)}$. For donors d, d' of one group c, equal or not,

$$\frac{W_{dd'}}{\sqrt{\pi_d \pi_{d'}}} = \frac{q_c p_{d|c} p_{d'|c}}{q_c \sqrt{p_{d|c} p_{d'|c}}} = \sqrt{p_{d|c} p_{d'|c}},$$

and $W_{dd'} = 0$ across groups. With $h_c = (\sqrt{p_{d|c}})_d$, this reads $\Pi^{-1/2} W \Pi^{-1/2} = \sum_c h_c h_c^\top$. The h_c are orthonormal, so $N = I - \sum_c h_c h_c^\top$ is an orthogonal projection and $\lambda_\pi = 1$. Then $y^\top N y = \|y\|_2^2$ for y orthogonal to the kernel, so the left inequality is an equality, and its left side is the right side at $u_j = \bar{v}_j$. □

Theorem A.1 (Two-sided identification bound)**.** *Assume finitely many donors and covariate groups, finite second moments, and $\sigma_{\min}(P) > 0$. Write $\Delta = Z - AS$, let $b_d = \mathbb{E}[\Delta \mid D = d]$ be its mean over donor d, and let $\Delta_w = \Delta - b_D$ be its within-donor part. Then $\varepsilon^2 = \mathbb{E}\|P\Delta_w\|_2^2 \geq 0$. The distance splits into a within-donor part and a donor-mean part,*

$$\operatorname{dist}^2(Z, \mathcal{Z}) = \mathbb{E}\|\Delta_w\|_2^2 + \sum_d \pi_d \|b_d - \bar{b}_{g(d)}\|_2^2. \tag{12}$$

Prediction controls the first part,

$$\varepsilon^2 / \sigma_{\max}^2 \leq \mathbb{E}\|\Delta_w\|_2^2 \leq \kappa^2 \varepsilon^2, \tag{13}$$

and B_Z and B_S control the Dirichlet energy of the donor means,

$$\Big[\sqrt{B_Z} - \kappa\big(\sqrt{B_S} + \varepsilon\big)\Big]_+ \leq \Big(\sum_{d<d'} w_{dd'} \|b_d - b_{d'}\|_2^2\Big)^{1/2} \leq \sqrt{B_Z} + \kappa\big(\sqrt{B_S} + \varepsilon\big). \tag{14}$$

Combining the three with Lemma A.1 gives

$$\frac{\varepsilon^2}{\sigma_{\max}^2} + \Big[\sqrt{B_Z} - \kappa\big(\sqrt{B_S} + \varepsilon\big)\Big]_+^2 \;\leq\; \operatorname{dist}^2(Z, \mathcal{Z}) \;\leq\; \kappa^2 \varepsilon^2 + \frac{1}{\lambda_\pi}\Big[\sqrt{B_Z} + \kappa\big(\sqrt{B_S} + \varepsilon\big)\Big]^2. \tag{15}$$

If $B_S = 0$, then $Z \in \mathcal{Z}$ if and only if $\varepsilon = 0$ and $B_Z = 0$, and

$$\frac{\varepsilon^2 + B_Z}{\sigma_{\max}^2 + (1 + \kappa \sigma_{\max})^2} \;\leq\; \operatorname{dist}^2(Z, \mathcal{Z}) \;\leq\; \frac{\max\{2, 3\kappa^2\}}{\lambda_\pi}\big(\varepsilon^2 + B_Z\big), \tag{16}$$

which is Equation (3) *with explicit constants.*

Proof. Within-donor part. Take donor means in $Z = AS + b_D + \Delta_w$, using $\mathbb{E}[\Delta_w \mid D] = 0$, and recall $T = S - \mathbb{E}[S \mid D]$:

$$\begin{aligned} \mathbb{E}[Z \mid D] &= A\,\mathbb{E}[S \mid D] + b_D, \\ Z - \mathbb{E}[Z \mid D] &= AT + \Delta_w, \\ R &= PAT + P\Delta_w - T \\ &= P\Delta_w - (I - PA)T. \end{aligned}$$

The two terms of R lie in the range of P and in its orthogonal complement, so

$$\mathbb{E}\|R\|_2^2 = \mathbb{E}\|P\Delta_w\|_2^2 + \mathbb{E}\|(I - PA)T\|_2^2, \qquad \text{hence} \qquad \varepsilon^2 = \mathbb{E}\|P\Delta_w\|_2^2.$$

Squaring $\sigma_{\min}\|x\|_2 \leq \|Px\|_2 \leq \sigma_{\max}\|x\|_2$ at $x = \Delta_w$ and taking expectations gives

$$\sigma_{\min}^2\, \mathbb{E}\|\Delta_w\|_2^2 \;\leq\; \varepsilon^2 \;\leq\; \sigma_{\max}^2\, \mathbb{E}\|\Delta_w\|_2^2,$$

which is Equation (13) since $\kappa = 1/\sigma_{\min}$.

Distance split. For component vectors v, $\Delta - v_{g(D)} = \Delta_w + (b_D - v_{g(D)})$. The second term is a function of D and $\mathbb{E}[\Delta_w \mid D] = 0$, so the cross term has mean zero, and

$$\begin{aligned}\mathbb{E}\|\Delta - v_{g(D)}\|_2^2 &= \mathbb{E}\|\Delta_w\|_2^2 + \sum_d \pi_d \|b_d - v_{g(d)}\|_2^2 \\ &\geq \mathbb{E}\|\Delta_w\|_2^2 + \sum_d \pi_d \|b_d - \bar{b}_{g(d)}\|_2^2 \qquad \text{by (M1) on each component,}\end{aligned}$$

with equality at $v = \bar{b}$. Minimizing over v gives Equation (12).

Donor means. Fix a group c and a donor d of c. Averaging $\Delta = Z - AS$ and $\Delta = b_D + \Delta_w$ over the cells of d in c gives

$$\mu_{dc}(Z) - A\mu_{dc}(S) = \mu_{dc}(\Delta) = b_d + \mu_{dc}(\Delta_w).$$

Subtracting the same identity for a second donor d' of c,

$$b_d - b_{d'} = \big[\mu_{dc}(Z) - \mu_{d'c}(Z)\big] - A\big[\mu_{dc}(S) - \mu_{d'c}(S)\big] - \big[\mu_{dc}(\Delta_w) - \mu_{d'c}(\Delta_w)\big].$$

For vectors $x_{cdd'}$ indexed by a group c and a pair $d < d'$ of its donors, let

$$\|x\|_w^2 = \sum_c q_c \sum_{d<d'} p_{d|c} p_{d'|c} \|x_{cdd'}\|_2^2.$$

By the definition of $w_{dd'}$, since $b_d - b_{d'}$ does not depend on c, the pairwise form of B_U, and $\|A\|_2 = \kappa$, the four terms have norms

$$\|b_d - b_{d'}\|_w^2 = \sum_{d<d'} w_{dd'} \|b_d - b_{d'}\|_2^2, \qquad \|\mu_{dc}(Z) - \mu_{d'c}(Z)\|_w^2 = B_Z,$$

$$\|A[\mu_{dc}(S) - \mu_{d'c}(S)]\|_w^2 \leq \kappa^2 B_S, \qquad \|\mu_{dc}(\Delta_w) - \mu_{d'c}(\Delta_w)\|_w^2 = B_{\Delta_w}.$$

The last term satisfies

$$\begin{aligned}B_{\Delta_w} &\leq \sum_c q_c \sum_d p_{d|c} \|\mu_{dc}(\Delta_w)\|_2^2 && \text{by (M1) at } u = 0 \\ &= \mathbb{E}\big\|\mathbb{E}[\Delta_w \mid D, C]\big\|_2^2 \\ &\leq \mathbb{E}\|\Delta_w\|_2^2 && \text{by conditional Jensen} \\ &\leq \kappa^2 \varepsilon^2 && \text{by Equation (13).}\end{aligned}$$

The triangle inequality for $\|\cdot\|_w$ in the identity gives the upper half of Equation (14). Solving the identity for the first bracket and applying the triangle inequality again gives

$$\sqrt{B_Z} \leq \Big(\sum_{d<d'} w_{dd'} \|b_d - b_{d'}\|_2^2\Big)^{1/2} + \kappa\sqrt{B_S} + \kappa\varepsilon,$$

and the middle term is nonnegative, which is the lower half.

The two bounds. By Equation (12), Equation (13), the left inequality of Equation (11) at $v = b$, and Equation (14),

$$\begin{aligned}\operatorname{dist}^2(Z, \mathcal{Z}) &= \mathbb{E}\|\Delta_w\|_2^2 + \sum_d \pi_d \|b_d - \bar{b}_{g(d)}\|_2^2 \\ &\leq \kappa^2\varepsilon^2 + \frac{1}{\lambda_\pi} \sum_{d<d'} w_{dd'} \|b_d - b_{d'}\|_2^2 \\ &\leq \kappa^2\varepsilon^2 + \frac{1}{\lambda_\pi}\Big[\sqrt{B_Z} + \kappa\big(\sqrt{B_S} + \varepsilon\big)\Big]^2.\end{aligned}$$

With the lower halves instead, and the right inequality of Equation (11) at $u_j = \bar{b}_j$,

$$\begin{aligned}\operatorname{dist}^2(Z, \mathcal{Z}) &\geq \frac{\varepsilon^2}{\sigma_{\max}^2} + \sum_{d<d'} w_{dd'} \|b_d - b_{d'}\|_2^2 \\ &\geq \frac{\varepsilon^2}{\sigma_{\max}^2} + \Big[\sqrt{B_Z} - \kappa\big(\sqrt{B_S} + \varepsilon\big)\Big]_+^2.\end{aligned}$$

Membership. Let $B_S = 0$. If $\varepsilon = B_Z = 0$, the upper bound vanishes. The minimum in the distance is attained at $v = \bar{b}$, so $Z = AS + \bar{b}_{g(D)}$ almost surely. Conversely, if $Z = AS + v_{g(D)}$, then $\Delta = v_{g(D)}$ is constant within donors, so $\Delta_w = 0$ and $\varepsilon = 0$. All donors of c lie in component $j(c)$, so $\mu_{dc}(Z) = A\mu_{dc}(S) + v_{j(c)}$, and for two donors of c

$$\|\mu_{dc}(Z) - \mu_{d'c}(Z)\|_2^2 = \|A[\mu_{dc}(S) - \mu_{d'c}(S)]\|_2^2 \le \kappa^2 \|\mu_{dc}(S) - \mu_{d'c}(S)\|_2^2.$$

The pairwise form of B_U then gives $B_Z \le \kappa^2 B_S = 0$.

Explicit constants. Let $B_S = 0$ and $\delta = \operatorname{dist}(Z, \mathcal{Z})$. The lower bound gives

$$\varepsilon \le \sigma_{\max}\delta, \qquad \sqrt{B_Z} \le \delta + \kappa\varepsilon \le (1 + \kappa\sigma_{\max})\delta.$$

Squaring and adding gives the left side of Equation (16). For the right side, the upper bound gives

$$\begin{aligned}
\delta^2 &\le \kappa^2\varepsilon^2 + \frac{(\sqrt{B_Z} + \kappa\varepsilon)^2}{\lambda_\pi} \\
&\le \kappa^2\varepsilon^2 + \frac{2B_Z + 2\kappa^2\varepsilon^2}{\lambda_\pi} && \text{since } (x+y)^2 \le 2x^2 + 2y^2 \\
&\le \frac{3\kappa^2\varepsilon^2 + 2B_Z}{\lambda_\pi} && \text{since } \lambda_\pi \le 1 \\
&\le \frac{\max\{2, 3\kappa^2\}}{\lambda_\pi}\big(\varepsilon^2 + B_Z\big). && \square
\end{aligned}$$

Both constants are attained. Let y be an eigenvector of $\Pi^{-1/2} L \Pi^{-1/2}$ with eigenvalue $\lambda > 0$ and set $x = \Pi^{-1/2} y$, so that $Lx = \lambda \Pi x$. Since $\mathbf{1}_j^\top L = 0$, we get $\lambda \sum_{g(d)=j} \pi_d x_d = 0$, so x has π-weighted component means zero. Take $P = I_k$, $S = 0$ and $Z = x_D u$ for a unit vector u. Then $\kappa = \sigma_{\max} = 1$, $b_d = x_d u$, $\Delta_w = 0$ and $\varepsilon = B_S = 0$. By the pairwise form of B_U and Equation (12),

$$\begin{aligned}
B_Z &= \sum_{d<d'} w_{dd'}(x_d - x_{d'})^2 \\
&= x^\top L x = \lambda\, x^\top \Pi x, \\
\operatorname{dist}^2(Z, \mathcal{Z}) &= x^\top \Pi x = B_Z/\lambda.
\end{aligned}$$

At $\lambda = \lambda_\pi$ this attains the upper bound of Equation (15). When 1 is an eigenvalue, $\lambda = 1$ attains the lower bound (Figure 10).

If every donor belongs to a single group, the lower bound is attained by every Z with $\varepsilon = B_S = 0$: Lemma A.1 gives $\lambda_\pi = 1$, and both sides of Equation (15) equal B_Z. In MINUET the covariate is constant within donor and the target is centered within donor, so every $\mu_{dc}(S) = \mathbb{E}[S \mid D = d]$ vanishes and $B_S = 0$.

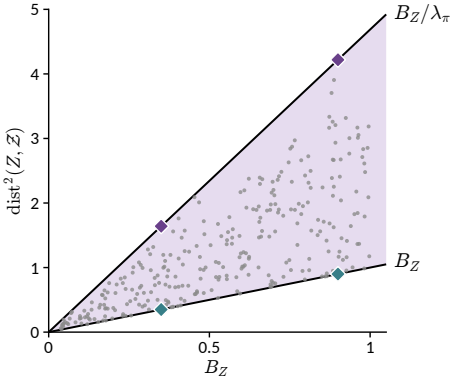

Figure 10: **Both sides of the bound are attained.** Gray points: $Z = x_D$ for random x on the graph of Figure 9b, with $P = I$ and $S = 0$. Black lines: the bounds B_Z and B_Z/λ_π. Diamonds: $Lx = \lambda\Pi x$, purple for $\lambda = \lambda_\pi$ and teal for $\lambda = 1$.

Relation to the donor-association penalty. With the one-hot donor vector e_D, column d of $\operatorname{Cov}(Z, e_D \mid C = c)$ is $p_{d|c}\{\mu_{dc}(Z) - \mu_c(Z)\}$. Under the empirical distribution of the n_c cells of group c, the numerator of Equation (9) is therefore

$$n_c^2 \,\|\operatorname{Cov}(Z, e_D \mid C = c)\|_F^2 = n_c^2 \sum_d p_{d|c}^2 \|\mu_{dc}(Z) - \mu_c(Z)\|_2^2.$$

It vanishes exactly when the group-c term of B_Z does, so all numerators vanish exactly when $B_Z = 0$.

Component-wise translations preserve every objective. Adding $v_{g(D_i)}$ to both Z_i^{R} and Z_i^{A} adds a vector that is constant within each donor and within each covariate group. Within-donor centering removes it from similarity scores and prediction residuals, it leaves the donor-association penalty and centered covariances unchanged, and reconstruction uses only the private variables. No objective distinguishes two members of $\mathcal{Z}$, so the identified set is exactly $\mathcal{Z}$.

A.4 Query donors

A query donor q is embedded by the frozen encoders, and its cells enter neither the fitting objective nor the reference distribution. On the cells of q, with the frozen head,

$$T_q = S - \mathbb{E}[S \mid D = q], \qquad R_q = P(Z - \mathbb{E}[Z \mid D = q]) - T_q,$$
$$\varepsilon_q^2 = \mathbb{E}[\|R_q\|_2^2 \mid D = q] - \mathbb{E}[\|(I - PA)T_q\|_2^2 \mid D = q].$$

Let $b_q = \mathbb{E}[Z - AS \mid D = q]$ and $p_{c|q} = \Pr(C = c \mid D = q)$. Suppose that every covariate group of q lies in reference component j, with the reference translation $\bar{b}_j$ of Theorem A.1.

Corollary A.1 (Query donors)**.** *With $\zeta_c = \mathbb{E}_{\mathrm{ref}}[Z - AS \mid C = c]$ and $o_c = \mathbb{E}[Z - AS \mid D = q, C = c] - \zeta_c$,*

$$\frac{\varepsilon_q^2}{\sigma_{\max}^2} + \|b_q - \bar{b}_j\|_2^2 \le \mathbb{E}\big[\|Z - AS - \bar{b}_j\|_2^2 \;\big|\; D = q\big] \le \kappa^2 \varepsilon_q^2 + \|b_q - \bar{b}_j\|_2^2,$$
$$b_q - \bar{b}_j = \textstyle\sum_c p_{c|q}\big(o_c + \zeta_c - \bar{b}_j\big).$$

Moreover $\sum_c q_c \|\zeta_c - \bar{b}_{j(c)}\|_2^2 \le \mathrm{dist}^2(Z, \mathcal{Z})$ over the reference donors.

Proof. On the cells of q, write $Z - AS = b_q + \Delta_w$, so that $\mathbb{E}[\Delta_w \mid D = q] = 0$. The within-donor step of the proof of Theorem A.1, run on the cells of q, gives

$$\frac{\varepsilon_q^2}{\sigma_{\max}^2} \le \mathbb{E}[\|\Delta_w\|_2^2 \mid D = q] \le \kappa^2 \varepsilon_q^2.$$

The vector $b_q - \bar{b}_j$ is fixed on the cells of q, so the cross term vanishes and

$$\begin{aligned}\mathbb{E}\big[\|Z - AS - \bar{b}_j\|_2^2 \;\big|\; D = q\big] &= \mathbb{E}\big[\|\Delta_w + (b_q - \bar{b}_j)\|_2^2 \;\big|\; D = q\big] \\ &= \mathbb{E}[\|\Delta_w\|_2^2 \mid D = q] + \|b_q - \bar{b}_j\|_2^2.\end{aligned}$$

Combining the two displays gives the two bounds. By the tower rule over the groups of q,

$$b_q = \sum_c p_{c|q}\, \mathbb{E}[Z - AS \mid D = q, C = c] = \sum_c p_{c|q}(o_c + \zeta_c),$$

and subtracting $\bar{b}_j = \sum_c p_{c|q} \bar{b}_j$ gives the second line. For the reference bound, $\bar{b}_{g(D)} = \bar{b}_{j(c)}$ whenever $C = c$, so

$$\begin{aligned}\sum_c q_c \|\zeta_c - \bar{b}_{j(c)}\|_2^2 &= \sum_c q_c \big\|\mathbb{E}_{\mathrm{ref}}[\Delta - \bar{b}_{g(D)} \mid C = c]\big\|_2^2 \\ &\le \mathbb{E}_{\mathrm{ref}}\|\Delta - \bar{b}_{g(D)}\|_2^2 \qquad \text{by conditional Jensen} \\ &= \mathrm{dist}^2(Z, \mathcal{Z}),\end{aligned}$$

because $v = \bar{b}$ attains the minimum in the distance. □

Covariates constant within donor. This case gives Corollary 3.1. The query q belongs to one group c, and the reference components are the groups, as in the proof of Lemma A.1, so component j is group c. The target is centered within donor, so $\mathbb{E}[S \mid D] = 0$ for every donor, and $\mathbb{E}_{\mathrm{ref}}[S \mid C = c] = 0$ because group c is a union of reference donors. With $\pi_d = q_c p_{d|c}$ for the donors of c and the tower rule,

$$b_q = \mathbb{E}[Z \mid D = q] = \mu_q, \qquad \zeta_c = \mathbb{E}_{\mathrm{ref}}[Z \mid C = c] = \mu_c, \qquad \bar{b}_j = \sum_d p_{d|c} b_d = \zeta_c.$$

With $r = \|\mu_q - \mu_c\|_2^2$, the two bounds of Corollary A.1 satisfy

$$\frac{\varepsilon_q^2}{\sigma_{\max}^2} + r \geq \min\{1, \sigma_{\max}^{-2}\}(\varepsilon_q^2 + r),$$
$$\kappa^2 \varepsilon_q^2 + r \leq \max\{1, \kappa^2\}(\varepsilon_q^2 + r),$$

which is Corollary 3.1.

Both terms are computed without labels. The target S is built from the inputs (Section 4), so ε_q, o_c and ζ_c are computed from cells, donors and covariates alone. In Corollary 3.1 they reduce to ε_q and $\|\mu_q - \mu_c\|_2$. The reference term $\zeta_c - \bar{b}_{j(c)}$ vanishes when $Z \in \mathcal{Z}$, by the last bound of Corollary A.1.

A.5 Donor suppression and phenotype information

Here Y is a finite-valued phenotype, such as disease status. Proposition A.3 and the paragraph after it follow from the chain rule of mutual information (Cover & Thomas, 2006, Thm. 2.5.2).

Proposition A.3 (Phenotype uncertainty under donor suppression)**.** *For finite-valued Y and D, $H(Y \mid Z, C) \geq I(Y; D \mid C) - I(Z; D \mid C)$.*

Proof. Expand $I(Y, Z; D \mid C)$ by the chain rule in both orders, and use $H(Y \mid D, Z, C) \geq 0$ for the finite-valued Y:

$$\begin{aligned} I(Y; D \mid C) &\leq I(Y; D \mid C) + I(Z; D \mid Y, C) \\ &= I(Y, Z; D \mid C) \\ &= I(Z; D \mid C) + I(Y; D \mid Z, C) \\ &= I(Z; D \mid C) + H(Y \mid Z, C) - H(Y \mid D, Z, C) \\ &\leq I(Z; D \mid C) + H(Y \mid Z, C). \end{aligned}$$

□

Donor-level phenotypes. Expanding $I(Z; D, Y)$ by the chain rule in both orders gives

$$I(Z; D) + I(Z; Y \mid D) = I(Z; Y) + I(Z; D \mid Y).$$

If $Y = f(D)$, then $I(Z; Y \mid D) = 0$, which gives the identity used in Section 3.

Proposition A.4 (What the identified representation keeps)**.** *If $Z = AS + v_{g(D)}$ almost surely, then for every donor d and group c, the conditional law of Z given $(D, C) = (d, c)$ is that of AS translated by a vector that depends on c alone. Consequently, $I(Z; D \mid C) = I(AS; D \mid C)$ and $I(Z; Y \mid C) = I(AS; Y \mid C)$ for every random variable Y.*

Proof. If $p_{d|c} > 0$, donor d belongs to group c, so $g(d) = j(c)$ and $Z = AS + v_{j(c)}$ on $\{D = d, C = c\}$. This proves the first claim, and $Z = AS + v_{j(C)}$ almost surely. Given $C = c$, the map $x \mapsto x + v_{j(c)}$ is a bijection, so Z and AS are functions of each other. The data-processing inequality (Cover & Thomas, 2006, Thm. 2.8.1), applied in both directions, gives

$$I(Z; Y \mid C = c) \leq I(AS; Y \mid C = c) \leq I(Z; Y \mid C = c).$$

Averaging over c gives $I(Z; Y \mid C) = I(AS; Y \mid C)$, and taking $Y = D$ gives $I(Z; D \mid C) = I(AS; D \mid C)$. □

B Model and Training Details

Architecture. Separate residual MLP encoders return a deterministic shared vector z^m and the mean and log variance of a Gaussian private variable v^m for each assay $m \in \{\mathrm{R}, \mathrm{A}\}$. The base trunk uses BatchNorm at its input and output, with GELU activation and dropout in each residual branch. For LODO transfer, LayerNorm precedes each residual branch and replaces the output BatchNorm, and the input BatchNorm is removed. BatchNorm relies on statistics estimated on the source datasets,

Table 1: **MINUET base configuration.** Used for all cohorts and protocols; LODO transfer replaces BatchNorm with LayerNorm.

Config key	Value
Encoder architecture	
`model.hidden_dim`	224
`model.residual_blocks`	3
`model.shared_dim`	64
`model.private_dim`	16
`model.decoder_rank`	64
`model.dropout`	0.2
`model.residual_normalization`	`batchnorm`
`model.persistent_parameters`	2,127,680
`model.training_only_head`	4,096
Optimization	
`train.optimizer`	Adam
`train.learning_rate`	3×10^{-4}
`train.weight_decay`	3×10^{-7}
`train.late_learning_rate_multiplier`	0.25
`train.gradient_clip`	10
`train.encoder_updates`	16,000
Objective	
`objective.pair_temperature`	0.1
`objective.pair_weight`	2.0
`objective.prediction_rank`	32
`objective.prediction_weight`	1.5
`objective.donor_association_weight`	0.1
`objective.conditional_covariance_weight`	0.35
`objective.covariance_start_fraction`	0.8
`objective.private_kl_weight`	7×10^{-5}
Sampling	
`sampler.contexts_per_batch`	2
`sampler.donors_per_context`	4
`sampler.cells_per_block`	12
`sampler.min_donors_per_context`	2
`sampler.batches_per_epoch`	100

which need not match an unseen target dataset. Table 2 compares the two normalizations under otherwise identical settings over three directions and three seeds. Independent $16 \to 64 \to p_m$ decoders reconstruct each assay from v^m alone. RNA uses a negative-binomial likelihood with observed library size and learned gene-wise inverse dispersion; ATAC uses a Bernoulli likelihood for binary accessibility. Only the private posterior contributes a KL term:

$$\mathcal{L}_{\text{rec}} = \sum_{m \in \{\text{R,A}\}} \mathbb{E}_{q_m(v_i^m \mid X_i^m)}[-\log p_m(X_i^m \mid v_i^m)] + \beta_v D_{\text{KL}}(q_m(v_i^m \mid X_i^m) \,\|\, \mathcal{N}(0, I)) \,. \quad (17)$$

Prediction target. For each assay, we subtract each reference donor's mean from the counts of its cells and fit a separate rank-32 randomized SVD to the centered reference cells. The scores are standardized with reference-cell means and standard deviations and then concatenated.

Covariance-alignment penalty. For the donor-centered joint covariance Σ_{cd} of donor d in group c,

$$\widetilde{\Sigma}_{cd} = \frac{\Sigma_{cd}}{\max\{\text{tr}(\Sigma_{cd}), \epsilon\}}, \qquad \mathcal{L}_{\text{cov}} = \frac{1}{|\mathcal{C}|} \sum_c \frac{1}{|\mathcal{D}_c|} \sum_{d \in \mathcal{D}_c} \|\widetilde{\Sigma}_{cd} - \overline{\Sigma}_c\|_F^2, \quad (18)$$

where $\overline{\Sigma}_c$ averages normalized covariances within group c. This extends correlation alignment (Sun & Saenko, 2016) to covariate groups. The penalty is switched on at update u of U by $\eta(u) = \mathbf{1}\{u \geq$

Table 2: **LODO normalization ablation.** Each entry gives the mean over three representation seeds with the standard deviation in parentheses. Bold marks the better result within each dataset and metric; darker purple indicates a better value.

Metric	Mannens		ROSMAP		Wang	
	BatchNorm	LayerNorm	BatchNorm	LayerNorm	BatchNorm	LayerNorm
ARI ↑	**0.652** (**0.042**)	0.591 (0.059)	0.785 (0.000)	**0.787** (**0.000**)	0.349 (0.024)	**0.543** (**0.054**)
NMI ↑	**0.640** (**0.030**)	0.601 (0.030)	0.836 (0.000)	**0.843** (**0.000**)	0.552 (0.008)	**0.623** (**0.025**)
CiLISI ↑	0.194 (0.002)	**0.199** (**0.006**)	0.256 (0.003)	**0.278** (**0.003**)	0.099 (0.001)	**0.109** (**0.001**)
Cell-type macro-F1 ↑	0.282 (0.010)	**0.284** (**0.040**)	0.634 (0.020)	**0.696** (**0.021**)	0.555 (0.037)	**0.562** (**0.023**)
FOSCTTM ↓	**0.251** (**0.002**)	0.270 (0.008)	0.203 (0.003)	**0.195** (**0.003**)	0.253 (0.003)	**0.236** (**0.001**)

$\mathrm{round}(\rho_{\mathrm{cov}} U)\}$. Both donor penalties weight covariate groups equally, and the covariance-alignment penalty also weights donors equally. A group contributes to a batch's penalties only when the batch contains at least two of its donors; the donor-association penalty further requires four cells in the group, and the covariance-alignment penalty two cells per donor.

Configuration. Table 1 gives the base configuration, which is used in all synthetic experiments and in every real-data cohort and protocol; each seed corresponds to an independently fitted model. With 2,048 genes and 4,096 peaks, the two encoders used at inference have 2.1 million parameters. Each real-data batch contains up to two covariate groups and four donors per group, with 12 cells from each donor. The normalization floor is $\epsilon_{\mathrm{norm}} = 10^{-6}$. The covariates that define the groups depend on the protocol (Section D).

Ablation variants and in-house controls. Each ablation changes one element of **`Full`** and keeps all other model, training, and evaluation settings fixed:

- **`+ pooled candidates`**: negatives are drawn across donors, with within-donor centering and the number of candidates unchanged;
- **`- pair matching`**: the same-cell contrastive term is removed;
- **`- prediction`**: the shared prediction head and target are removed;
- **`- donor association`**: the donor-association penalty is removed;
- **`- covariance`**: the covariance-alignment penalty is removed;
- **`Unconditional`**: both donor penalties use a single group, with batch sampling unchanged.

Real-data variants share three seeds per cohort and synthetic variants share 20 seeds, so every comparison is a seed-paired difference from **`Full`**, with Student-t intervals on real data and paired bootstrap intervals on synthetic data. Two in-house controls use MINUET's encoders but omit shared prediction and the donor penalties:

- **`Pooled`** (Pooled contrastive): negatives are drawn across donors;
- **`Within-donor`** (Within-donor contrastive): the centered within-donor loss in Equation (7).

Baseline models. We compared seven published methods and DC-PCA under the protocols in Section D. All methods share data, selected features, splits, seeds, and metrics; preprocessing, architectures, and training schedules are method-specific. MultiVI and scGLUE receive the donor label as their batch covariate; the other published methods do not accept one.

- **Apollo** separates shared and assay-specific information using an autoencoder with a partially shared latent space (Zhang et al., 2026). Code: `github.com/uhlerlab/APOLLO`.
- **scPair** predicts one assay from another to learn cell representations and select informative features (Hu & Quon, 2024). Code: `github.com/quon-titative-biology/scPair`.

- **scPairing** uses contrastive learning to align paired cells from low-dimensional assay representations (Niu et al., 2025). Code: `github.com/Ding-Group/scPairing`.
- **scCLIP** aligns RNA and ATAC representations using transformer encoders and same-cell contrastive learning (Xiong et al., 2023a). Code: `github.com/jsxlei/scCLIP`.
- **IndiSeek** separates shared and assay-specific information through reconstruction and independence objectives (Gui et al., 2025). Code: author-provided OpenReview supplement.
- **MultiVI** is a variational generative model that jointly represents RNA expression and chromatin accessibility (Ashuach et al., 2023). Code: `github.com/scverse/scvi-tools`.
- **scGLUE** couples assay-specific generative models through a regulatory guidance graph to learn a common embedding (Cao & Gao, 2022). Code: `github.com/gao-lab/GLUE`.
- **DC-PCA** is the within-donor level of multilevel component analysis (Timmerman, 2006): it uses the prediction target of MINUET (Section B) as the representation. Each assay is centered within donor and reduced by a randomized SVD with the rank of MINUET's target in the same run, and the standardized scores of the two assays are concatenated. A query donor is centered with the mean of its own cells, as in Section A.4, and projected with the reference fit. Unlike MINUET, DC-PCA therefore cannot embed a cell without the other cells of its donor, and this centering favors it on donor mixing and donor prediction. In phenotype prediction the target is refitted in every fold, like MINUET. The RNA and ATAC scores lie in different spaces, so paired-cell matching is N/A for DC-PCA.

C SYNTHETIC EXPERIMENTS

Each seed generates one paired dataset, which all methods receive unchanged, together with the same initialization seed. A dataset contains eight donors, 768 cells, 64 RNA features, 96 ATAC features, and a six-dimensional latent state. The shortcut and transformation generators place 48 cells in each donor–covariate group, whereas the conditional generator splits the 96 cells of a donor between the two phenotype groups according to ρ. RNA counts follow a Poisson observation model and ATAC entries a Bernoulli model. Latent states are used only for evaluation. In the shortcut and conditional generators, six donors are reference donors and two are query donors; the transformation generator instead holds out cells of reference donors. All synthetic classification probes are scored by macro-F1. MINUET uses the configuration in Table 1 for all 20 seeds. DC-PCA uses MINUET's rank-32 target fitted on the same cells. Query donors are centered with their own means, and held-out cells of reference donors are centered with the means of their donors.

Same-donor shortcut. The latent state contains four asymmetric cell-state components and a shift along the biological covariate. A five-dimensional donor code, independent of the latent state, enters the observation maps of both assays. The cell-state mixture has the same distribution in every donor, and the phenotype groups have equal sizes. A linear phenotype probe is fitted on the reference donors and evaluated on the query donors. Within-donor recall@1 is the fraction of query cells whose paired measurement is the nearest candidate among cells of the same donor. Donor-prediction macro-F1 is computed with a cross-validated logistic classifier over all eight donors (random guess: $1/8$).

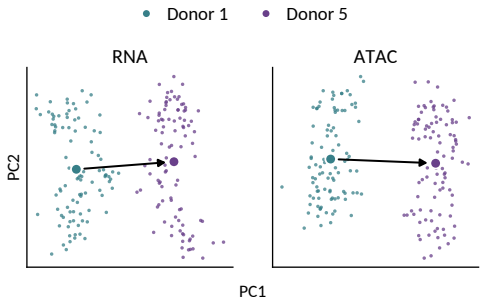

Figure 11: **Same-donor shortcut generator.** PCA of RNA and ATAC for donors 1 and 5 of seed 0; points are cells and arrows join the donor means. One donor code shifts both assays in the same direction.

Donor-specific transformations. Within each biological-covariate group, four Gaussian components have counts $(5, 9, 14, 20)$ and fixed asymmetric six-dimensional means; the two groups shift all four means by $-h$ and $+h$, respectively. Coordinate-wise biological noise has standard deviations $(.18, .22, .26, .30, .34, .38)$, and paired noise has standard deviation $.15$. Every donor receives an independently drawn orthogonal transformation before the RNA and ATAC observation maps. The four component weights are distinct, so the components are identifiable. The diagonal covariance has distinct eigenvalues, and one component mean has a nonzero projection on every eigendirection; together, these conditions make the identity the only orthogonal transformation that preserves both the mixture distribution and that mean. Each donor–covariate group holds out $(1, 2, 4, 5)$ cells of the four components for evaluation and uses the remaining 36 cells for fitting. The evaluation therefore tests new cells from reference donors rather than query donors, whose transformations are drawn independently and never seen during fitting. For a learned representation $Z \in \mathbb{R}^{n\times p}$ with $p \geq 6$ and latent target $Y \in \mathbb{R}^{n\times 6}$, let Z_F, Y_F denote all fitting rows in one covariate group. We center both matrices and scale them to unit root mean square using only those rows, which gives Z_F°, Y_F°. The shared map solves the rectangular orthogonal Procrustes problem

$$\min_{W^\top W = I_6} \|Z_F^\circ W - Y_F^\circ\|_F^2.$$

The same centering and scaling are applied to the held-out rows. Donor-specific maps solve the same problem on the fitting rows of one donor at a time, so the shared and donor-specific maps differ only in which rows determine W; both are scored on held-out cells only. The shared map gives the R^2 of the shared cross-donor readout. The registration gap is the R^2 of the donor-specific maps minus the shared R^2; it is zero when a single map registers every donor. The R^2 of latent-state recovery is the variance-weighted R^2 of a ridge regression from the representation to the simulated latent state, also evaluated on held-out cells. Representations with $p < 6$ cannot span the target and are not scored on this metric.

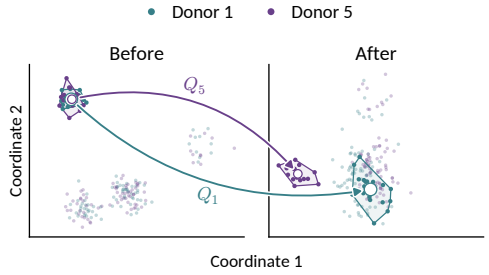

Figure 12: **Donor-specific transformation generator.** The first two latent coordinates of donors 1 and 5 of seed 0, before and after their orthogonal transformations Q_1 and Q_5. One highlighted cell state starts at the same place in both donors and ends at two different places.

Conditional suppression. For donors in the first and second halves, respectively, the target phenotype fractions are $.5 - .5\rho$ and $.5 + .5\rho$, with $\rho \in \{0, .25, .50, .75, .90\}$. For each donor, the corresponding count among 96 cells is clipped to $[4, 92]$ and rounded, and the phenotype vector is then shuffled; at $\rho = .9$ this yields 5 and 91 phenotype-positive cells in the two donor groups. Datasets are shared across methods at each ρ but are generated independently across values of ρ. The simulated phenotype is supplied as the biological covariate, and the donor code of the shortcut generator enters both assays with fixed strength 1.5. Unconditional MINUET applies the donor penalties without conditioning on phenotype.

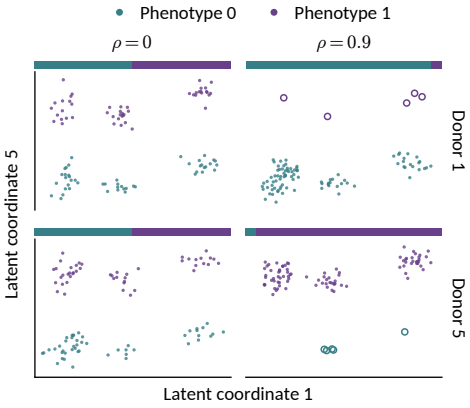

Figure 13: **Conditional suppression generator.** Latent coordinates 1 and 5, colored by phenotype, for donors 1 and 5 of seed 0 at $\rho = 0$ and $\rho = 0.9$; bars give each donor's phenotype fractions. Phenotype-1 counts go from 48 per donor to 5 and 91.

Reference-donor metrics use a donor- and phenotype-balanced sample of cells from the six reference donors. A linear phenotype probe is evaluated out of fold over four donor- and phenotype-balanced folds, and donor-prediction macro-F1 is evaluated separately within each phenotype group (random guess: $1/6$). Both metrics are summa-

rized by their trapezoidal mean over $\rho \in [0, .9]$.
Phenotype-prediction macro-F1 on query donors is evaluated at $\rho = .9$.

Ablations. Each panel of Figure 14 evaluates the ablation variants of Section B that act on one failure mode. In the same-donor shortcut, **`+ pooled candidates`** and **`– pair matching`** leave phenotype macro-F1, recall@1, and donor macro-F1 within 0.015 of **`Full`**, because within-donor centering already removes the additive donor cue. Under donor-specific transformations, removing shared prediction lowers latent-state R^2 by 0.119 and raises the registration gap by 0.052; removing pair matching also raises the gap, by 0.048. Under conditional suppression, removing the donor-association penalty raises within-phenotype donor macro-F1 by 0.145 and lowers phenotype macro-F1 by 0.024 on both reference and query donors. **`Unconditional`** lowers reference-donor phenotype macro-F1 by 0.038 and raises within-phenotype donor macro-F1 by 0.057, averaged over ρ.

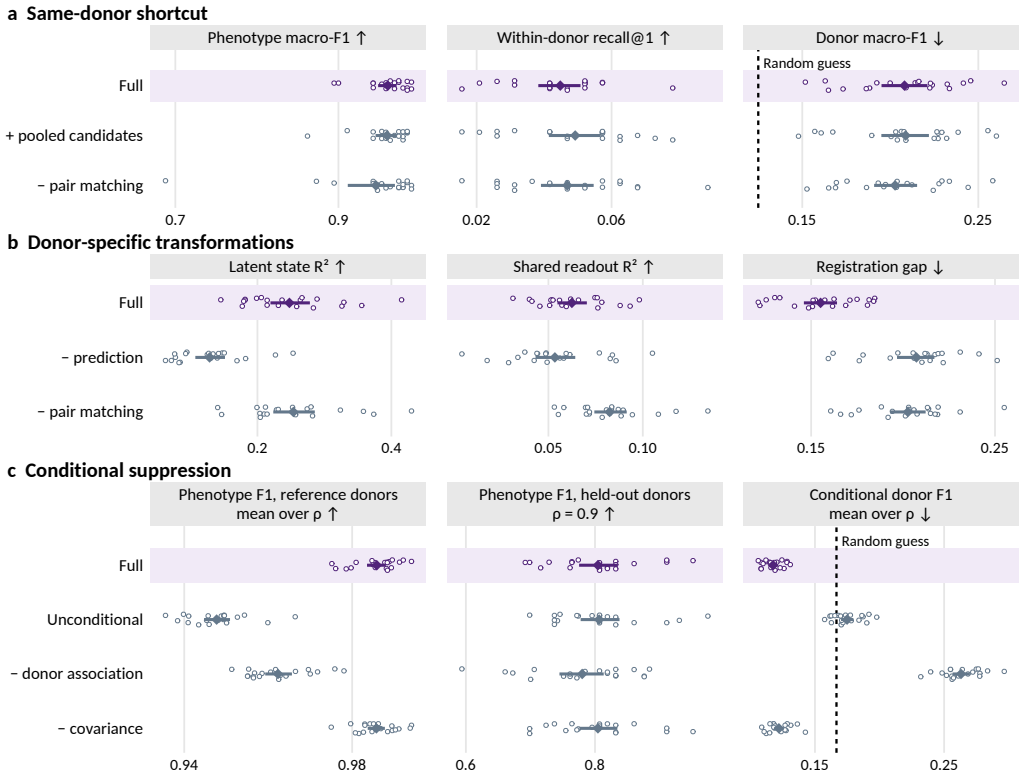

Figure 14: **Ablation study in the synthetic experiments.** Open circles show 20 shared seeds; diamonds and horizontal lines show means and 95% seed-bootstrap intervals for the same-donor shortcut, donor-specific transformations, and conditional suppression; shared-readout R^2 and registration gap are shown side by side. Dashed lines mark the random-guess level of donor macro-F1.

D Evaluation Protocols

Table 3 lists the cohorts, covariates, and phenotypes. The covariates define the groups of the donor penalties for observed-cell integration and held-out-donor transfer. Each donor contributes one sample in ROSMAP, alopecia, and rheumatoid arthritis. In Wang, 11 of 27 donors contribute two regional samples; in Mannens, six of eight donors contribute up to five dissections; and one Crohn donor was sequenced in two pools. In these cohorts, the donor penalties act on donor and sample effects jointly, with region or tissue among the covariates. Covariate groups containing a single donor, to which the donor penalties do not apply, occur only in Wang (one of ten groups) and Mannens (three of eight).

Table 3: **Paired multiome cohorts.** Biological covariates define the donor-penalty groups for observed-cell integration and held-out-donor transfer. Predicted phenotypes are withheld from representation fitting in the phenotype-prediction protocol. Total: 636,509 cells and 110 donors.

Cohort	System	Cells	Donors	Covariates	Phenotype
Wang	developing neocortex	232,328	27	stage × region	not evaluated
Mannens	first-trimester brain	166,785	8	tissue	not evaluated
ROSMAP	adult prefrontal cortex	97,070	19	pathology	Alzheimer status
Crohn	peripheral blood	72,457	28	disease status	disease status
Alopecia	peripheral blood	36,322	16	severity	severity
Rheumatoid arthritis	synovium	31,547	12	tissue	not evaluated

Reference fitting and query encoding. Feature selection, preprocessing, and fitting use reference cells only. Query embeddings average the outputs of the two frozen assay encoders. Decoders, the prediction head, the target bases and standardization constants, and covariates are used only during fitting; query mapping uses no labels, parameter updates, or centering across query cells. The exception is DC-PCA, which centers each query donor with the mean of that donor's cells.

Observed-cell and held-out-donor evaluation. Observed-cell integration evaluates the cells used for fitting with a scMultiBench-compatible metric panel over three representation seeds (Liu et al., 2025; Luecken et al., 2022; Liu, 2025). Its average rank combines ARI, NMI, CiLISI, and FOSCTTM; held-out-donor evaluation adds cell-type transfer macro-F1. For held-out-donor evaluation, all cells from query donors are reserved for evaluation. Cell-type labels are transferred by logistic regression, with regularization selected by donor-grouped cross-validation on reference cells. ARI and NMI measure cell-type conservation; CiLISI measures donor mixing through cell-type-adjusted donor iLISI. FOSCTTM measures paired-cell matching, with lower values indicating better matching (Liu et al., 2019; Klein et al., 2025).

LODO transfer. Wang, Mannens, and ROSMAP define three $A + B \to C$ directions. The two source datasets determine the shared gene and peak sets, preprocessing, representation, and cell-type classifier. Each direction is reported separately with the five metrics of held-out-donor evaluation, and uncertainty is estimated by resampling query donors.

Phenotype prediction. We predicted Alzheimer disease status in ROSMAP, disease status in Crohn, and severity in alopecia areata by five-fold cross-validation over donors, stratified by phenotype, with three representation seeds. In each fold, the representation and the classifier are fitted on the training donors; the classifier uses donor-level summaries of cell embeddings and cell-type proportions. The covariates in Table 3 coincide with or determine the predicted phenotype and are therefore withheld, so the reference donors of each fold form a single group for the donor penalties. Out-of-fold probabilities are averaged across seeds before computing macro-AUROC and macro-F1. The protocol average rank combines these metrics with CiLISI and FOSCTTM, computed from the out-of-fold cell embeddings.

Statistical analysis. Donors, rather than cells, folds, or representation seeds, are treated as the independent units. Intervals for query mapping resample donors, and phenotype comparisons use paired bootstraps over donors' seed-averaged out-of-fold probabilities. Protocol summaries weight cohorts equally, and a method's average rank leaves out the metrics reported as N/A for it, which is FOSCTTM for Apollo and DC-PCA. Ranks are computed within cohort and metric before averaging over all displayed metrics, including FOSCTTM.

Annotation sensitivity. For alopecia, Azimuth annotation agreed with 99.4% of the level-2 cell-type labels. Excluding the 2,808 of 20,000 sampled cells whose labels were not recovered changed the CiLISI of scPair from 0.0767 to 0.0789 and that of scPairing from 0.0566 to 0.0560, without changing ranks; the main comparisons use labels for all cells. ROSMAP cell-type and pathology labels were rederived from the raw data.

Table 4: **Unspecified covariates in the conditional suppression experiment.** Each entry gives the macro-F1 mean over 20 shared seeds, after averaging each seed over the five values of ρ, with the 95% seed-bootstrap interval below. Reference-donor Y_2 uses leave-one-donor-out probing; query-donor Y_2 applies a probe fitted on all reference cells. ΔY is the seed-paired change in reference-donor Y macro-F1 relative to $C = \{Y\}$. Bold marks the best setting; darker purple indicates a better value.

Covariates C	Y_2 reference donors ↑	Y_2 query donors ↑	Y reference donors ↑	ΔY ↑
$\{Y\}$	0.491 [0.484, 0.496]	0.644 [0.628, 0.661]	**0.985** [0.983, 0.987]	reference
$\{Y, Y_2\}$	**0.690** [0.683, 0.697]	**0.887** [0.875, 0.897]	0.963 [0.960, 0.966]	**−0.022** [−0.024, −0.020]
Unconditional	0.490 [0.484, 0.496]	0.636 [0.621, 0.652]	0.947 [0.943, 0.951]	−0.038 [−0.042, −0.035]

E What Donor Information Contains

The donor label D is an index, not a cause. What a representation can learn about a donor is the donor-level vector $W_d = (A_d, G_d, U_d, Y_d, B_d)$ of Figure 15, and the question is which parts of W the penalties remove.

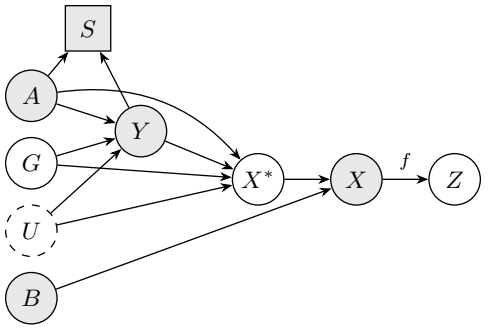

Figure 15: **What a donor label carries.** Donor-level demographics A, genotype G, unmeasured environment U (dashed), phenotype Y, and technical factors B; cell state X^*, paired measurements $X = (X^{\mathrm{R}}, X^{\mathrm{A}})$, and representation Z. Recruitment S is conditioned on by design. Gray: recorded in at least some cohorts.

Decomposition. Every path from the donor to a cell passes through W_D, so $I(Z; D) = I(Z; W_D) = I(Z; C) + I(Z; W_{-C} \mid C)$, where W_{-C} collects the unrecorded components. The conditional penalties act on the second term and leave the first (Section A.5). That term mixes unrecorded biology (G, U, and any part of A or Y outside C) with technical variation B; within a covariate group both are donor-level shifts, and no criterion based on (X, D, C) alone can separate them. As in Theorem 3.1, variation not tied to C is not identified as biology.

Confounding among covariates. Recruitment S depends on both A and Y, so conditioning on it can make them dependent among enrolled donors (Pearl, 2009). With $C = \{Y\}$, the penalties then remove A within phenotype groups but keep its imbalance between groups, which a downstream analysis would attribute to Y. Adding A to C splits the groups and shrinks the cliques of the donor graph. Technical factors nested in donors (Section D) are aliased with D and handled like any other donor effect; if B is aligned with Y by design, no representation fitted to such data can separate it from disease.

Unspecified covariates in simulation. A variant of the conditional suppression generator (Section C) adds a second donor-level phenotype Y_2, balanced across donors and independent of Y, that shifts a separate latent coordinate by the same magnitude as Y; everything else is unchanged. Table 4 compares conditioning on Y alone, on both phenotypes, and on neither, over 20 seeds. When Y_2 is not included in C, its reference-donor macro-F1 is 0.49, the same as without conditioning, whereas including it raises this value to 0.69 on all 20 seeds, at a cost of 0.022 in reference-donor Y macro-F1. On the two query donors, Y_2 remains partly predictable without conditioning on it (0.64 versus 0.89), because the penalties equalize reference-donor means rather than remove the direction.

Covariate choice on Crohn. Table 5 varies the covariates that define the groups of both donor penalties on Crohn, with everything else fixed: none, disease status, disease status crossed with sex, and sex alone. Disease status reaches leave-one-donor-out macro-F1 of 0.585 and 0.562 when included, versus 0.510 and 0.513 when not; sex reaches 0.527 and 0.541 when included, versus 0.490

Table 5: **Choice of conditioning covariates on Crohn.** Only the covariates defining the groups of the two donor penalties change; the configuration, seeds, and evaluation follow the ablation study. Each entry gives the mean over three seeds with the standard deviation in parentheses and, in brackets, the seed-paired Student-t 95% interval for the difference from the disease-status setting. Disease-status and sex macro-F1 are leave-one-donor-out probes on reference donors; donor macro-F1 within disease status measures excess donor information. Bold marks the best setting; darker purple indicates a better value.

Covariates C	Disease status macro-F1 ↑	Sex macro-F1 ↑	Donor macro-F1 within disease status ↓
None	0.510 (0.011) [−0.070, −0.036]	0.488 (0.014) [−0.021, 0.017]	0.382 (0.007) [−0.001, 0.027]
Disease status	**0.585** (0.017) reference	0.490 (0.011) reference	**0.369** (0.006) reference
Disease status and sex	0.562 (0.002) [−0.039, 0.037]	0.527 (0.010) [0.025, 0.050]	0.388 (0.013) [−0.027, 0.064]
Sex	0.513 (0.010) [−0.069, −0.031]	**0.541** (0.017) [0.033, 0.070]	0.381 (0.006) [−0.014, 0.038]

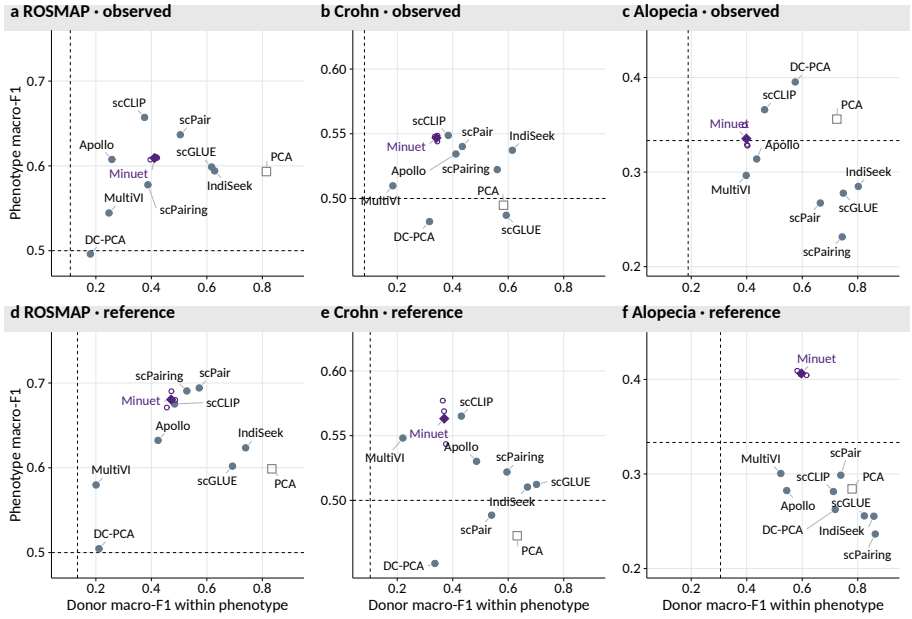

Figure 16: **Phenotype and excess donor information.** Each point is one method's seed-0 fitted representation (MINUET: three seeds and their mean; open square: unintegrated PCA), probed on a donor-balanced sample of 192 cells per donor. The horizontal axis is within-phenotype donor macro-F1, which measures the excess donor information targeted by the conditional penalties; the vertical axis is phenotype macro-F1 under leave-one-donor-out probing. Dashed lines mark the random-guess level. Rows are observed cells and the reference donors of the held-out protocol; columns are ROSMAP Alzheimer disease status, Crohn disease status, and alopecia severity.

and 0.488. A covariate left out of C is therefore treated as excess donor information, even when, like sex, it is biological. Crossing disease status with sex slightly raises donor macro-F1 within disease status (0.388 versus 0.369), because finer groups contain fewer donors; exact stratification on many covariates soon leaves groups too small to compare. Cell-type transfer macro-F1, CiLISI, and FOSCTTM are unchanged across settings.

Donor information on real data. The probes in Figure 16 are logistic regressions on standardized features with balanced class weights, fitted on a donor-balanced sample of 192 cells per donor that is drawn once and shared across methods. Donor identity is probed with five stratified folds; excess donor information with four folds within each phenotype group, averaged over groups; and phenotype with leave-one-donor-out folds, so that no donor is used for both fitting and evaluation.

Methods that remove donor information marginally have low within-phenotype donor macro-F1 but predict phenotype near the random-guess level: on ROSMAP, DC-PCA predicts Alzheimer disease status with macro-F1 0.50 (random guess 0.5), whereas MINUET reaches 0.61 on observed cells and 0.68 on the reference donors of the held-out-donor protocol. For alopecia severity on the reference donors, MINUET reaches 0.41 (random guess 0.33), and every other method is at or below that level; for Crohn, all methods are near the random-guess level. Within-phenotype donor prediction stays above the random-guess level for every method and panel, consistent with penalties that attenuate excess donor information rather than remove it.

Toward exogenous donor variation. Recorded covariates are themselves shaped by recruitment and environment. Germline genotype is not: it is fixed at conception, and autosomal variants are independent of sex. Mendelian randomization uses this property to treat genetic variants as instruments for confounded exposures (Davey Smith & Ebrahim, 2003; Lawlor et al., 2008; Angrist et al., 1996). Describing donors by genotype instead of phenotype labels would let the penalties keep only donor variation upstream of A, U, and S. Genotype effects on cell state are detectable at single-cell resolution (Yazar et al., 2022), the relevance condition of an instrument; ancestry must enter C (Price et al., 2006), and selection on disease can still link G to U. Because genotype and age are continuous, discrete groups would be replaced by a kernel similarity $k(C_d, C_{d'})$, giving a weighted donor graph and a kernel conditional-dependence penalty (Fukumizu et al., 2007). Whether Equation (3) extends to this graph is left for future work.

F DETAILED RESULTS

Table 6 summarizes each evaluation protocol with cohorts weighted equally. The tables that follow report every metric by cohort, ordered by protocol.

Metrics. The tables by cohort add secondary metrics of biological conservation, donor mixing, and paired-cell matching, and two measures of residual donor information. Donor prediction fits a standardized logistic regression with balanced class weights on half of a donor-balanced sample of up to 192 query cells per donor and scores the other half with macro-F1, overall and within each covariate group. Within-group donor dispersion is the mean squared distance of each donor's mean embedding from the average of donor means within a covariate group, averaged over groups and donors. Silhouette scores are reported but excluded from average ranks, because they can misrepresent both batch removal and biological conservation (Rautenstrauch & Ohler, 2026).

Reading the tables. Each entry gives the mean over three representation seeds above its standard deviation; phenotype-prediction entries use probabilities averaged over the seeds, and donor-prediction entries use one seed. Within a cohort and metric, bold and underlining mark the best and second-best results, darker purple a better score, and arrows the preferred direction; gray rows are the unranked in-house controls. Paired-cell matching is N/A for Apollo, whose second training stage fits both assay encoders to one shared code per cell rather than training them to match each other, and for DC-PCA, which has no shared RNA and ATAC space. Metrics that a protocol cannot estimate are omitted. Graph conn.: graph connectivity.

Table 6: **Numerical summaries across evaluation protocols.** Metric values are means with cohorts weighted equally. Avg. rank averages within-cohort ranks over all displayed metrics, including FOSCTTM. Darker purple indicates a better value. Bold and underlining mark the best and second-best results. Apollo uses binary inputs except for observed-cell and held-out-donor Crohn and rheumatoid arthritis and for phenotype prediction on Crohn. FOSCTTM is N/A for Apollo and DC-PCA, whose average ranks use the other metrics.

Method	Observed cells					Held-out donors					
	ARI ↑	NMI ↑	CiLISI ↑	FOSCTTM ↓	Avg. rank ↓	ARI ↑	NMI ↑	CiLISI ↑	F1 ↑	FOSCTTM ↓	Avg. rank ↓
Minuet	**0.465**	**0.677**	0.223	**0.027**	**1.67**	**0.440**	**0.642**	0.272	0.640	0.191	**2.83**
DC-PCA	0.357	0.518	0.207	N/A	5.11	0.337	0.519	0.264	0.560	N/A	5.71
IndiSeek	0.444	0.626	0.133	0.463	5.54	0.436	0.630	0.230	**0.737**	0.473	4.27
scGLUE	0.357	0.552	0.176	0.284	5.42	0.359	0.578	0.263	0.655	0.233	4.37
scPairing	0.397	0.610	0.133	0.159	4.33	0.413	0.623	0.232	0.565	**0.165**	4.00
Apollo	0.395	0.555	0.214	N/A	4.06	0.336	0.510	0.367	0.511	N/A	4.88
MultiVI	0.356	0.516	**0.256**	0.217	4.88	0.303	0.508	**0.368**	0.442	0.211	5.13
scPair	0.374	0.589	0.137	0.169	5.54	0.344	0.552	0.249	0.553	0.211	5.30
scCLIP	0.229	0.452	0.135	0.084	6.50	0.222	0.446	0.278	0.286	0.315	7.23

Method	LODO transfer						Phenotype prediction				
	ARI ↑	NMI ↑	CiLISI ↑	F1 ↑	FOSCTTM ↓	Avg. rank ↓	AUROC ↑	F1 ↑	CiLISI ↑	FOSCTTM ↓	Avg. rank ↓
Minuet	**0.640**	**0.689**	0.195	0.514	**0.234**	**2.73**	**0.623**	**0.617**	0.090	0.171	**2.33**
DC-PCA	0.242	0.364	0.204	0.374	N/A	7.83	0.592	0.518	**0.191**	N/A	3.67
IndiSeek	0.550	0.650	0.138	**0.540**	0.403	5.33	0.591	0.536	0.064	0.493	5.75
scGLUE	0.582	0.657	0.189	0.516	0.316	3.67	0.502	0.468	0.079	0.284	5.75
scPairing	0.593	0.675	0.180	0.507	0.274	3.87	0.525	0.552	0.064	**0.170**	4.42
Apollo	0.508	0.598	**0.233**	0.431	N/A	5.42	0.594	0.339	0.188	N/A	4.56
MultiVI	0.500	0.611	0.220	0.483	0.279	4.27	0.389	0.429	0.098	0.259	5.58
scPair	0.461	0.576	0.216	0.489	0.327	5.53	0.545	0.492	0.072	0.211	5.25
scCLIP	0.486	0.584	0.210	0.414	0.293	5.60	0.588	0.482	0.090	0.357	5.08

Table 7: **Observed cells: biological conservation.**

Method	Wang			Mannens			ROSMAP			Crohn			Alopecia			Rheumatoid arthritis		
	cLISI ↑	Graph conn. ↑	Cell-type ASW ↑	cLISI ↑	Graph conn. ↑	Cell-type ASW ↑	cLISI ↑	Graph conn. ↑	Cell-type ASW ↑	cLISI ↑	Graph conn. ↑	Cell-type ASW ↑	cLISI ↑	Graph conn. ↑	Cell-type ASW ↑	cLISI ↑	Graph conn. ↑	Cell-type ASW ↑
Minuet	0.988 (0.000)	0.932 (0.000)	0.550 (0.003)	0.997 (0.000)	0.976 (0.000)	0.546 (0.003)	**1.000** (0.000)	0.938 (0.016)	0.555 (0.002)	0.986 (0.000)	0.877 (0.001)	0.532 (0.003)	**0.987** (0.000)	0.800 (0.011)	0.546 (0.001)	0.998 (0.000)	**0.964** (0.001)	0.539 (0.014)
Apollo	0.973 (0.001)	0.929 (0.001)	0.519 (0.001)	0.992 (0.001)	**0.982** (0.004)	0.524 (0.001)	0.923 (0.005)	0.587 (0.010)	0.467 (0.006)	0.946 (0.002)	**0.972** (0.003)	0.503 (0.001)	0.940 (0.004)	0.833 (0.005)	0.494 (0.009)	0.939 (0.002)	0.955 (0.007)	0.508 (0.003)
scPair	0.968 (0.001)	0.868 (0.003)	0.610 (0.007)	0.997 (0.000)	0.961 (0.002)	0.605 (0.012)	**1.000** (0.000)	0.872 (0.016)	0.585 (0.011)	**0.989** (0.003)	0.943 (0.006)	0.620 (0.007)	0.979 (0.001)	0.698 (0.018)	0.562 (0.007)	0.998 (0.000)	0.947 (0.004)	0.703 (0.050)
scPairing	0.973 (0.000)	0.895 (0.003)	0.576 (0.004)	**0.998** (0.000)	0.975 (0.000)	0.617 (0.005)	**1.000** (0.000)	0.886 (0.006)	0.558 (0.006)	0.987 (0.000)	0.956 (0.001)	0.552 (0.007)	0.980 (0.000)	0.732 (0.008)	0.545 (0.003)	0.998 (0.000)	0.957 (0.001)	0.625 (0.003)
scCLIP	0.953 (0.000)	0.687 (0.004)	**0.683** (0.004)	0.997 (0.001)	0.910 (0.029)	0.602 (0.014)	0.992 (0.002)	0.430 (0.004)	**0.686** (0.008)	0.970 (0.001)	0.665 (0.029)	**0.800** (0.005)	0.977 (0.000)	0.404 (0.011)	**0.836** (0.002)	0.975 (0.009)	0.663 (0.014)	0.640 (0.026)
IndiSeek	**0.991** (0.000)	**0.936** (0.001)	0.536 (0.003)	0.997 (0.001)	0.934 (0.029)	0.523 (0.014)	**1.000** (0.000)	**0.962** (0.010)	0.534 (0.008)	0.975 (0.005)	0.946 (0.006)	0.533 (0.003)	0.969 (0.003)	**0.850** (0.012)	0.508 (0.001)	**0.999** (0.000)	**0.964** (0.004)	0.575 (0.014)
MultiVI	0.943 (0.004)	0.687 (0.043)	0.602 (0.006)	0.993 (0.000)	0.877 (0.007)	**0.775** (0.017)	0.999 (0.000)	0.536 (0.010)	0.606 (0.005)	0.941 (0.009)	0.668 (0.039)	0.629 (0.005)	0.947 (0.008)	0.429 (0.013)	0.590 (0.006)	0.998 (0.000)	0.912 (0.010)	**0.810** (0.031)
DC-PCA	0.970 (0.000)	0.888 (0.002)	0.420 (0.004)	0.973 (0.002)	0.976 (0.001)	0.512 (0.003)	0.983 (0.001)	0.832 (0.004)	0.465 (0.027)	0.943 (0.002)	0.871 (0.002)	0.435 (0.001)	0.963 (0.001)	0.702 (0.003)	0.478 (0.008)	0.875 (0.005)	0.918 (0.003)	0.482 (0.024)
scGLUE	0.969 (0.001)	0.883 (0.003)	0.589 (0.003)	0.980 (0.021)	0.961 (0.004)	0.581 (0.018)	0.991 (0.006)	0.887 (0.007)	0.560 (0.004)	0.986 (0.002)	0.946 (0.003)	0.574 (0.004)	0.983 (0.001)	0.735 (0.030)	0.556 (0.001)	0.978 (0.016)	0.957 (0.005)	0.594 (0.020)

Method	Wang		Mannens		ROSMAP		Crohn		Alopecia		Rheumatoid arthritis	
	ARI ↑	NMI ↑	ARI ↑	NMI ↑	ARI ↑	NMI ↑	ARI ↑	NMI ↑	ARI ↑	NMI ↑	ARI ↑	NMI ↑
Minuet	**0.440** (0.008)	**0.650** (0.004)	0.541 (0.105)	**0.618** (0.036)	0.324 (0.011)	0.759 (0.006)	**0.499** (0.027)	**0.667** (0.004)	0.219 (0.005)	**0.531** (0.003)	0.766 (0.036)	0.835 (0.022)
Apollo	0.339 (0.023)	0.540 (0.009)	0.434 (0.002)	0.524 (0.002)	0.230 (0.023)	0.421 (0.010)	0.424 (0.036)	0.623 (0.015)	**0.337** (0.048)	0.523 (0.016)	0.604 (0.052)	0.698 (0.024)
scPair	0.327 (0.013)	0.535 (0.008)	0.422 (0.011)	0.519 (0.008)	0.199 (0.007)	0.582 (0.009)	0.398 (0.037)	0.610 (0.016)	0.179 (0.013)	0.472 (0.008)	0.718 (0.024)	0.818 (0.007)
scPairing	0.349 (0.003)	0.573 (0.003)	0.415 (0.044)	0.549 (0.019)	0.252 (0.013)	0.601 (0.003)	0.413 (0.034)	0.621 (0.018)	0.193 (0.008)	0.481 (0.006)	0.759 (0.037)	0.833 (0.014)
scCLIP	0.220 (0.003)	0.464 (0.003)	0.201 (0.028)	0.375 (0.028)	0.128 (0.001)	0.433 (0.008)	0.301 (0.011)	0.508 (0.005)	0.157 (0.003)	0.421 (0.005)	0.366 (0.142)	0.512 (0.102)
IndiSeek	0.439 (0.007)	0.599 (0.000)	0.361 (0.127)	0.491 (0.056)	**0.422** (0.049)	**0.768** (0.018)	0.388 (0.027)	0.585 (0.031)	0.180 (0.014)	0.424 (0.015)	**0.876** (0.067)	**0.890** (0.025)
MultiVI	0.272 (0.012)	0.447 (0.016)	**0.567** (0.018)	0.593 (0.006)	0.176 (0.008)	0.473 (0.006)	0.268 (0.045)	0.459 (0.042)	0.091 (0.015)	0.295 (0.023)	0.763 (0.041)	0.830 (0.032)
DC-PCA	0.436 (0.016)	0.567 (0.005)	0.273 (0.002)	0.461 (0.000)	0.300 (0.005)	0.496 (0.004)	0.394 (0.018)	0.544 (0.015)	0.192 (0.017)	0.415 (0.006)	0.544 (0.002)	0.629 (0.006)
scGLUE	0.347 (0.021)	0.520 (0.005)	0.334 (0.114)	0.481 (0.068)	0.282 (0.010)	0.530 (0.034)	0.404 (0.014)	0.613 (0.008)	0.180 (0.003)	0.491 (0.007)	0.596 (0.185)	0.676 (0.095)

Table 8: **Observed cells: donor mixing.**

Method	Wang		Mannens		ROSMAP		Crohn		Alopecia		Rheumatoid arthritis	
	Donor iLISI ↑	Donor ASW ↑	Donor iLISI ↑	Donor ASW ↑	Donor iLISI ↑	Donor ASW ↑	Donor iLISI ↑	Donor ASW ↑	Donor iLISI ↑	Donor ASW ↑	Donor iLISI ↑	Donor ASW ↑
Minuet	0.087 (0.000)	0.900 (0.001)	0.161 (0.001)	0.915 (0.000)	0.262 (0.001)	0.878 (0.003)	0.206 (0.002)	0.893 (0.004)	0.355 (0.001)	0.865 (0.003)	0.295 (0.005)	0.906 (0.003)
Apollo	0.053 (0.002)	**0.939** (0.000)	0.053 (0.001)	0.953 (0.001)	**0.360** (0.003)	**0.939** (0.002)	0.213 (0.001)	**0.933** (0.002)	**0.397** (0.007)	**0.947** (0.002)	0.251 (0.004)	**0.959** (0.002)
scPair	0.058 (0.001)	0.847 (0.007)	0.019 (0.003)	0.821 (0.007)	0.187 (0.014)	0.812 (0.022)	0.138 (0.008)	0.754 (0.002)	0.248 (0.008)	0.780 (0.006)	0.164 (0.014)	0.855 (0.004)
scPairing	0.059 (0.000)	0.903 (0.001)	0.061 (0.001)	0.864 (0.002)	0.234 (0.002)	0.829 (0.002)	0.120 (0.001)	0.889 (0.001)	0.174 (0.002)	0.874 (0.002)	0.128 (0.002)	0.912 (0.003)
scCLIP	0.047 (0.001)	0.809 (0.002)	0.004 (0.002)	0.589 (0.016)	0.232 (0.002)	0.699 (0.010)	0.149 (0.003)	0.597 (0.017)	0.267 (0.005)	0.623 (0.009)	0.130 (0.007)	0.648 (0.010)
IndiSeek	0.041 (0.003)	0.928 (0.002)	0.009 (0.001)	0.908 (0.004)	0.190 (0.010)	0.893 (0.004)	0.143 (0.004)	0.855 (0.005)	0.235 (0.014)	0.888 (0.010)	0.168 (0.006)	0.888 (0.006)
MultiVI	**0.141** (0.014)	0.751 (0.018)	**0.248** (0.008)	0.734 (0.014)	0.310 (0.001)	0.696 (0.003)	**0.245** (0.004)	0.585 (0.008)	0.347 (0.009)	0.593 (0.016)	**0.308** (0.003)	0.786 (0.007)
DC-PCA	0.074 (0.001)	0.851 (0.000)	0.176 (0.003)	**0.954** (0.001)	0.302 (0.001)	0.840 (0.002)	0.171 (0.001)	0.825 (0.002)	0.301 (0.002)	0.856 (0.001)	0.262 (0.005)	0.843 (0.002)
scGLUE	0.060 (0.002)	0.890 (0.003)	0.156 (0.020)	0.904 (0.009)	0.236 (0.007)	0.896 (0.005)	0.149 (0.012)	0.845 (0.010)	0.235 (0.005)	0.836 (0.002)	0.236 (0.017)	0.916 (0.005)

Method	Wang		Mannens		ROSMAP		Crohn		Alopecia		Rheumatoid arthritis	
	Donor mixing ARI ↑	Donor mixing NMI ↑	Donor mixing ARI ↑	Donor mixing NMI ↑	Donor mixing ARI ↑	Donor mixing NMI ↑	Donor mixing ARI ↑	Donor mixing NMI ↑	Donor mixing ARI ↑	Donor mixing NMI ↑	Donor mixing ARI ↑	Donor mixing NMI ↑
Minuet	0.892 (0.004)	0.699 (0.004)	0.944 (0.005)	0.861 (0.005)	0.961 (0.004)	0.926 (0.002)	0.962 (0.002)	0.906 (0.001)	0.978 (0.001)	**0.949** (0.001)	0.952 (0.002)	0.902 (0.001)
Apollo	0.847 (0.005)	0.602 (0.004)	0.854 (0.019)	0.715 (0.006)	0.939 (0.001)	0.896 (0.002)	0.899 (0.010)	0.823 (0.003)	0.964 (0.001)	0.931 (0.003)	0.915 (0.001)	0.835 (0.005)
scPair	0.814 (0.006)	0.562 (0.005)	0.833 (0.029)	0.707 (0.031)	0.895 (0.005)	0.813 (0.013)	0.919 (0.004)	0.786 (0.009)	0.956 (0.003)	0.902 (0.003)	0.902 (0.013)	0.808 (0.021)
scPairing	0.836 (0.003)	0.590 (0.007)	0.883 (0.018)	0.778 (0.033)	0.964 (0.004)	0.915 (0.000)	0.908 (0.006)	0.762 (0.011)	0.922 (0.004)	0.833 (0.008)	0.881 (0.005)	0.782 (0.008)
scCLIP	0.775 (0.010)	0.510 (0.007)	0.622 (0.035)	0.493 (0.023)	0.910 (0.003)	0.810 (0.005)	0.884 (0.001)	0.755 (0.004)	0.940 (0.006)	0.854 (0.001)	0.860 (0.023)	0.708 (0.018)
IndiSeek	0.808 (0.010)	0.556 (0.011)	0.876 (0.027)	0.751 (0.043)	0.938 (0.015)	0.870 (0.033)	0.928 (0.004)	0.822 (0.009)	0.943 (0.010)	0.872 (0.010)	0.938 (0.004)	0.867 (0.004)
MultiVI	**0.932** (0.006)	**0.800** (0.026)	0.940 (0.007)	**0.887** (0.005)	0.966 (0.003)	0.926 (0.001)	**0.976** (0.003)	**0.928** (0.005)	**0.981** (0.001)	**0.949** (0.002)	0.946 (0.003)	0.896 (0.004)
DC-PCA	0.872 (0.002)	0.664 (0.004)	**0.993** (0.001)	0.864 (0.010)	**0.986** (0.001)	**0.948** (0.007)	0.972 (0.001)	0.905 (0.003)	0.969 (0.000)	0.895 (0.003)	**0.964** (0.002)	0.894 (0.005)
scGLUE	0.829 (0.007)	0.580 (0.005)	0.920 (0.038)	0.842 (0.044)	0.974 (0.017)	0.937 (0.030)	0.934 (0.016)	0.834 (0.027)	0.960 (0.000)	0.905 (0.002)	0.953 (0.008)	**0.904** (0.010)

Table 9: **Observed cells: paired-cell matching.**

Method	Wang			Mannens			ROSMAP			Crohn			Alopecia			Rheumatoid arthritis		
	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑
Minuet	**0.072** (0.001)	**0.208** (0.002)	**0.101** (0.002)	**0.036** (0.001)	**0.472** (0.001)	**0.335** (0.001)	**0.036** (0.000)	**0.305** (0.004)	**0.171** (0.004)	**0.007** (0.000)	**0.684** (0.011)	**0.560** (0.013)	**0.002** (0.000)	**0.806** (0.025)	**0.693** (0.035)	**0.007** (0.001)	**0.600** (0.015)	**0.452** (0.015)
Apollo	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
scPair	0.188 (0.010)	0.069 (0.002)	0.021 (0.001)	0.139 (0.008)	0.103 (0.006)	0.037 (0.004)	0.184 (0.005)	0.058 (0.003)	0.015 (0.001)	0.190 (0.010)	0.063 (0.005)	0.017 (0.002)	0.154 (0.004)	0.079 (0.002)	0.026 (0.002)	0.159 (0.003)	0.069 (0.003)	0.019 (0.003)
scPairing	0.188 (0.001)	0.063 (0.001)	0.018 (0.000)	0.139 (0.003)	0.086 (0.000)	0.027 (0.002)	0.225 (0.003)	0.040 (0.001)	0.009 (0.001)	0.134 (0.001)	0.071 (0.001)	0.019 (0.001)	0.125 (0.001)	0.094 (0.000)	0.032 (0.001)	0.145 (0.001)	0.075 (0.001)	0.022 (0.001)
scCLIP	0.122 (0.001)	0.119 (0.001)	0.049 (0.000)	0.096 (0.006)	0.151 (0.006)	0.069 (0.004)	0.077 (0.001)	0.167 (0.004)	0.065 (0.003)	0.071 (0.002)	0.223 (0.009)	0.114 (0.007)	0.054 (0.001)	0.271 (0.005)	0.138 (0.005)	0.088 (0.010)	0.187 (0.019)	0.088 (0.010)
IndiSeek	0.378 (0.009)	0.020 (0.003)	0.003 (0.001)	0.426 (0.005)	0.018 (0.001)	0.003 (0.000)	0.495 (0.000)	0.014 (0.000)	0.002 (0.000)	0.488 (0.002)	0.014 (0.000)	0.002 (0.000)	0.492 (0.004)	0.014 (0.000)	0.002 (0.000)	0.498 (0.001)	0.015 (0.000)	0.002 (0.000)
MultiVI	0.235 (0.007)	0.042 (0.002)	0.010 (0.001)	0.184 (0.004)	0.052 (0.002)	0.012 (0.001)	0.217 (0.003)	0.040 (0.001)	0.009 (0.001)	0.250 (0.010)	0.037 (0.002)	0.009 (0.001)	0.268 (0.005)	0.043 (0.004)	0.011 (0.001)	0.147 (0.002)	0.071 (0.003)	0.020 (0.002)
DC-PCA	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
scGLUE	0.251 (0.014)	0.047 (0.003)	0.012 (0.001)	0.376 (0.106)	0.034 (0.011)	0.009 (0.003)	0.414 (0.049)	0.021 (0.002)	0.004 (0.001)	0.172 (0.017)	0.057 (0.009)	0.014 (0.004)	0.150 (0.011)	0.074 (0.007)	0.023 (0.004)	0.339 (0.075)	0.033 (0.011)	0.007 (0.003)
Pooled	0.099 (0.003)	0.140 (0.005)	0.057 (0.004)	0.065 (0.002)	0.236 (0.010)	0.123 (0.009)	0.057 (0.002)	0.257 (0.016)	0.141 (0.016)	0.020 (0.000)	0.489 (0.004)	0.353 (0.004)	0.003 (0.000)	0.757 (0.002)	0.632 (0.002)	0.011 (0.001)	0.582 (0.017)	0.446 (0.020)
Within-donor	0.136 (0.001)	0.080 (0.002)	0.022 (0.001)	0.093 (0.003)	0.117 (0.007)	0.039 (0.004)	0.095 (0.003)	0.109 (0.003)	0.035 (0.002)	0.037 (0.005)	0.248 (0.035)	0.121 (0.027)	0.011 (0.000)	0.431 (0.010)	0.252 (0.011)	0.023 (0.001)	0.373 (0.014)	0.236 (0.012)

Table 10: **Held-out donors: paired-cell matching.**

Method	Wang			Mannens			ROSMAP			Crohn			Alopecia			Rheumatoid arthritis		
	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑
Minuet	**0.157** (0.005)	**0.089** (0.002)	**0.030** (0.003)	**0.133** (0.006)	**0.113** (0.004)	**0.044** (0.003)	0.243 (0.003)	**0.064** (0.003)	**0.019** (0.002)	0.215 (0.007)	0.071 (0.006)	**0.023** (0.003)	0.168 (0.005)	**0.104** (0.004)	**0.040** (0.004)	0.229 (0.006)	**0.069** (0.002)	**0.023** (0.003)
Apollo	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
scPair	0.256 (0.012)	0.049 (0.004)	0.014 (0.002)	0.209 (0.011)	0.056 (0.004)	0.017 (0.004)	0.169 (0.004)	0.060 (0.002)	0.016 (0.004)	0.203 (0.013)	0.055 (0.005)	0.013 (0.003)	0.197 (0.032)	0.059 (0.006)	0.017 (0.003)	0.233 (0.010)	0.045 (0.001)	0.010 (0.001)
scPairing	0.213 (0.002)	0.051 (0.000)	0.014 (0.000)	0.176 (0.004)	0.056 (0.001)	0.015 (0.001)	**0.157** (0.001)	0.061 (0.003)	0.016 (0.004)	**0.121** (0.001)	**0.074** (0.001)	0.021 (0.000)	0.142 (0.003)	0.080 (0.001)	0.025 (0.001)	**0.182** (0.011)	0.057 (0.003)	0.015 (0.003)
scCLIP	0.270 (0.003)	0.046 (0.002)	0.013 (0.002)	0.266 (0.009)	0.044 (0.006)	0.013 (0.003)	0.301 (0.005)	0.047 (0.002)	0.010 (0.002)	0.298 (0.005)	0.044 (0.002)	0.011 (0.001)	0.364 (0.001)	0.044 (0.002)	0.011 (0.002)	0.391 (0.010)	0.032 (0.002)	0.007 (0.004)
IndiSeek	0.435 (0.010)	0.018 (0.002)	0.003 (0.001)	0.437 (0.015)	0.018 (0.001)	0.004 (0.000)	0.489 (0.002)	0.014 (0.001)	0.002 (0.000)	0.487 (0.002)	0.014 (0.000)	0.002 (0.000)	0.494 (0.005)	0.014 (0.000)	0.002 (0.000)	0.499 (0.002)	0.014 (0.000)	0.002 (0.000)
MultiVI	0.244 (0.006)	0.039 (0.004)	0.010 (0.003)	0.208 (0.016)	0.044 (0.007)	0.013 (0.003)	0.185 (0.004)	0.046 (0.004)	0.009 (0.003)	0.227 (0.032)	0.038 (0.003)	0.008 (0.000)	0.177 (0.024)	0.059 (0.008)	0.015 (0.003)	0.226 (0.012)	0.046 (0.005)	0.010 (0.002)
DC-PCA	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
scGLUE	0.265 (0.005)	0.045 (0.001)	0.012 (0.000)	0.273 (0.041)	0.041 (0.007)	0.010 (0.002)	0.208 (0.011)	0.048 (0.001)	0.011 (0.002)	0.252 (0.131)	0.040 (0.014)	0.008 (0.003)	**0.140** (0.002)	0.081 (0.001)	0.026 (0.001)	0.263 (0.025)	0.035 (0.002)	0.008 (0.002)
Pooled	0.182 (0.006)	0.075 (0.004)	0.026 (0.001)	0.152 (0.007)	0.076 (0.004)	0.024 (0.006)	0.181 (0.004)	0.066 (0.003)	0.018 (0.002)	0.149 (0.001)	0.075 (0.000)	0.021 (0.001)	0.166 (0.004)	0.081 (0.002)	0.028 (0.002)	0.187 (0.009)	0.068 (0.004)	0.021 (0.002)
Within-donor	0.209 (0.012)	0.050 (0.003)	0.011 (0.001)	0.149 (0.005)	0.070 (0.004)	0.020 (0.005)	0.156 (0.003)	0.064 (0.004)	0.016 (0.002)	0.135 (0.003)	0.075 (0.003)	0.021 (0.002)	0.136 (0.001)	0.096 (0.003)	0.033 (0.001)	0.178 (0.005)	0.067 (0.004)	0.018 (0.002)

Table 11: **Held-out donors: cell-type transfer and donor dispersion.**

Method	Wang			Mannens		ROSMAP			Crohn			Alopecia			Rheumatoid arthritis	
	Cell-type macro-F1 ↑	Covariate macro-F1 ↑	Within group donor dispersion ↓	Cell-type macro-F1 ↑	Covariate macro-F1 ↑	Cell-type macro-F1 ↑	Covariate macro-F1 ↑	Within group donor dispersion ↓	Cell-type macro-F1 ↑	Covariate macro-F1 ↑	Within group donor dispersion ↓	Cell-type macro-F1 ↑	Covariate macro-F1 ↑	Within group donor dispersion ↓	Cell-type macro-F1 ↑	Within group donor dispersion ↓
Minuet	0.409 (0.042)	0.083 (0.004)	0.015 (0.001)	0.607 (0.076)	0.198 (0.032)	0.847 (0.035)	0.162 (0.008)	0.020 (0.002)	0.604 (0.024)	**0.371** (0.007)	0.049 (0.004)	0.503 (0.021)	0.153 (0.004)	0.035 (0.002)	0.868 (0.032)	0.016 (0.000)
Apollo	0.279 (0.017)	0.090 (0.004)	0.008 (0.001)	0.605 (0.124)	0.213 (0.025)	0.429 (0.040)	0.164 (0.003)	0.016 (0.001)	0.566 (0.016)	0.364 (0.010)	0.011 (0.000)	0.365 (0.006)	**0.180** (0.004)	0.009 (0.000)	0.822 (0.018)	0.018 (0.000)
scPair	0.289 (0.028)	0.137 (0.010)	0.014 (0.001)	0.582 (0.043)	0.237 (0.034)	0.692 (0.033)	0.141 (0.011)	0.005 (0.001)	0.581 (0.013)	0.357 (0.006)	0.086 (0.011)	0.393 (0.033)	0.155 (0.007)	0.058 (0.010)	0.785 (0.023)	0.158 (0.039)
scPairing	0.341 (0.017)	0.137 (0.016)	0.014 (0.002)	0.543 (0.096)	0.151 (0.019)	0.714 (0.012)	0.137 (0.016)	0.018 (0.005)	0.601 (0.034)	0.370 (0.005)	0.059 (0.003)	0.372 (0.014)	0.153 (0.000)	0.084 (0.005)	0.819 (0.002)	0.121 (0.020)
scCLIP	0.179 (0.021)	0.132 (0.006)	0.015 (0.001)	0.508 (0.021)	0.182 (0.003)	0.273 (0.009)	0.166 (0.003)	0.034 (0.001)	0.291 (0.007)	0.370 (0.007)	0.026 (0.003)	0.083 (0.015)	0.171 (0.008)	0.022 (0.000)	0.383 (0.048)	0.040 (0.008)
IndiSeek	**0.472** (0.006)	**0.191** (0.037)	0.018 (0.001)	**0.652** (0.114)	**0.269** (0.021)	**0.919** (0.014)	0.166 (0.006)	0.016 (0.003)	**0.829** (0.030)	0.358 (0.004)	0.025 (0.002)	**0.639** (0.031)	0.174 (0.003)	0.038 (0.003)	**0.911** (0.032)	0.081 (0.004)
MultiVI	0.239 (0.012)	0.082 (0.014)	0.014 (0.002)	0.508 (0.113)	0.117 (0.007)	0.421 (0.028)	**0.187** (0.008)	**0.001** (0.000)	0.455 (0.013)	0.345 (0.004)	0.063 (0.009)	0.291 (0.023)	0.167 (0.013)	0.059 (0.005)	0.740 (0.038)	0.096 (0.013)
DC-PCA	0.363 (0.004)	0.058 (0.006)	**0.006** (0.001)	0.492 (0.042)	0.220 (0.012)	0.835 (0.026)	0.158 (0.015)	**0.001** (0.000)	0.622 (0.025)	0.338 (0.011)	**0.001** (0.000)	0.351 (0.018)	0.165 (0.024)	**0.002** (0.001)	0.695 (0.012)	**0.000** (0.000)
scGLUE	0.350 (0.025)	0.181 (0.021)	0.024 (0.001)	0.589 (0.141)	0.227 (0.040)	0.844 (0.034)	0.146 (0.003)	0.036 (0.009)	0.781 (0.018)	0.365 (0.003)	0.043 (0.008)	0.507 (0.047)	0.161 (0.001)	0.058 (0.006)	0.856 (0.025)	0.058 (0.002)

Covariate classification is omitted for rheumatoid arthritis because only one tissue class is recorded.
Within-group donor dispersion is omitted for Mannens because its query split does not support within-group donor comparisons.

Table 12: **Held-out donors: donor prediction on query donors.**

Method	Wang		Mannens	ROSMAP		Crohn		Alopecia		Rheumatoid arthritis	
	Donor macro-F1 ↓	Within-group donor macro-F1 ↓	Donor macro-F1 ↓	Donor macro-F1 ↓	Within-group donor macro-F1 ↓	Donor macro-F1 ↓	Within-group donor macro-F1 ↓	Donor macro-F1 ↓	Within-group donor macro-F1 ↓	Donor macro-F1 ↓	Within-group donor macro-F1 ↓
Minuet	0.729	0.755	0.917	0.650	0.911	0.571	0.720	0.522	0.819	0.844	0.769
Apollo	0.646	**0.627**	0.786	0.414	0.651	0.351	0.559	**0.324**	**0.658**	0.838	0.776
scPair	0.697	0.697	0.917	0.573	0.786	0.648	0.780	0.511	0.826	0.839	0.802
scPairing	0.773	0.758	0.922	0.628	0.854	0.587	0.738	0.640	0.861	0.864	0.849
scCLIP	0.665	0.710	0.865	0.508	0.822	0.452	0.611	0.442	0.790	0.771	0.807
IndiSeek	0.901	0.888	0.969	0.779	0.880	0.509	0.756	0.617	0.833	0.818	0.802
MultiVI	0.667	0.663	0.781	**0.353**	0.582	**0.297**	**0.552**	0.340	0.706	0.828	0.796
DC-PCA	**0.320**	**0.627**	**0.489**	0.357	**0.567**	0.468	0.638	0.416	0.797	**0.697**	**0.755**
scGLUE	0.866	0.867	0.937	0.794	0.911	0.610	0.779	0.623	0.875	0.859	0.833

Each value uses the seed-0 representation, refitted with the same configuration and seed where the original was not stored.
Within-group donor prediction is omitted for Mannens because its query split does not support within-group comparisons.

Table 13: **LODO transfer: cell-type transfer and donor dispersion.**

Method	Wang held out		Mannens held out		ROSMAP held out	
	Cell-type macro-F1 ↑	Within group donor dispersion ↓	Cell-type macro-F1 ↑	Within group donor dispersion ↓	Cell-type macro-F1 ↑	Within group donor dispersion ↓
Minuet	0.562 (0.023)	0.116 (0.005)	0.284 (0.040)	0.069 (0.003)	0.696 (0.021)	0.072 (0.002)
Apollo	0.502 (0.029)	0.090 (0.006)	0.221 (0.017)	0.025 (0.002)	0.571 (0.032)	0.054 (0.001)
scPair	0.531 (0.024)	0.163 (0.012)	0.281 (0.033)	0.056 (0.006)	0.657 (0.023)	0.124 (0.006)
scPairing	**0.591** (0.036)	0.236 (0.006)	0.241 (0.043)	0.083 (0.006)	0.689 (0.042)	0.089 (0.015)
scCLIP	0.372 (0.025)	0.196 (0.006)	0.286 (0.019)	0.080 (0.003)	0.585 (0.016)	0.080 (0.002)
IndiSeek	0.551 (0.056)	0.131 (0.019)	**0.334** (0.040)	0.078 (0.004)	**0.735** (0.054)	0.081 (0.005)
MultiVI	0.529 (0.030)	0.170 (0.038)	0.287 (0.011)	0.047 (0.003)	0.632 (0.064)	0.215 (0.023)
DC-PCA	0.374 (0.011)	**0.002** (0.000)	0.277 (0.021)	**0.001** (0.000)	0.471 (0.021)	**0.002** (0.000)
scGLUE	0.544 (0.022)	0.144 (0.019)	0.287 (0.014)	0.063 (0.008)	0.718 (0.032)	0.104 (0.014)

Table 14: **LODO transfer: donor prediction on query donors.**

Method	Wang held out		Mannens held out		ROSMAP held out	
	Donor macro-F1 ↓	Within-group donor macro-F1 ↓	Donor macro-F1 ↓	Within-group donor macro-F1 ↓	Donor macro-F1 ↓	Within-group donor macro-F1 ↓
Minuet	0.391	0.389	0.557	0.557	0.311	0.312
Apollo	0.282	0.275	0.428	0.439	0.262	0.266
scPair	0.336	0.352	0.423	0.379	0.173	0.177
scPairing	0.361	0.359	0.538	0.542	0.246	0.240
scCLIP	0.295	0.301	0.488	0.468	0.251	0.255
IndiSeek	0.557	0.552	0.680	0.692	0.477	0.480
MultiVI	0.254	0.258	0.414	0.431	0.207	0.212
DC-PCA	**0.088**	**0.094**	**0.163**	**0.155**	**0.078**	**0.071**
scGLUE	0.431	0.431	0.604	0.585	0.398	0.405

Each value uses the seed-0 representation, refitted with the same configuration and seed where the original was not stored.

Table 15: **LODO transfer: paired-cell matching.**

Method	Wang held out			Mannens held out			ROSMAP held out		
	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑	FOSCTTM ↓	MRR ↑	Recall@1 ↑
Minuet	**0.236** (0.001)	**0.056** (0.000)	**0.016** (0.000)	**0.270** (0.008)	**0.045** (0.002)	**0.012** (0.001)	**0.195** (0.003)	**0.058** (0.002)	**0.016** (0.001)
Apollo	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
scPair	0.323 (0.001)	0.036 (0.001)	0.009 (0.000)	0.407 (0.009)	0.028 (0.003)	0.007 (0.001)	0.250 (0.007)	0.042 (0.002)	0.010 (0.000)
scPairing	0.303 (0.004)	0.035 (0.001)	0.009 (0.001)	0.292 (0.003)	0.033 (0.002)	0.008 (0.002)	0.228 (0.003)	0.042 (0.001)	0.010 (0.000)
scCLIP	0.345 (0.005)	0.031 (0.000)	0.007 (0.000)	0.296 (0.007)	0.037 (0.002)	0.009 (0.001)	0.238 (0.008)	0.043 (0.003)	0.010 (0.002)
IndiSeek	0.449 (0.006)	0.015 (0.000)	0.002 (0.000)	0.401 (0.008)	0.020 (0.001)	0.003 (0.000)	0.358 (0.023)	0.023 (0.003)	0.004 (0.001)
MultiVI	0.283 (0.003)	0.037 (0.001)	0.009 (0.000)	0.327 (0.011)	0.031 (0.001)	0.007 (0.001)	0.226 (0.004)	0.043 (0.001)	0.010 (0.001)
DC-PCA	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
scGLUE	0.383 (0.102)	0.027 (0.010)	0.006 (0.003)	0.317 (0.034)	0.035 (0.005)	0.009 (0.001)	0.248 (0.026)	0.044 (0.004)	0.011 (0.002)
`Pooled`	0.259 (0.002)	0.045 (0.000)	0.012 (0.001)	0.252 (0.007)	0.045 (0.001)	0.012 (0.000)	0.212 (0.007)	0.052 (0.002)	0.013 (0.001)
`Within-donor`	0.260 (0.002)	0.040 (0.002)	0.009 (0.001)	0.288 (0.013)	0.032 (0.003)	0.006 (0.001)	0.223 (0.003)	0.045 (0.001)	0.010 (0.001)

Table 16: **Phenotype prediction on held-out donors.** Retrieval@1: participant retrieval at rank one.

Method	ROSMAP			Crohn			Alopecia		
	Macro AUROC ↑	Macro-F1 ↑	Retrieval @1 ↑	Macro AUROC ↑	Macro-F1 ↑	Retrieval @1 ↑	Macro AUROC ↑	Macro-F1 ↑	Retrieval @1 ↑
Minuet	0.693	**0.708**	0.649	**0.627**	0.683	**0.711**	0.548	0.460	0.378
Apollo	0.705	0.296	0.439	0.567	0.562	0.607	0.511	0.159	0.354
scPair	0.648	0.591	0.632	0.579	0.562	0.607	0.410	0.325	0.292
scPairing	0.420	0.578	0.561	0.620	**0.701**	0.595	0.535	0.378	0.250
scCLIP	0.670	0.521	0.579	0.614	0.591	0.702	0.480	0.335	0.292
IndiSeek	0.636	0.521	0.614	0.579	0.620	0.631	**0.559**	**0.468**	0.312
MultiVI	0.273	0.284	0.386	0.509	0.581	0.667	0.387	0.423	0.250
DC-PCA	**0.727**	0.676	**0.667**	0.550	0.520	0.512	0.498	0.356	**0.417**
scGLUE	0.670	0.568	0.649	0.480	0.591	0.655	0.356	0.244	0.312

G Representation Visualizations

Each cohort block shows UMAPs of the joint embeddings, with one panel per method, colored by cell type in the upper row and by donor in the lower row. All blocks use seed 0 and up to 20,000 donor-balanced cells, with the same cells, colors, and method order within a cohort; cell-type colors are fixed across protocols, and donor colors are arbitrary. For the held-out-donor and LODO protocols, UMAP is fitted on reference embeddings and query cells are projected without refitting. Phenotype-prediction panels instead show a donor-balanced sample of out-of-fold embeddings, because the five folds each have their own reference set; each fold's model has its own coordinates, so the folds of one method can occupy separate regions.

G.1 Observed-cell integration

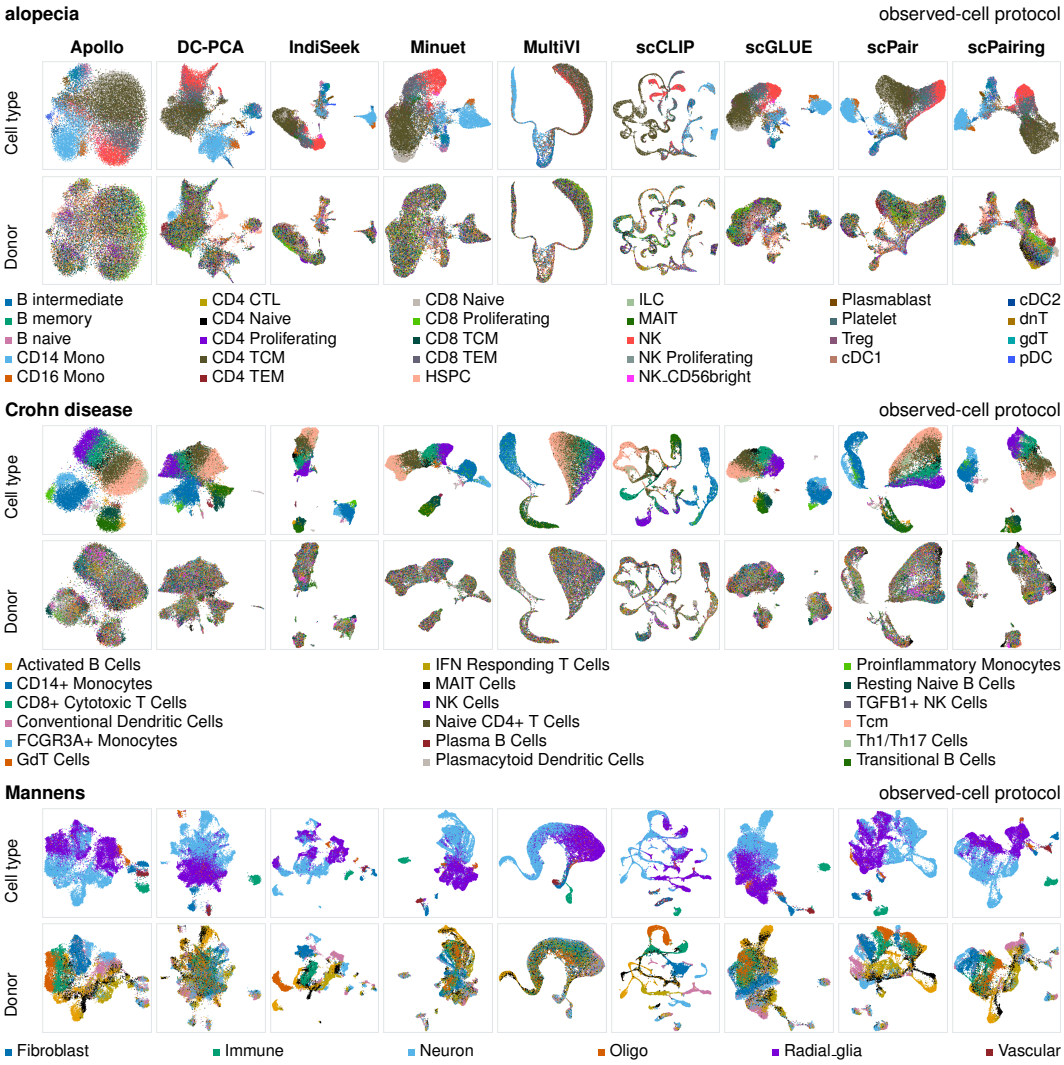

Figure 17: **Observed-cell embeddings.** Each cohort shows cell-type colors above donor colors in the same method order.

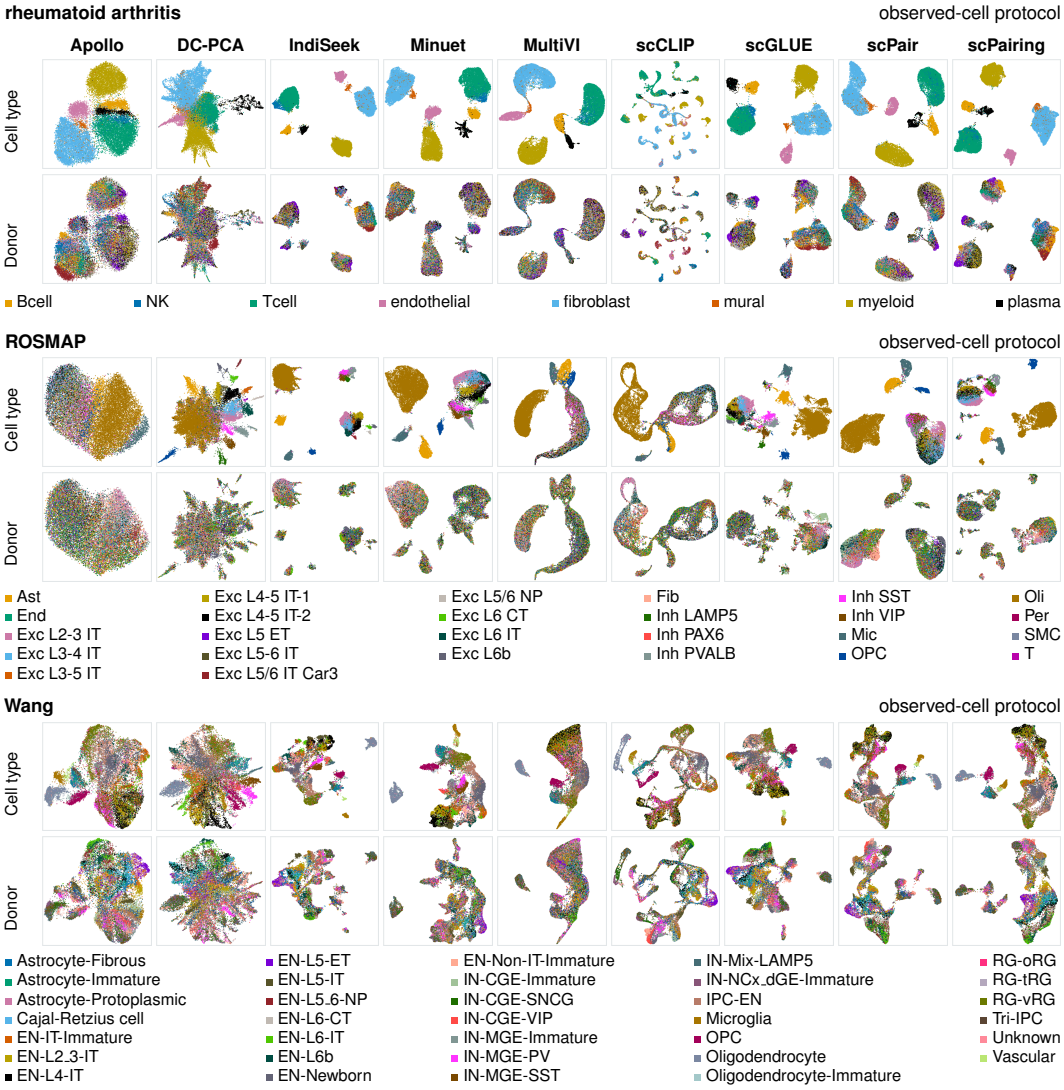

Figure 18: **Observed-cell embeddings.** Each cohort shows cell-type colors above donor colors in the same method order.

G.2 HELD-OUT-DONOR TRANSFER

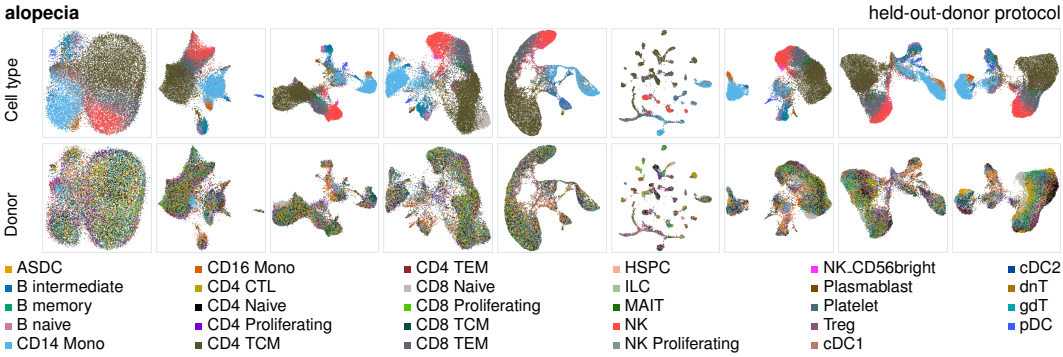

Figure 19: **Held-out-donor embeddings.** Each cohort shows cell-type colors above donor colors in the same method order.

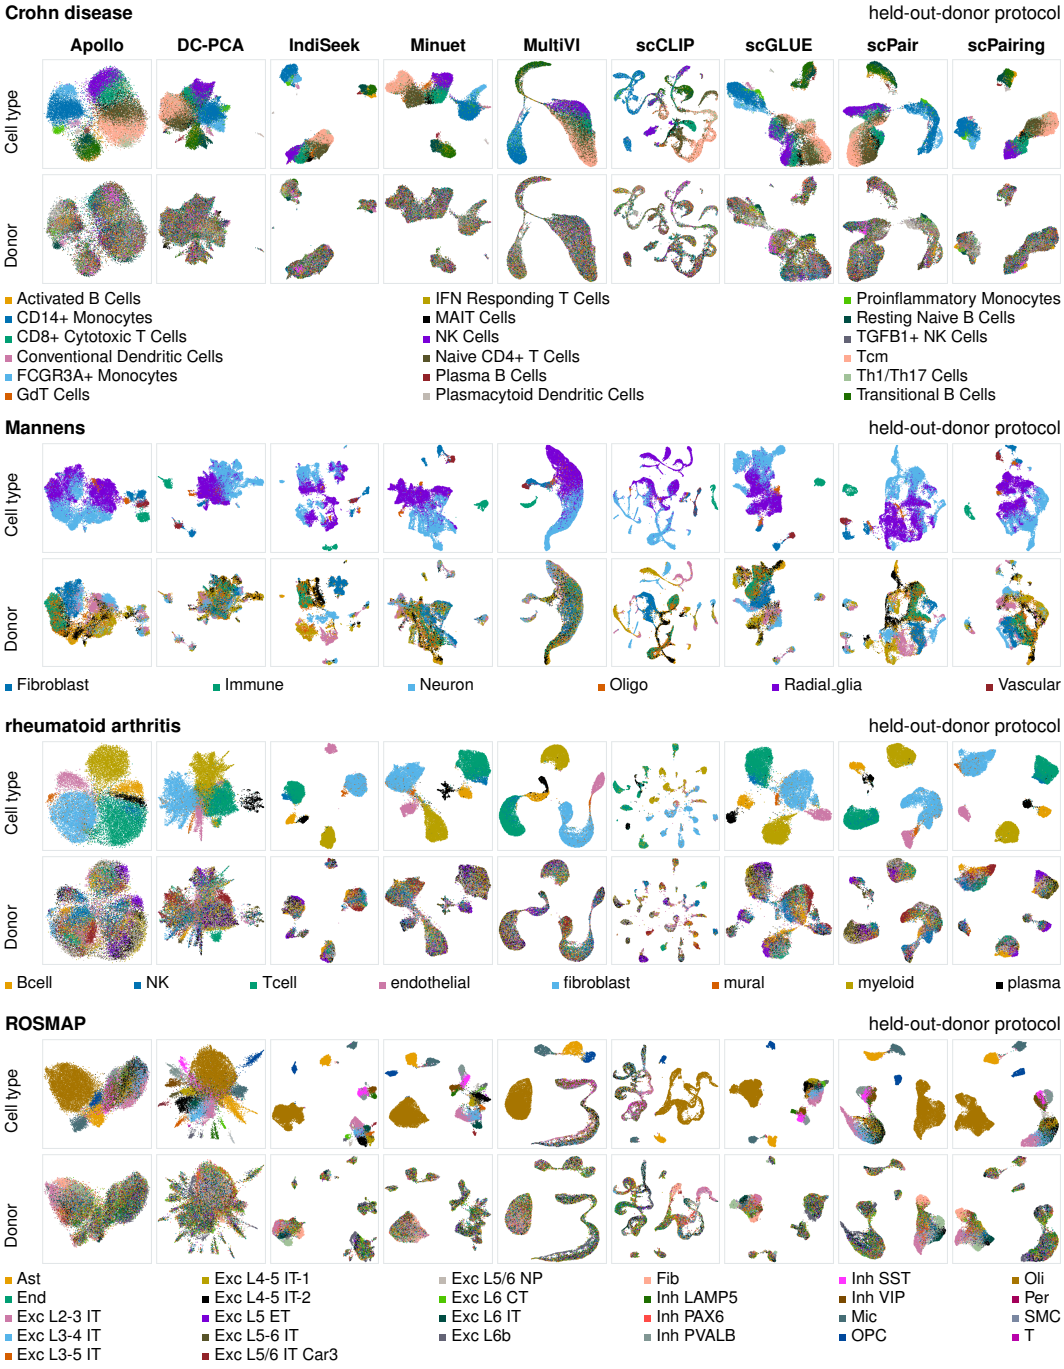

Figure 20: **Held-out-donor embeddings.** Each cohort shows cell-type colors above donor colors in the same method order.

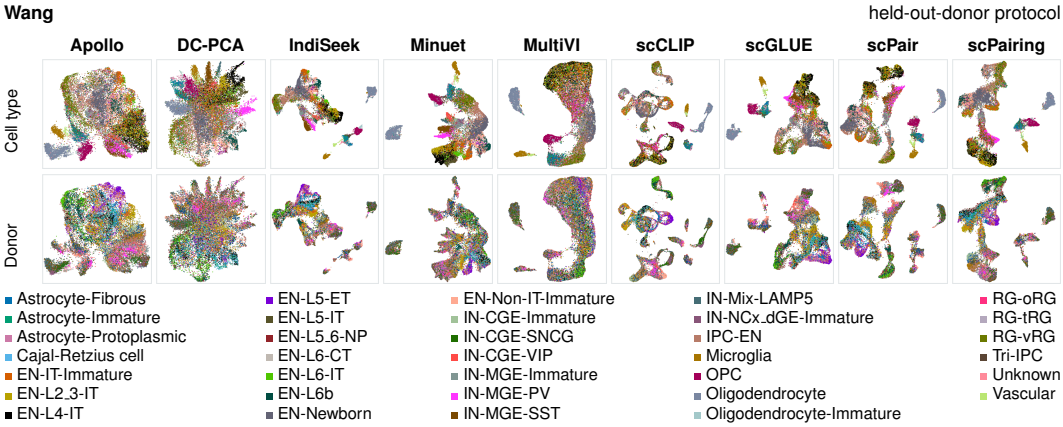

Figure 21: **Held-out-donor embeddings.** Each cohort shows cell-type colors above donor colors in the same method order.

G.3 Leave-one-dataset-out transfer

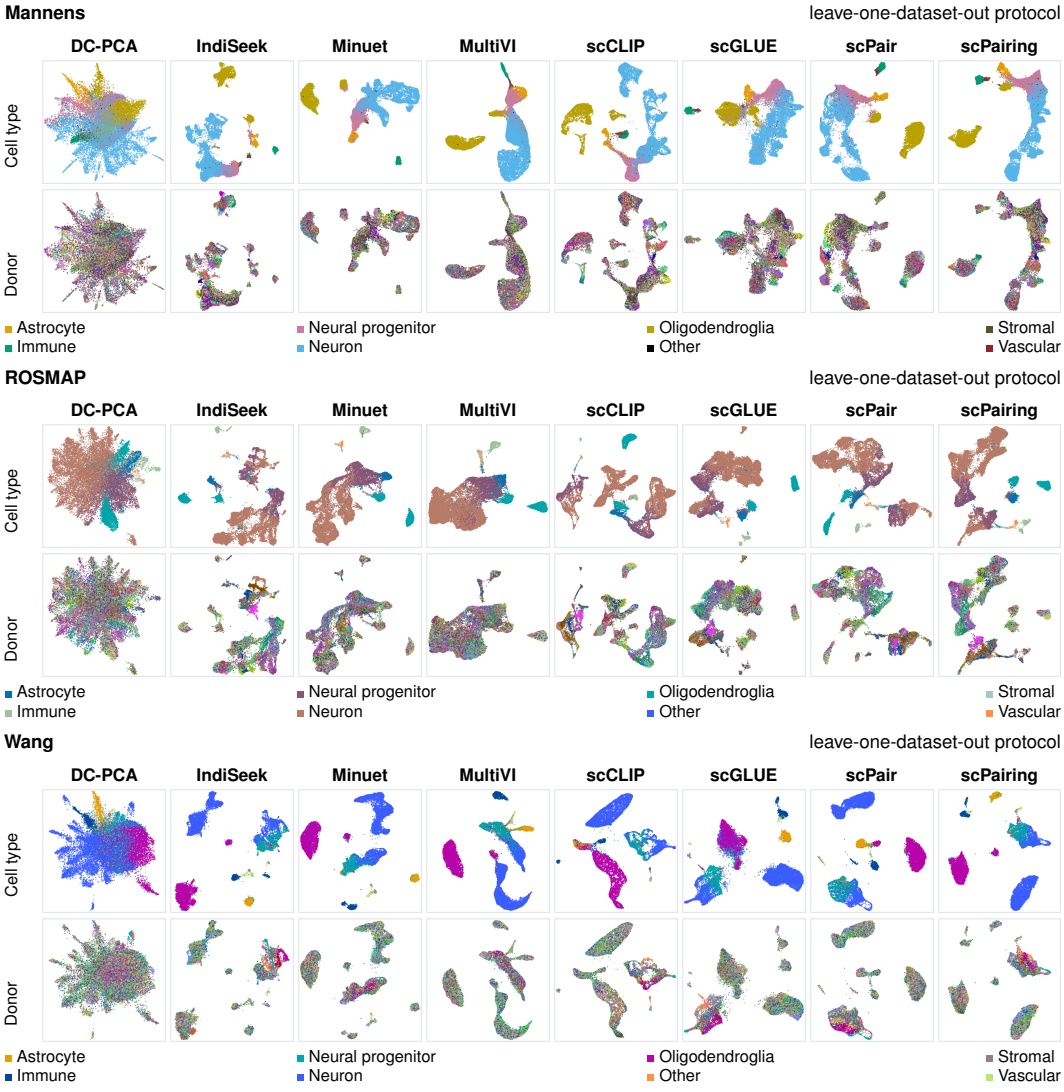

Figure 22: **Leave-one-dataset-out embeddings.** Each cohort shows cell-type colors above donor colors in the same method order.

G.4 PHENOTYPE PREDICTION

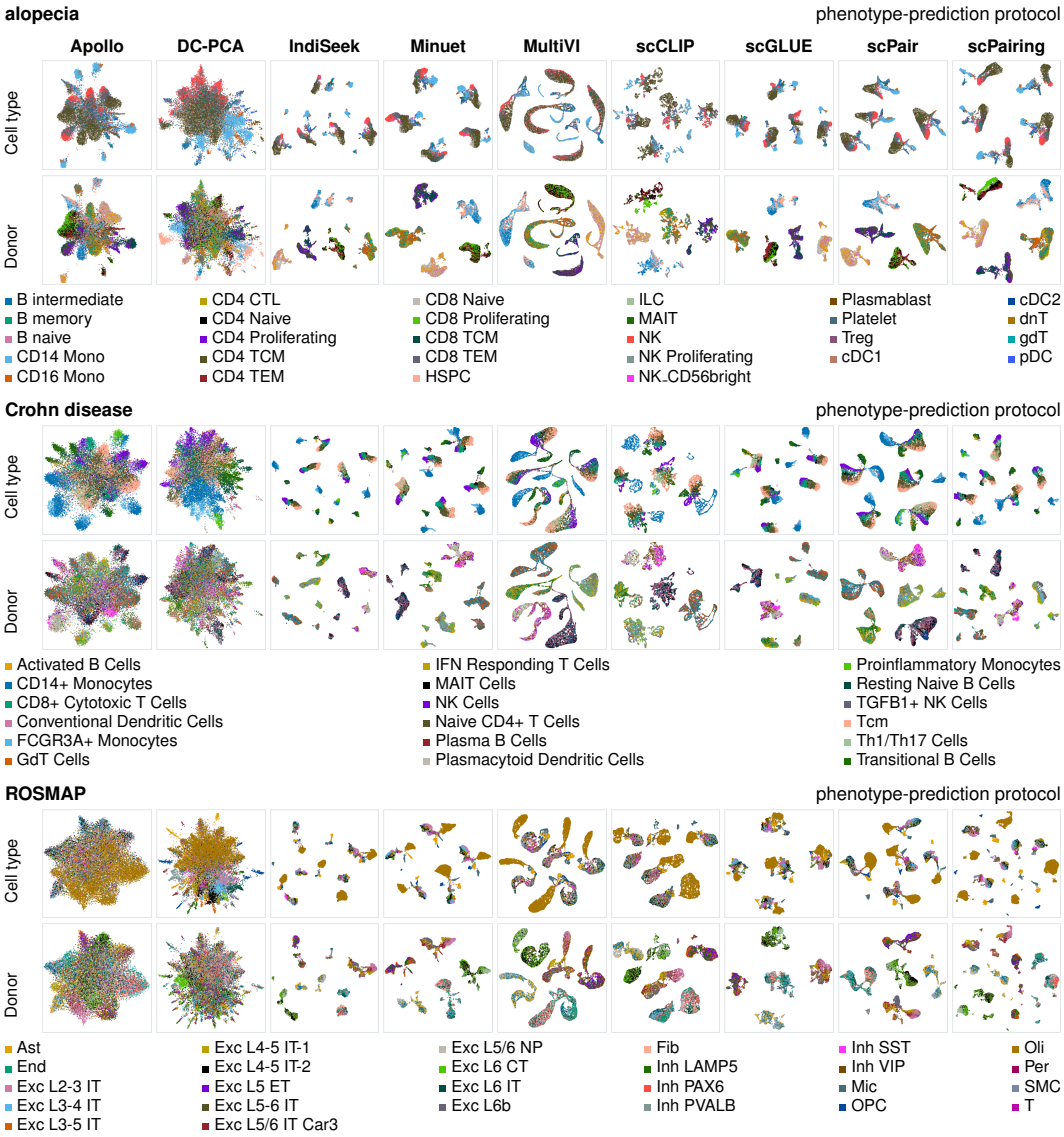

Figure 23: **Phenotype-prediction embeddings.** Each cohort shows cell-type colors above donor colors in the same method order.