

000 GROUP CONTRASTIVE LEARNING FOR WEAKLY 001 PAIRED MULTIMODAL DATA 002 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009 010

011 We present GROOVE, a semi-supervised multi-modal representation learning
012 approach for high-content perturbation data where samples across modalities are
013 weakly paired through shared perturbation labels but lack direct correspondence.
014 Our primary contribution is GroupCLIP, a novel group-level contrastive loss that
015 bridges the gap between CLIP for paired cross-modal data and SupCon for uni-
016 modal supervised contrastive learning, addressing a fundamental gap in contrastive
017 learning for weakly-paired settings. We integrate GroupCLIP with an on-the-fly
018 backtranslating autoencoder framework to encourage cross-modally entangled
019 representations while maintaining group-level coherence within a shared latent
020 space. Critically, we introduce a comprehensive combinatorial evaluation frame-
021 work that systematically assesses representation learners across multiple optimal
022 transport aligners, addressing key limitations in existing evaluation strategies.
023 This framework includes novel simulations that systematically vary shared versus
024 modality-specific perturbation effects enabling principled assessment of method
025 robustness. Our combinatorial benchmarking reveals that there is not yet an aligner
026 that uniformly dominates across settings or modality pairs. Across simulations
027 and two real single-cell genetic perturbation datasets, GROOVE performs on par
028 with or outperforms existing approaches for downstream cross-modal matching
029 and imputation tasks. Our ablation studies demonstrate that GroupCLIP is the key
030 component driving performance gains. These results highlight the importance of
031 leveraging group-level constraints for effective multi-modal representation learning
032 in scenarios where only weak pairing is available.

033 1 INTRODUCTION

034 Perturbation screens have gained major prominence in recent years for their ability to elucidate causal
035 gene regulatory networks (Dixit et al., 2016), identify candidate therapeutic targets (Rood et al., 2024),
036 and enable small molecule repurposing (Bhandari et al., 2022). Each given modality (e.g. RNA-Seq,
037 ATAC-Seq, or high-content imaging) only observes a subset of the underlying biology of a system (Cui
038 et al., 2025), therefore recent efforts have shifted toward multi-modal investigation of perturbation
039 effects via paired profiling approaches (Frangieh et al., 2021; Martin-Rufino et al., 2025). While
040 promising, this type of profiling remains feasible only for specific combinations of modalities, such as
041 gene expression paired with chromatin accessibility (Martin-Rufino et al., 2025) or gene expression
042 paired with surface protein measurements (Frangieh et al., 2021). Notably, it is not currently feasible
043 to obtain both perturbed microscopy images from cell painting assays (Feldman et al., 2019) and
044 perturbed gene expression profiles from the *same* individual cells, as both measurements are inherently
045 destructive assays. In this setting, we do not have access to paired samples across modalities and can
046 only broadly *group* cells by their perturbation (other experimental) *labels*.

047 Consequently, recent efforts in multi-modal perturbation screens have shifted toward developing
048 computational approaches for post-hoc “pairing” (even though true pairs don’t actually exist) of
049 cells across modalities or cross-modal imputation (See Section 2). Both these objectives depend on
050 learning a useful joint representation of the non-paired multi-modal data with group-level information
051 only. Such a setting immediately rules out existing standard contrastive learning approaches for joint
052 cross-modal inference (see Section 2). Cross-modal contrastive approaches like CLIP (Radford et al.,
053 2021) need paired data while uni-modal label-based contrastive methods like SupCon (Khosla et al.,
2020) are not natively compatible with multi-modal data. Moreover, existing multi-modal single-cell

054 deep learning approaches, such as uncoupled autoencoders (Samaran et al., 2024; Ashuach et al.,
 055 2023), rely on either strong human-defined priors to establish putative cell correspondences or access
 056 to paired data. No current framework can effectively learn from *native* weakly paired multi-modal
 057 data for a *well-mixed* multi-modal latent representation (See Section 2).

058 **Contributions.** In this work, we develop GROOVE (GROUp cOntrastive learning for weakly paired
 059 multi-modal data) to address these challenges in weakly paired multi-modal data. Our approach
 060 makes three key contributions. First, we introduce a group-level semi-supervised contrastive loss,
 061 bridging the gap between CLIP (Radford et al., 2021) and SupCon (Khosla et al., 2020) for weakly
 062 paired multi-modal data. Second, GROOVE integrates this loss with *on-the-fly* backtranslating
 063 autoencoders adapted from neural machine translation (Artetxe et al., 2017), creating a unified
 064 architecture for learning from weakly paired single-cell data. Finally, we develop a comprehensive
 065 evaluation framework consisting of: (i) novel simulations that systematically vary the proportion of
 066 shared versus modality-specific information, and (ii) combinatorial benchmarking that pairs different
 067 representation learners with various alignment algorithms to assess both matching and cross-modal
 068 imputation performance.

069 **Notations.** In standard constructions of multi-modal learning, one works with a setting where all
 070 the data modalities are observed for all the samples such that $\mathcal{D} = \{(\mathbf{x}_i^{(1)}, \mathbf{x}_i^{(2)})\}_{i=1}^N$. However, the
 071 focus of this work is on settings where such paired data does not exist, i.e., we only have access to
 072 one data modality per sample. We instead have access to a common state, environment, intervention
 073 or *perturbation* label $t \in \mathcal{T} \subset \mathbb{Z}$ that is shared across modalities and samples. This additional
 074 information renders our data as *weakly paired* such that any data instance across the modalities without
 075 the same label t are strictly unrelated. We can now re-formulate multi-modal learning in the *weakly*
 076 *paired* setting as having data $\mathcal{D}^{(m)}$ from two disjoint data modalities indexed by $m \in \mathbb{M} = \{1, 2\}$.
 077 Each modality-specific dataset is a collection of $N^{(m)}$ samples $\mathcal{D}^{(m)} = \{(\mathbf{x}_i^{(m)}, t_i)\}_{i=1}^{N^{(m)}}$, where
 078 each $\mathbf{x}_i^{(m)} \in \mathcal{X}^{(m)} \subseteq \mathbb{R}^{k^{(m)}}$ is the data instance for modality m and its corresponding label t_i . Given
 079 this, our multimodal representation learning problem is to learning an embedding $\mathbf{z} \in \mathcal{Z} \subseteq \mathbb{R}^d$ for
 080 each sample in a *shared* low-dimensional representation space, such that $d \ll \min(k^{(1)}, k^{(2)})$. And
 081 let $\mathcal{D}_z^{(m)} = \{(\mathbf{z}_i^{(m)}, t_i)\}_{i=1}^{N_m}$ represent the collection of latent representations and labels for modality
 082 m . We define \bar{m} to denote the *other* modality.

085 2 BACKGROUND

086 In this section, we review related work from both the broader contrastive representation learning
 087 literature and the single-cell community, highlighting the gap our contributions aim to address. We
 088 also briefly review unsupervised neural translation which underpins our architecture.

089 **Contrastive representation learning.** Contrastive representation learning has emerged as a pow-
 090 erful general paradigm (Le-Khac et al., 2020), with theoretical underpinnings (Alshammari et al.,
 091 2025; Van Assel et al., 2025; HaoChen et al., 2021) and a wide range of practical instantiations. The
 092 foundational InfoNCE loss (Oord et al., 2018) maximizes mutual information between positives
 093 while minimizing it for negatives, and has since inspired numerous influential extensions (Chen et al.,
 094 2020; He et al., 2020; Grill et al., 2020). Despite their success, these methods suffer from a key
 095 limitation: they often misclassify samples from the same class as negatives. Supervised contrastive
 096 learning (Khosla et al., 2020) alleviates this by leveraging labels to form multiple positives per
 097 anchor, but remains uni-modal in design. Similarly, Yao et al. (2024) leverage perturbation labels
 098 to pair samples across batches, but operates at the group level as a uni-modal denoising approach.
 099 For multimodal learning, CLIP (Radford et al., 2021) and CMC (Tian et al., 2020) are the most
 100 influential approaches, aligning image–text pairs or extending contrastive objectives across multiple
 101 modalities. **Other works, S-CLIP (Mo et al., 2023) and SemiCLIP (Gan et al., 2025), have explored**
 102 **supervised and semi-supervised extensions of CLIP that leverage *limited paired data* with additional**
 103 **label information.** However, these methods still require some instance-level pairs between modalities
 104 **and cannot operate in the purely weakly-paired regime where only group labels connect modalities.**

105 Their reliance on strictly paired data remains a critical bottleneck. Thus, a key gap in the literature is
 106 the absence of a supervised extension of CLIP that can exploit weak pairing (Figure 1a).

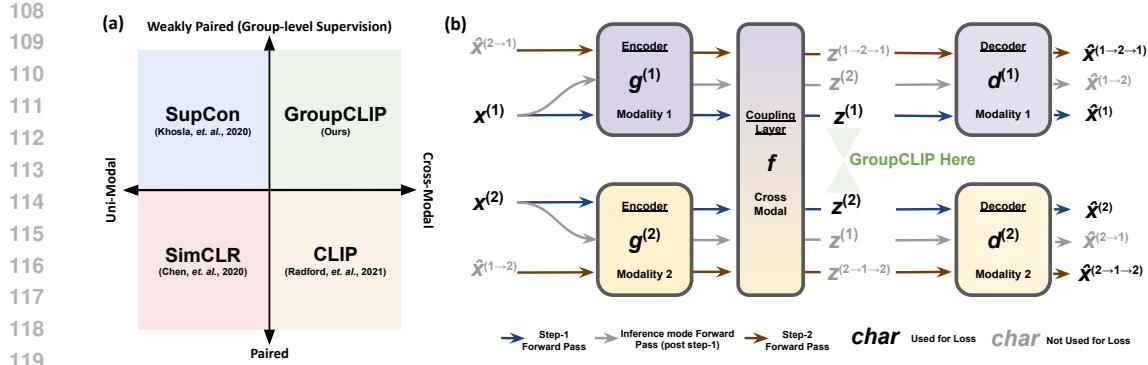


Figure 1: (a) GroupCLIP in the context for broader contrastive learning can be viewed as the multi-modal generalization of SupCon. (b) GROOVE Architecture and training step illustration. Each iteration consists of two steps: (1) optimize reconstruction and GroupCLIP losses, then (2) generate cross-modal pseudo-samples in inference mode and optimize the backtranslation loss.

Weakly paired learning for multimodal single-cell data. Single-cell data is inherently unpaired because measurements are destructive, preventing multiple modalities from being captured from the same cell. This limitation has motivated the development of computational methods for unpaired multimodal integration (Yang et al., 2021; Samaran et al., 2024) and weakly paired learning (Xi et al., 2024; Ryu et al., 2025), particularly in settings involving perturbations, which are the focus of our work. A particularly important downstream application of these methods is cross-modal imputation, where one modality is predicted from observations of another. Next, we briefly describe the most relevant works tackling this problem. First, Xi et al. (2024) leverage perturbation labels via propensity score matching. Their method trains independent classifiers to predict perturbations in each modality and then uses the resulting logits to define a common support for alignment, following the classical balancing scores of Rubin (1974). Although conceptually appealing, this approach has notable limitations: it assumes that all perturbation-induced variation is perfectly shared across modalities, ignoring modality-specific effects; and the learned latent representations (classifier logits) capture only perturbation-predictive information, discarding potentially valuable *intra*-perturbation variation that may be critical for downstream tasks. Second, Ryu et al. (2025) address the alignment stage of the pipeline by introducing a label-constrained variant of the Gromov–Wasserstein optimal transport (GW-OT) problem (Mémoli, 2011). Unlike approaches that rely on direct sample-wise correspondences, GW-OT aligns representations by minimizing structural discrepancies between the metric spaces induced by the two modalities (Sebbouh et al., 2024; Van Assel et al., 2024). This formulation allows alignment across modality-specific latent spaces of different dimensions, each obtained via PCA. The method’s effectiveness, however, is constrained by (i) the quality and robustness of the latent representations, (ii) the cubic computational complexity of GW-OT (Peyré et al., 2016), and (iii) its non-convexity, which makes it susceptible to local minima (Vayer, 2020). Finally, Samaran et al. (2024) propose the only method that integrates alignment and imputation within a unified framework. Yet, their approach requires a predefined set of aligned features shared across modalities—a strong assumption that is rarely satisfied in practice—and it does not leverage perturbation labels that could substantially enhance alignment.

Multi-lingual Neural Translation. Neural machine translation has demonstrated that high-quality translation systems can be trained using only unpaired monolingual data through unsupervised approaches that employ shared encoder architectures, dual training objectives combining reconstruction and back-translation losses, and cross-lingual initialization strategies (Sennrich et al., 2015; Artetxe et al., 2017; Lample et al., 2018; Conneau et al., 2018). Central to these methods is an *on-the-fly* backtranslation strategy designed to encourage entangled, cross-modally¹ informative latent representations. During training, a sample from one modality is encoded to the shared latent space, then this representation is used to generate a corresponding translation in the *other* modality. Parameters are updated via a two step optimization procedure based on a *reconstruction* (Equation 1) and *backtranslation* (Equation 2) loss. The reconstruction ($\hat{x}_i^{(m)}$) loss computes the error in encoding and decoding $x_i^{(m)}$ from the same modality and the backtranslation loss computes the error when

¹We use modal and lingual interchangeably here.

162 encoding $\mathbf{x}_i^{(m)}$ and decoding to the other modality ($\hat{\mathbf{x}}_i^{(m \rightarrow \bar{m})}$), followed by encoding and decoding
 163 $\hat{\mathbf{x}}_i^{(m \rightarrow \bar{m})}$ back to its original modality. This backtranslation strategy encourages the latents to be
 164 well-mixed across modalities. And as the model improves higher-quality pseudo-pairs are generated,
 165 creating a positive feedback loop that enhances both cross-modal alignment and within-modality
 166 representations. See Appendix A.1 for extended details on this procedure and the architecture. To the
 167 best of our knowledge, iterative back-translation from multilingual neural machine translation has
 168 not yet been applied to weakly paired single-cell multi-modal data.
 169

$$\mathcal{L}_{\text{reconstruction}} = \frac{1}{|\mathbb{M}|} \sum_{m \in \mathbb{M}} \frac{1}{|\mathcal{D}^{(m)}|} \sum_{\mathbf{x}_i^{(m)} \in \mathcal{D}^{(m)}} \left\| \hat{\mathbf{x}}^{(m)} - \mathbf{x}_i^{(m)} \right\|_2^2 \quad (1)$$

$$\mathcal{L}_{\text{backtranslation}} = \frac{1}{|\mathbb{M}|} \sum_{m \in \mathbb{M}} \frac{1}{|\mathcal{D}^{(m)}|} \sum_{\mathbf{x}_i^{(m)} \in \mathcal{D}^{(m)}} \left\| \hat{\mathbf{x}}^{(m \rightarrow \bar{m} \rightarrow m)} - \mathbf{x}_i^{(m)} \right\|_2^2 \quad (2)$$

3 PROPOSED METHOD

180 We instantiate our base architecture using uncoupled autoencoders, which are standard in the single-
 181 cell literature (Samaran et al., 2024; Lopez et al., 2018), with an added shared linear (coupling)
 182 projection layer across modalities. We train using the two-step optimization procedure from unsu-
 183 pervised machine translation (Appendix A for more details). This base model, however, does not
 184 leverage the additional supervisory signal present in the sample associated (perturbation) labels \mathcal{T} .
 185 We can make use of this additional information to increase the level of supervision of our *on-the-fly*
 186 autoencoder from unsupervised to semi-supervised. The label information allows us to update our
 187 latent representation with the following desiderata: (1) sample across modalities with the same label
 188 should be pushed close together and (2) any samples without the same label should be as far away
 189 as possible (repulsed). Such a formulation invokes contrastive learning as a natural solution. This
 190 motivates us to develop a novel contrastive loss for weakly paired multi-modal data. Our approach
 191 leverages the weak pairing structure by treating samples from different modalities that share the same
 192 label as positive pairs. For a latent representation $\mathbf{z}_i^{(m)}$ from modality m , we define its attractors as
 193 all latent representations from the other modality \bar{m} that share the same label.

194 Formally, we define $\mathcal{P}_i^{(m)} = \{ \mathbf{z}_j^{(\bar{m})} \in \mathcal{D}_z^{(\bar{m})} : t_j = t_i \}$ as the collection of *positive* samples for
 195 anchor $\mathbf{z}_i^{(m)}$ (same label, opposite modality) and $\mathcal{A}_i^{(m)} = \mathcal{D}_z^{(\bar{m})}$ as the collection of all *candidates*
 196 from the other modality (includes both positives and negatives). Following CLIP (Radford et al.,
 197 2021), we normalize over *all* candidates from the other modality (including positives). Given the
 198 anchor $\mathbf{z}_i^{(m)}$, the loss is:

$$\ell_i^{(m)} = -\log \frac{\sum_{\mathbf{z}_p \in \mathcal{P}_i^{(m)}} \exp(\text{sim}(\mathbf{z}_i^{(m)}, \mathbf{z}_p) / \tau)}{\sum_{\mathbf{z}_a \in \mathcal{A}_i^{(m)}} \exp(\text{sim}(\mathbf{z}_i^{(m)}, \mathbf{z}_a) / \tau)} \quad (3)$$

204 where $\text{sim}(\mathbf{u}, \mathbf{v})$ is a similarity function between two embeddings (such as cosine similarity) and
 205 $\tau > 0$ is the temperature parameter used to scale similarities. Note that in addition to the standard
 206 cosine similarity, we also experimented with t -distribution based similarity kernels, see Appendix A.4.
 207 Averaging over all anchors and modalities gives:

$$\mathcal{L}_{\text{GroupCLIP}} = \frac{1}{|\mathbb{M}|} \sum_{m \in \mathbb{M}} \frac{1}{|\mathcal{D}_z^{(m)}|} \sum_{\mathbf{z}_i^{(m)} \in \mathcal{D}_z^{(m)}} \ell_i^{(m)} \quad (4)$$

212 The resulting GroupCLIP loss encourages latent representations of the same label to cluster together
 213 across modalities while pushing apart representations of all other labels. Because contrastive learning
 214 is sensitive to batch composition, we employ a balanced under-sampling strategy that maintains equal
 215 sample counts per label in each mini-batch, preventing class imbalance without oversampling rare
 labels, see Appendix A.5.

Our approach fills a crucial gap in the contrastive learning landscape by extending supervised contrastive learning to the cross-modal setting (Figure 1a). Just as CLIP (Radford et al., 2021) represents the canonical cross-modal extension of SimCLR’s (Chen et al., 2020) unsupervised contrastive framework, GroupCLIP serves as the natural cross-modal extension of SupCon’s (Khosla et al., 2020) supervised contrastive approach. While CLIP leverages naturally paired data without explicit label supervision, GroupCLIP leverages label supervision without requiring natural pairings between modalities. This distinction is critical for biological applications where true cross-modal pairs are often unavailable, but perturbation labels provide rich supervisory signal. By bringing supervised contrastive learning to the multi-modal domain, GroupCLIP bridges the gap between uni-modal supervised methods and cross-modal unsupervised approaches, offering a principled framework for scenarios with weak pairing but strong label information.

Finally, by integrating our novel GroupCLIP loss (Equation 4) with the *on-the-fly* backtranslating autoencoder framework (Equation 1,2), we develop the GROOVE architecture – a novel method that learns a unified representations from weakly paired multi-modal data with the following losses:

$$\mathcal{L}_{\text{step-1}} = \alpha \cdot \mathcal{L}_{\text{GroupCLIP}} + \beta \cdot \mathcal{L}_{\text{reconstruction}} \quad (5) \quad \mathcal{L}_{\text{step-2}} = \beta \cdot \mathcal{L}_{\text{backtranslation}} \quad (6)$$

α, β are hyperparameter that balance the loss components. Algorithm 1 in Appendix A.6 sketches the overall training loop of GROOVE and Figure 1b is a visual illustration of this architecture.

4 EVALUATION FRAMEWORK

We evaluate the quality of our learned latent representations through two approaches. First, we assess their utility for OT-based cross-modal sample matching, which is a standard evaluation (Ryu et al., 2025; Xi et al., 2024) since high-quality latent representations should enable the recovery of meaningful transport plans that correctly identify and match similar samples across modalities. We then evaluate performance on a key downstream task leveraging the transport plan: uni-directional imputation, where we predict one modality given samples from another.

4.1 COMBINATORIAL EVALUATION

We identify a key limitation in previous evaluations: they either hold the representation learner fixed while comparing various OT methods (Ryu et al., 2025), or fix the OT approach while comparing different representation learners (Xi et al., 2024). This evaluation paradigm can produce systematically biased results because the performance of representation learning methods is inherently coupled with the choice of downstream alignment algorithm. The coupling arises from fundamental differences in the geometric assumptions underlying various OT formulations.

For instance, Entropic Optimal Transport (EOT) implicitly assumes that both modalities can be embedded into a shared feature space, such that cross-modal correspondences are meaningful when representations are compared using the same metric. In other words, the learned embeddings are expected to preserve distances consistently across modalities in a common coordinate system. By contrast, Entropic Gromov–Wasserstein Optimal Transport (EGWOT) does not assume a shared space; instead, it treats the modalities as potentially distinct metric spaces and aligns them by comparing their internal distance structures rather than absolute coordinates. These differing assumptions suggest that representation learners optimized for one geometric framework may not perform optimally under alternative OT formulations. To address this limitation, we propose a combinatorial evaluation framework that systematically tests each representation learner against all available OT aligners. This approach allows us to assess how representation learners and alignment algorithms interact, yielding more robust performance evaluations that account for the sensitivity of downstream choices.

4.2 BASELINES

We focus our evaluation on approaches that can natively operate on weakly-paired data without requiring pre-specified feature correspondences or paired samples. This includes the following representation learning baselines: Propensity Score (PS) matching (Xi et al., 2024) and domain-adversarial variational autoencoder with label adaptation (DAVAE), a custom modification we made of Ashuach et al. (2023) as described in Ryu et al. (2025). PS represents the most recent approach

270 for weakly paired cross-modal matching and is the only existing method that directly addresses this
 271 problem in its native form. DVAE is conceptually closest to our proposed approach but differs in
 272 employing modality identification adversarial loss and using linear probe supervision for latent space
 273 regularization. All methods use identical architectures with consistent latent embedding dimensions,
 274 and we evaluate each representation learner in conjunction with both standard (EOT, EGWOT) and
 275 label-constrained optimal transport approaches (LabeledEOT, LabeledEGWOT, labeledCOOT). See
 276 Appendix B for details.

277 For unidirectional imputation, we trained a 2-layer MLP interleaved with ReLU activations and a final
 278 linear projection layer. This model was trained using the transport plan returned by the OT aligner
 279 which defines the training sampling strategy. Specifically, when training to predict modality 1 from
 280 modality 2 and given a transport plan $\mathbf{T} \in [0, 1]^{N^{(1)} \times N^{(2)}}$, we sample index- j from the modality 1 for
 281 each sample index- i from modality 2 as function of $j \sim \text{Multinomial}(\frac{\mathbf{T}_{:,i}}{\sum \mathbf{T}_{:,i}})$, for each mini-batch.
 282

283 4.3 DATASETS

285 **Simulations.** A key limitation in previous evaluations is the assumption that variation from perturba-
 286 tions are fully shared between modalities, making modality-specific variation completely independent
 287 of labels, an unrealistic assumption for biological data (Argelaguet et al., 2020; Lin & Zhang, 2023).
 288 To address this, we developed a probabilistic simulation framework that captures both shared and
 289 modality-specific latent variation affected by perturbations. Our framework models each cell through
 290 shared latent factors (affecting both modalities identically) and modality-specific factors (unique to
 291 each molecular layer), with two perturbation types: shared perturbations affecting both modalities
 292 jointly, and modality-specific perturbations targeting individual modalities independently. We sys-
 293 tematically evaluate method performance across different coupling levels by varying the proportion
 294 of shared versus modality-specific dimensions, testing scenarios with 100%, 80%, and 50% shared
 295 variation. Full simulation details are provided in Appendix C. For each scenario, we simulate 10
 296 replicates with perturbation-balanced 80-20 train-test splits for evaluation.
 297

298 **Perturb-Multiome** We use the paired gene expression and chromatin accessibility (multiome) dataset
 299 from Martin-Rufino et al. (2025), consisting of transcription factor based perturbations for a more
 300 realistic evaluation. The original data was pre-processed and subset to result in 2560 cells over 20
 301 perturbation with 128 cells per perturbation with 512 genes and 256 gene accessibility scores. See
 302 details in Appendix D. This dataset was analyzed under two frameworks: (1) perturbation balanced
 303 5-fold splits, and (2) leave-one-perturbation-out (LOPO). All imputation results for this dataset focus
 304 on predicting gene expression from accessibility scores.
 305

306 **Perturb-CITE-seq.** We also use the Frangieh et al. (2021) Pertub-CITE-seq data, consisting of
 307 genetic perturbation with paired RNA-seq and surface protein measurements, for a more realistic
 308 evaluation. The original data was pre-processed and subset to result in 3689 cells over 19 genetic
 309 perturbation (cells per condition: median [min, max] = 201 [72, 270]) with 20 proteins and 500 genes,
 310 which were directly used as input. See details in Appendix E. This dataset was also analyzed under
 311 balanced 5-fold splits and LOPO cross-validation. All imputation results for this dataset focus on
 312 predicting gene expression from protein level measurements.
 313

314 4.4 METRICS

315 Since our datasets contain true cell pairings, we assess OT matching accuracy using two ground-truth-
 316 based metrics. The *trace* of the normalized transport plan measures the proportion of true pairs being
 317 perfectly matched (Xi et al., 2024), e.g., it is 1.0 when all weight is assigned to true pairs and $1/N$
 318 for uninformative uniform assignments. The symmetric Barycentric Fraction of Cells Closer Than
 319 True Match (*Bary. FOSCTTM*) measures how much weight the transport plan incorrectly assigns to
 320 false pairs relative to true pairs (Demetci et al., 2022; Liu et al., 2019), where 0.0 indicates perfect
 321 matching and 0.5 represents random uniform assignment. We report the symmetric *Bary. FOSCTTM*
 322 by averaging performance across both modalities. Next, we quantify imputation accuracy using 6
 323 metrics: *MSE*, 1-Wasserstein distance (*WD*), Cosine Similarity (*Cos-sim*), *KNN Recall*, KNN average
 324 precession-recall (*KNN PR*) and KNN area under the receiver-operating characteristic curve (*KNN*
 325 *ROC*). *MSE* and *WD* are standard metrics in the imputation literature (Gorla et al., 2025), where

324
 325 Table 1: Matching performance metrics for top 5 method combinations in each shared proportion
 326 settings in simulations. SEs follow \pm ; best in bold, second-best underlined.

327 328	329 Shared Prop.	330 Method	331 Mean Rank	332 Trace	333 Bary. FOSCTTM
330 100	331 GROOVE (cosine)+ LabeledCOOT 332 GROOVE (cosine)+ LabeledEOT 333 DVAE+ LabeledCOOT 334 DVAE+ LabeledEOT 335 PS+LabeledEOT	336 1.5 3.0 3.5 4.5 5.5	337 0.856\pm0.027 0.466 \pm 0.020 0.669 \pm 0.034 0.453 \pm 0.023 0.366 \pm 0.006	338 0.027 \pm 0.006 0.026\pm0.004 0.066 \pm 0.009 0.042 \pm 0.006 0.054 \pm 0.006	339 3.5 4.5 5.5 6.0
334 80	335 DVAE+ LabeledEOT 336 GROOVE (cosine)+ LabeledCOOT 337 GROOVE (cosine)+ LabeledEOT 338 PS+LabeledEOT 339 DVAE+ LabeledCOOT	340 2.5 3.5 4.5 4.5 6.0	341 0.165 \pm 0.008 0.237\pm0.015 0.148 \pm 0.008 0.156 \pm 0.005 0.202 \pm 0.014	342 0.146\pm0.014 0.180 \pm 0.012 0.150 \pm 0.013 0.161 \pm 0.011 0.196 \pm 0.011	343 4.5 5.0 5.5 6.0
337 50	338 DVAE+ LabeledEOT 339 GROOVE (cosine)+ LabeledEOT 340 GROOVE (cosine)+ LabeledCOOT 341 GROOVE (tdist)+ LabeledEOT 342 DVAE+ LabeledCOOT	343 2.5 4.0 5.0 5.5 6.0	344 0.125 \pm 0.007 0.118 \pm 0.007 0.184\pm0.015 0.105 \pm 0.006 0.164 \pm 0.011	345 0.162\pm0.006 0.162\pm0.009 0.206 \pm 0.008 0.168 \pm 0.009 0.214 \pm 0.007	346 4.0 5.0 5.5 6.0

341 the former assess global error and the latter quantifies the distributional alignment. Cos-sim is a
 342 common metric in perturbation prediction tasks (Littman et al., 2025; Adduri et al., 2025). The
 343 KNN-based metrics are motivated by recent work from (Littman et al., 2025), which demonstrated
 344 that KNN metrics provide robust assessment of local similarity preservation in perturbation effect
 345 prediction tasks. KNN metrics avoid bias from single gene failures by evaluating neighborhood
 346 preservation rather than global prediction errors. All metrics are reported on held-out data (either test
 347 sets or cross-validation folds) with standard errors (SE; across folds or replicates). Standard errors are
 348 rounded to three decimal places; values reported as 0.000 indicate $SE < 0.0005$. For further details
 349 on metrics, see Appendix F.

350 5 RESULTS

351 5.1 OVERALL SIMULATED DATA PERFORMANCE

354 Table 1 presents the matching performance of the top-5 ranked methods across different shared
 355 proportion settings in simulated data. Under the standard 100% shared perturbation-relevant variation
 356 setting, GROOVE (combined with LabeledCOOT) achieves superior performance compared to all
 357 other method combinations. At 80% and 50% shared proportions, GROOVE ranks second on average;
 358 however, it maintains the highest trace-based matching performance across all shared proportion
 359 conditions. Since matching performance serves as an intermediate step toward downstream objectives,
 360 Table 2 evaluates downstream imputation performance in the same simulated data. GROOVE -based
 361 approaches again achieve the highest rankings under both 100% and 80% shared conditions, while
 362 a PS-based combination is ranked marginally higher under 50% shared conditions. **Statistical tests**
 363 (**Appendix G**) **confirm GROOVE achieves significant performance gains ($p < 0.05$) under 100%**
 364 **shared conditions across both matching and imputation tasks, with cross metric average win rates**
 365 **competitive with the next-highest ranking methods in other settings.** A closer inspection of the
 366 metrics reveals that GROOVE -based combinations consistently achieves top performance on nearly
 367 every individual metric except KNN PR, where they trail PS+LabeledEOT by 0.001 (much less than
 368 SE). Comparing the rankings between Tables 1 and 2, we notice that better matching performance
 369 does not guarantee optimal downstream task performance. For instance, DVAE demonstrates
 370 better average matching performance for shared proportion $< 100\%$, but PS-based combinations
 371 yield better imputation performance under the same conditions. Despite these variations, GROOVE
 372 -based combination deliver top performance under 100% shared variation and demonstrates robust,
 373 consistent performance, securing *atleast* top-two ranks across all evaluated scenarios.

374 5.2 MULTIMODAL SINGLE-CELL PERFORMANCE

375 We now evaluate performance on the more realistic perturb-CITE-seq dataset. Table 3 presents
 376 imputation performance for the top-10 method combinations under 5-fold cross-validation. In
 377 contrast to the simulated data results, GROOVE -based approaches occupy the top-5 ranks while

378

379
380
Table 2: Imputation performance metrics for top 5 method combinations in each shared proportion
settings in simulations. SEs follow \pm ; best in bold, second-best underlined.

Shared Prop.	Method	Mean Rank	MSE	Cos-sim	KNN Recall	KNN PR	KNN ROC	WD
100	GROOVE (cosine)+LabeledCOOT	1.00	0.099\pm0.012	<u>0.949\pm0.005</u>	0.444\pm0.029	<u>0.272\pm0.024</u>	0.706\pm0.015	<u>0.065\pm0.003</u>
	DAVAE+LabeledCOOT	2.00	0.121 \pm 0.013	0.931 \pm 0.006	0.412 \pm 0.027	0.244 \pm 0.021	0.689 \pm 0.014	0.092 \pm 0.005
	PS+LabeledCOOT	3.50	0.148 \pm 0.012	0.920 \pm 0.006	0.382 \pm 0.023	0.221 \pm 0.017	0.673 \pm 0.012	0.107 \pm 0.004
	GROOVE (tdist)+LabeledCOOT	4.00	0.146 \pm 0.014	0.931 \pm 0.004	0.372 \pm 0.020	0.216 \pm 0.014	0.668 \pm 0.011	0.097 \pm 0.005
	DAVAE+LabeledEOT	4.67	0.156 \pm 0.010	0.916 \pm 0.005	0.379 \pm 0.021	0.220 \pm 0.015	0.672 \pm 0.011	0.142 \pm 0.004
80	GROOVE (cosine)+LabeledEOT	2.83	0.623 \pm 0.119	0.837\pm0.011	0.241\pm0.014	0.137\pm0.008	0.598\pm0.008	0.199 \pm 0.027
	PS+LabeledEOT	3.00	0.593 \pm 0.115	0.836 \pm 0.010	0.239 \pm 0.013	0.135 \pm 0.007	0.597 \pm 0.007	0.198\pm0.027
	GROOVE (cosine)+LabeledCOOT	5.00	0.740 \pm 0.197	0.832 \pm 0.016	0.240 \pm 0.017	0.134 \pm 0.008	0.598\pm0.009	0.126 \pm 0.010
	PS+LabeledEGWOT	5.00	0.614 \pm 0.113	0.834 \pm 0.009	0.238 \pm 0.014	0.135 \pm 0.007	0.597 \pm 0.007	0.256 \pm 0.028
	PS+EOT	5.33	0.601\pm0.111	0.832 \pm 0.010	0.236 \pm 0.014	0.134 \pm 0.007	0.596 \pm 0.007	0.187 \pm 0.022
50	PS+LabeledEOT	3.83	0.516 \pm 0.040	0.830 \pm 0.011	0.215\pm0.011	0.123\pm0.006	0.585\pm0.006	<u>0.183\pm0.011</u>
	GROOVE (cosine)+EOT	4.00	0.519 \pm 0.039	0.838\pm0.009	0.215\pm0.009	0.122 \pm 0.006	0.585\pm0.005	0.208 \pm 0.012
	GROOVE (cosine)+LabeledEOT	4.00	0.513\pm0.040	0.835 \pm 0.011	0.212 \pm 0.010	0.121 \pm 0.006	0.583 \pm 0.005	0.194 \pm 0.012
	DAVAE+LabeledEOT	5.33	0.517 \pm 0.041	0.833 \pm 0.009	0.211 \pm 0.012	0.121 \pm 0.007	0.582 \pm 0.006	0.189 \pm 0.009
	GROOVE (cosine)+LabeledCOOT	6.00	0.559 \pm 0.049	0.832 \pm 0.012	0.213 \pm 0.010	0.121 \pm 0.006	0.584 \pm 0.005	0.134\pm0.009

394

PS-based approaches rank in the bottom two positions. This pattern supports the robustness of GROOVE -based combinations across diverse experimental conditions. Next, Table 8 (Appendix H) further shows that GROOVE -based approaches achieve the highest matching performance rankings. Notably, a PS-based approach ties with GROOVE -based methods for top matching performance; however, consistent with our simulation findings, superior matching performance does not necessarily translate to improved downstream imputation results. We further report performance of the top-10 methods under LOPO cross-validation, with matching and imputation performance results presented in Tables 9 and 10 (Appendix H), respectively. Even under this setting, GROOVE -based methods take the top-2 ranks in imputation. We also attain the single highest Trace and Bary. FOSCTTM metrics.

404

We next evaluate performance on the perturb-Multiome dataset. Tables 14 and 12 (Appendix I) summarize matching and imputation results, respectively, for the top 10 methods under five-fold cross-validation. Consistent with previous datasets, GROOVE attains the highest overall rank in both tasks. These sets of results in conjunction with the simulations indicates that cosine similarity kernel is often a good default but there are situations where the t -distribution kernel is helpful. We also again observe that high matching performance does not necessitate good imputation performance; DVAE’s better matching performance compared to PS-based approaches, yet DVAE fails to rank within the top-10 for imputation performance. Similarly, GROOVE -based methods occupy the the top-3 ranks of matching (Appendix I, Table. 13) and imputation performance under LOPO evaluation. Interestingly, EOT variants enable superior imputation performance in the perturb-Multiome dataset when holding the underlying representation learner fixed, while EGWOT variants (of which COOT is a member) are preferred for imputation tasks in the perturb-CITE-seq dataset. This pattern suggests that optimal OT aligner selection depends on dataset characteristics, including the specific modality pairs, independent of the chosen representation learner.

418

419

420

421
422
423
Table 3: Imputation performance metrics for top 10 method combinations in Perturb-CITE-seq dataset
with 5-fold evaluation. SEs follow \pm ; best in bold, second-best underlined, homogeneous metrics
unannotated.

Method	Mean Rank	MSE	Cos-sim	KNN Recall	KNN PR	KNN ROC	WD
GROOVE (cosine)+ LabeledEGWOT	6.83	0.261\pm0.001	0.049 \pm 0.002	0.020 \pm 0.001	0.017 \pm 0.000	0.503 \pm 0.000	0.353 \pm 0.001
GROOVE (cosine)+ LabeledCOOT	7.33	0.282 \pm 0.000	<u>0.016\pm0.002</u>	0.021\pm0.001	0.017 \pm 0.000	0.503 \pm 0.000	<u>0.297\pm0.000</u>
GROOVE (tdist)+ LabeledCOOT	7.50	0.282 \pm 0.001	0.018 \pm 0.001	0.020 \pm 0.001	0.017 \pm 0.000	0.503 \pm 0.000	<u>0.297\pm0.000</u>
GROOVE (tdist)+ LabeledEGWOT	7.67	0.261\pm0.000	0.046 \pm 0.002	<u>0.019\pm0.001</u>	0.017 \pm 0.000	0.502 \pm 0.000	<u>0.353\pm0.001</u>
GROOVE (tdist)+ LabeledEOT	7.67	0.262 \pm 0.000	0.042 \pm 0.002	0.020 \pm 0.001	0.017 \pm 0.000	0.503 \pm 0.000	0.347 \pm 0.002
DAVAE+ LabeledEGWOT	8.00	0.261\pm0.000	0.059\pm0.002	<u>0.018\pm0.001</u>	0.017 \pm 0.000	0.502 \pm 0.000	0.352 \pm 0.000
DAVAE+ LabeledCOOT	8.33	0.295 \pm 0.000	0.019 \pm 0.001	<u>0.020\pm0.001</u>	0.017 \pm 0.000	0.503 \pm 0.001	0.281\pm0.000
GROOVE (cosine)+ LabeledEOT	9.00	0.262 \pm 0.001	0.044 \pm 0.002	0.019 \pm 0.001	0.017 \pm 0.000	0.502 \pm 0.000	0.350 \pm 0.001
PS+ LabeledCOOT	9.67	0.295 \pm 0.001	0.018 \pm 0.002	0.019 \pm 0.001	0.017 \pm 0.000	0.502 \pm 0.000	0.281\pm0.001
PS+ LabeledEOT	10.00	<u>0.262\pm0.001</u>	0.040 \pm 0.001	0.019 \pm 0.001	0.017 \pm 0.000	0.502 \pm 0.000	0.344 \pm 0.001

432
433
434Table 4: Imputation performance metrics for top 10 method combinations in Perturb-Multiome dataset with leave one perturbation out evaluation. SEs follow \pm ; best in bold, second-best underlined.435
436
437
438
439
440
441
442
443
444

Method	Mean Rank	MSE	Cos-sim	KNN Recall	KNN PR	KNN ROC	WD
GROOVE (tdist)+ LabeledEOT	2.17	0.306\pm0.003	0.140\pm0.015	0.192 \pm 0.007	0.132 \pm 0.004	0.558 \pm 0.004	0.425 \pm 0.003
GROOVE (cosine)+ LabeledEOT	4.83	0.310 \pm 0.002	0.078 \pm 0.008	0.198\pm0.008	<u>0.133\pm0.004</u>	0.561\pm0.004	0.432 \pm 0.001
GROOVE (tdist)+ EOT	5.00	<u>0.309\pm0.002</u>	<u>0.100\pm0.012</u>	0.180 \pm 0.008	0.126 \pm 0.004	0.551 \pm 0.005	0.431 \pm 0.002
PS+ LabeledEOT	6.67	0.310 \pm 0.002	0.074 \pm 0.008	0.177 \pm 0.007	0.122 \pm 0.003	0.550 \pm 0.004	0.431 \pm 0.002
GROOVE (cosine)+ EOT	7.17	0.311 \pm 0.002	0.073 \pm 0.006	0.183 \pm 0.006	0.127 \pm 0.003	0.553 \pm 0.004	0.433 \pm 0.001
GROOVE (cosine)+ LabeledCOOT	7.67	0.357 \pm 0.014	0.072 \pm 0.030	0.170 \pm 0.008	0.122 \pm 0.003	0.546 \pm 0.004	0.326 \pm 0.004
DAVAE+ LabeledEOT	8.33	0.310 \pm 0.003	0.080 \pm 0.004	0.117 \pm 0.006	0.101 \pm 0.002	0.517 \pm 0.003	0.428 \pm 0.002
PS+ EOT	9.50	0.310 \pm 0.002	0.077 \pm 0.008	0.129 \pm 0.006	0.104 \pm 0.002	0.524 \pm 0.003	0.432 \pm 0.002
GROOVE (tdist)+ LabeledCOOT	10.33	0.363 \pm 0.013	0.050 \pm 0.030	0.161 \pm 0.008	0.119 \pm 0.004	0.541 \pm 0.004	0.324\pm0.005
DAVAE+ LabeledCOOT	10.50	0.344 \pm 0.010	0.062 \pm 0.023	0.148 \pm 0.007	0.113 \pm 0.003	0.534 \pm 0.004	0.348 \pm 0.004

444

5.3 ABLATION ANALYSIS

445

We conduct ablation analyses to assess the relative importance of key components within the GROOVE approach. These analyses fix the similarity kernel to cosine similarity and the aligner. We examine two ablation configurations: (1) 'No GroupCLIP': uses *on-the-fly* backtranslation with reconstruction and backtranslation losses but removes the GroupCLIP loss, and (2) 'Autoencoder only': uses standard reconstruction loss without GroupCLIP or backtranslation. We evaluate both configurations under the 80% shared proportion simulation setting (with LabeledEOT) and on the real Perturb-Multiome dataset (with EOT). Table H presents the ablation results. Across both simulated and real data, removing the GroupCLIP loss results in the largest performance decrease across all metrics. While both ablations show substantial degradation relative to full GROOVE, the difference between standard autoencoder and backtranslation (without GroupCLIP) is smaller by comparison, underscoring that GroupCLIP is the primary driver of performance gains. The additional supervision provided by GroupCLIP adds meaningful soft constraints to the latent representation. This supports our hypothesis that backtranslation alone is insufficient for encouraging group-level coherence and cross-group discrimination.

446

In the Pertub-Multiome dataset (under five-fold cross-validation), we see a meaningful increase across most metrics by adding *on-the-fly* backtranslation to the standard autoencoder framework. However, in simulations, there difference between these two approaches is minimal. We see two possible, but not mutually exclusive, explanations for this dependency. First, our simulations still do not capture the full complexity of real multi-modal single-cell data that backtranslation can leverage. Second, the original *on-the-fly* backtranslation framework (Artetxe et al., 2017) utilized a shared, pre-trained multi-modal encoder, which was not available within the scope of this work. These factors may jointly or independently lead to underestimating both the utility of backtranslation and GROOVE overall potential in simulations. Note that we try to performe ablations in the Pertub-CITE-seq dataset (Appendix H), but the metrics showed insufficient dynamic range to support meaningful conclusions.

447

We have additionally conducted a hyperparameter sweep to quantitatively assess sensitivity to key parameters. Figure 2 presents contour plots of average matching performance (Trace and Bary. FOSCTTM) across the 100%, 80%, and 50% shared proportion settings as functions of temperature (τ) and and reconstruction/backtranslation weight (β). The results reveal that GROOVE is relatively robust to τ but more sensitive to β . As with all deep learning approaches, we expect the hyperparameter sensitivity landscape to vary substantially across datasets and objectives. We recommend users perform dataset-specific hyperparameter optimization for best results.

448

449

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

472

473

474

475

476

477

478

479

Table 5: GROOVE ablation analysis performance metrics under 80% shared variation simulations and Perturb-Multiome dataset. SEs follow \pm ; best in bold, second-best underlined.480
481
482
483
484
485

Dataset	Abation Type	Bary. FOSCTTM	MSE	Cos-sim	KNN Recall	KNN PR	KNN ROC	WD
Perturb-Multiome	GROOVE (cosine)	0.489\pm0.007	0.308\pm0.003	0.102\pm0.028	0.044\pm0.007	0.028\pm0.002	0.512\pm0.004	0.428\pm0.005
	No GroupCLIP	0.497 \pm 0.001	0.310 \pm 0.001	0.076 \pm 0.007	0.036 \pm 0.002	0.025 \pm 0.000	0.507 \pm 0.001	0.432 \pm 0.002
	Autoencoder only	0.500 \pm 0.000	0.311 \pm 0.001	0.069 \pm 0.002	0.029 \pm 0.002	0.025 \pm 0.000	0.504 \pm 0.001	0.433 \pm 0.001
Simulations	GROOVE (cosine)	0.143\pm0.012	0.622\pm0.122	0.836\pm0.011	0.239\pm0.015	0.137\pm0.008	0.597\pm0.008	0.195\pm0.021
	No GroupCLIP	0.321 \pm 0.013	0.965 \pm 0.225	0.778 \pm 0.016	0.187 \pm 0.014	0.106 \pm 0.007	0.570 \pm 0.008	0.226 \pm 0.007
	Autoencoder only	0.280 \pm 0.013	0.846 \pm 0.175	0.793 \pm 0.013	0.195 \pm 0.012	0.111 \pm 0.006	0.574 \pm 0.006	0.240 \pm 0.012

486

6 DISCUSSION

488 This work introduces a multi-modal semi-supervised representation learning approach for weakly
 489 paired data leveraging a novel group contrastive loss (GroupCLIP) inside an *on-the-fly* backtranslating
 490 autoencoder. We also introduce combinatorial (with respect to OT aligners) benchmarking for multi-
 491 modal single cell perturbation alignment and cross-modal imputation. Empirical evaluation shows that
 492 GROOVE enables learning more useful latent representations for cell-level matching and downstream
 493 imputation. Ablations demonstrate our GroupCLIP is a crucial component for aligning the multi-
 494 modal representations and for good alignment performance. It is likely that GroupCLIP can have
 495 utility beyond just single cell data since weakly paired multi-modal data is ubiquitous in many
 496 domains (Yang et al., 2020; Mei et al., 2024; Sun et al., 2024). Our empirical results in Tables 1-4,
 497 powered by our robust evaluation framework, reveal that GroupCLIP-derived representations achieve
 498 particularly strong performance when paired with label-constrained OT methods (extended discussion
 499 on this topic in Appendix J.1). Furthermore, for the first time, we empirically show that optimal
 500 aligner choice varies across data modality pairs, methods, and shared variation percentages.

501 **Limitations.** While GROOVE demonstrates consistent performance across evaluated datasets,
 502 several limitations warrant discussion. First, like all contrastive approaches, our method is sensitive
 503 to hyperparameters; poor choices can lead to representation collapse or failure to leverage group
 504 structure. Second, our ablations reveal that backtranslation contributes are mixed, with greater
 505 effect in real data than in simulations, without pre-trained encoders, which were unavailable for
 506 our modalities. Third, the method assumes accurate perturbation labels, a standard assumption
 507 in single-cell perturbation screens where mislabeled samples are removed during quality control
 508 preprocessing. Label noise would corrupt any method using this as supervision signal, making
 509 this an experimental design consideration rather than a method-specific limitation. Fourth, neither
 510 GROOVE nor comparable methods have been evaluated under extreme class imbalance, which could
 511 significantly degrade performance. This should be more thoroughly evaluated in future work.

512 **Future directions.** Our work highlights several important directions for the community. First,
 513 developing more realistic simulation frameworks that incorporate variable proportions of shared
 514 versus modality-specific perturbation effects is critical, as the true regime in real datasets remains
 515 unknown. Second, there is a need for the community to establish consensus on evaluation priorities:
 516 whether sample matching or downstream tasks like imputation should serve as the primary perfor-
 517 mance criterion. Our results demonstrate that superior matching performance does not guarantee
 518 effective downstream imputation, indicating these objectives may require different methodological
 519 approaches. This prioritization decision is critical, as the optimal solutions for matching tasks likely
 520 differ from those that excel in downstream applications. Indeed our empirical results in both sim-
 521 ulated and real data likely point to a tradeoff between matching and imputation when perturbation
 522 effects are not fully shared across modalities (extended discussion in Appendix J.2). Third, more
 523 sensitive imputation metrics with greater dynamic range are needed for robust perturbation prediction
 524 evaluation. Development of these metrics necessitates close partnership with biologists to define
 525 relevant research goals, facilitating the creation of evaluation frameworks that accurately capture
 526 performance on scientifically meaningful tasks. Fourth, to our knowledge, this is the first application
 527 of backtranslation architectures to single-cell data. While our ablations show some gains without
 528 pre-trained encoders, this framework provides a natural foundation for future multi-modal single-cell
 529 applications and should be further explored, particularly as suitable pre-trained models become avail-
 530 able. Fifth, GroupCLIP and GROOVE do not make any strong assumptions limiting it to bimodal
 531 data and should generalize naturally to > 2 modalities, empirical validation of GROOVE and related
 532 methods in such setting remains valuable future work. Finally, systematic evaluation under extreme
 533 class imbalance would clarify method robustness boundaries.

534
 535
 536
 537
 538
 539

540 REFERENCES
541

542 Abhinav K Adduri, Dhruv Gautam, Beatrice Bevilacqua, Alishba Imran, Rohan Shah, Mohsen
543 Naghipourfar, Noam Teyssier, Rajesh Ilango, Sanjay Nagaraj, Mingze Dong, et al. Predicting
544 cellular responses to perturbation across diverse contexts with state. *bioRxiv*, pp. 2025–06, 2025.

545 Shaden Alshammari, John Hershey, Axel Feldmann, William T. Freeman, and Mark Hamilton. I-con:
546 A unifying framework for representation learning, 2025. URL <https://arxiv.org/abs/2504.16929>.

547

548 Ricard Argelaguet, Damien Arnol, Danila Bredikhin, Yonatan Deloro, Britta Velten, John C Marioni,
549 and Oliver Stegle. Mofa+: a statistical framework for comprehensive integration of multi-modal
550 single-cell data. *Genome biology*, 21(1):111, 2020.

551

552 Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. Unsupervised neural machine
553 translation. *arXiv preprint arXiv:1710.11041*, 2017.

554

555 Tal Ashuach, Mariano I Gabitto, Rohan V Koodli, Giuseppe-Antonio Saldi, Michael I Jordan, and
556 Nir Yosef. Multivi: deep generative model for the integration of multimodal data. *Nature Methods*,
557 20(8):1222–1231, 2023.

558

559 Ashish Bhandari, Rebecca Sarto Basso, Daria Beshnova, Aurora Blucher, Carl Brooks, Lu Chen,
560 China Chien, Jacob Cooper, Michael Cuccarese, Hayley Donnella, Bryan Ellis, Marie Evangelista,
561 Kevin Fales, Imran Haque, Aimee Iberg, Weston Judd, Kiran Nadella, Chase Neumann, Janet
562 Paulsen, Paul Rearden, Shane Rowley, Jenny Rudnick, Ashraf Saeed, Bahar Shamloo, and Lee
563 Zamparo. Identification and optimization of novel small molecule modulators of immune check-
564 point resistance with a unified representation space for genomic and chemical perturbations. In
565 *AACR*, 2022.

566

567 Jan Niklas Böhm, Philipp Berens, and Dmitry Kobak. Unsupervised visualization of image datasets
568 using contrastive learning. *arXiv preprint arXiv:2210.09879*, 2022.

569

570 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
571 contrastive learning of visual representations. In *International conference on machine learning*, pp.
572 1597–1607. PMLR, 2020.

573

574 Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrau, and Marco Baroni. What
575 you can cram into a single vector: Probing sentence embeddings for linguistic properties. *arXiv
576 preprint arXiv:1805.01070*, 2018.

577

578 Haotian Cui, Alejandro Tejada-Lapuerta, Maria Brbić, Julio Saez-Rodriguez, Simona Cristea, Hani
579 Goodarzi, Mohammad Lotfollahi, Fabian J. Theis, and Bo Wang. Towards multimodal foundation
580 models in molecular cell biology. *Nature*, 2025.

581

582 Pinar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara Singh.
583 Scot: single-cell multi-omics alignment with optimal transport. *Journal of computational biology*,
584 29(1):3–18, 2022.

585

586 Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P Fulco, Livnat Jerby-Arnon, Nemanja D
587 Marjanovic, Danielle Dionne, Tyler Burks, Rakimra Raychndhury, Britt Adamson, Thomas M
588 Norman, Eric S Lander, Jonathan S Weissman, Nir Friedman, and Aviv Regev. Perturb-seq:
589 Dissecting molecular circuits with scalable single-cell rna profiling of pooled genetic screens. *Cell*,
590 2016.

591

592 David Feldman, Avtar Singh, Jonathan L Schmid-Burgk, Rebecca J Carlson, Anja Mezger, Anthony J
593 Garrity, Feng Zhang, and Paul C Blainey. Optical pooled screens in human cells. *Cell*, 2019.

594

595 Chris J Frangieh, Johannes C Melms, Pratiksha I Thakore, Kathryn R Geiger-Schuller, Patricia
596 Ho, Adrienne M Luoma, Brian Cleary, Livnat Jerby-Arnon, Shruti Malu, Michael S Cuoco, et al.
597 Multimodal pooled perturb-cite-seq screens in patient models define mechanisms of cancer immune
598 evasion. *Nature genetics*, 53(3):332–341, 2021.

594 Kai Gan, Bo Ye, Min-Ling Zhang, and Tong Wei. Semi-supervised clip adaptation by enforc-
 595 ing semantic and trapezoidal consistency. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu
 596 (eds.), *International Conference on Representation Learning*, volume 2025, pp. 12142–12163,
 597 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/2058b90b1e0699770840471fc294d1e7-Paper-Conference.pdf.

598

599 Aditya Gorla, Ryan Wang, Zhengtong Liu, Ulzee An, and Sriram Sankararaman. CACTI: Leveraging
 600 copy masking and contextual information to improve tabular data imputation. In *Proceedings of*
 601 *the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine*
 602 *Learning Research*, pp. 20187–20225. PMLR, 13–19 Jul 2025.

603

604 Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena
 605 Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar,
 606 et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural*
 607 *information processing systems*, 33:21271–21284, 2020.

608

609 Jeff Z HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised
 610 deep learning with spectral contrastive loss. *Advances in neural information processing systems*,
 611 34:5000–5011, 2021.

612

613 Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for
 614 unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on*
 615 *computer vision and pattern recognition*, pp. 9729–9738, 2020.

616

617 Tianyang Hu, Zhili Liu, Fengwei Zhou, Wenjia Wang, and Weiran Huang. Your contrastive learning
 618 is secretly doing stochastic neighbor embedding. *arXiv preprint arXiv:2205.14814*, 2022.

619

620 Leonid V Kantorovich. Mathematical methods of organizing and planning production. *Management*
 621 *science*, 6(4):366–422, 1960.

622

623 Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron
 624 Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural*
 625 *information processing systems*, 33:18661–18673, 2020.

626

627 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*
 628 *arXiv:1312.6114*, 2013.

629

630 Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. Phrase-
 631 based & neural unsupervised machine translation. *arXiv preprint arXiv:1804.07755*, 2018.

632

633 Phuc H Le-Khac, Graham Healy, and Alan F Smeaton. Contrastive representation learning: A
 634 framework and review. *Ieee Access*, 8:193907–193934, 2020.

635

636 Kevin Z Lin and Nancy R Zhang. Quantifying common and distinct information in single-cell
 637 multimodal data with tilted canonical correlation analysis. *Proceedings of the National Academy*
 638 *of Sciences*, 120(32):e2303647120, 2023.

639

640 Russell Littman, Jacob Levine, Sepideh Maleki, Yongju Lee, Vladimir Ermakov, Lin Qiu, Alexander
 641 Wu, Kexin Huang, Romain Lopez, Gabriele Scialia, et al. Gene-embedding-based prediction and
 642 functional evaluation of perturbation expression responses with presage. *bioRxiv*, pp. 2025–06,
 643 2025.

644

645 Jie Liu, Yuanhao Huang, Ritambhara Singh, Jean-Philippe Vert, and William Stafford Noble. Jointly
 646 embedding multiple single-cell omics measurements. In *Algorithms in bioinformatics:... International*
 647 *Workshop, WABI..., proceedings. WABI (Workshop)*, volume 143, pp. 10, 2019.

648

649 Romain Lopez, Jeffrey Regier, Michael B Cole, Michael I Jordan, and Nir Yosef. Deep generative
 650 modeling for single-cell transcriptomics. *Nature methods*, 15(12):1053–1058, 2018.

651

652 Jorge Diego Martin-Rufino, Alexis Caulier, Seayoung Lee, Nicole Castano, Emily King, Samantha
 653 Joubran, Marcus Jones, Seth R Goldman, Uma P Arora, Lara Wahlster, et al. Transcription factor
 654 networks disproportionately enrich for heritability of blood cell phenotypes. *Science*, 388(6742):
 655 52–59, 2025.

648 Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D Plumbley,
 649 Yuexian Zou, and Wenwu Wang. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning
 650 dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech, and*
 651 *Language Processing*, 32:3339–3354, 2024.

652 Facundo Mémoli. Gromov–wasserstein distances and the metric approach to object matching.
 653 *Foundations of computational mathematics*, 11(4):417–487, 2011.

654 Sangwoo Mo, Minkyu Kim, Kyungmin Lee, and Jinwoo Shin. S-clip: semi-supervised vision-
 655 language learning using few specialist captions. In *Proceedings of the 37th International Confer-
 656 ence on Neural Information Processing Systems*, NIPS ’23, Red Hook, NY, USA, 2023. Curran
 657 Associates Inc.

658 Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive
 659 coding. *arXiv preprint arXiv:1807.03748*, 2018.

660 Gabriel Peyré, Marco Cuturi, and Justin Solomon. Gromov-wasserstein averaging of kernel and
 661 distance matrices. In *International conference on machine learning*, pp. 2664–2672. PMLR, 2016.

662 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 663 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 664 models from natural language supervision. In *International conference on machine learning*, pp.
 665 8748–8763. PmLR, 2021.

666 Jennifer E. Rood, Anna Hupalowska, and Aviv Regev. Toward a foundation model of causal cell and
 667 tissue biology with a perturbation cell and tissue atlas. *Cell*, 2024.

668 Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies.
 669 *Journal of educational Psychology*, 66(5):688, 1974.

670 Jayoung Ryu, Charlotte Bunne, Luca Pinello, Aviv Regev, and Romain Lopez. Cross-modality
 671 matching and prediction of perturbation responses with labeled gromov-wasserstein optimal
 672 transport, 2025. URL <https://arxiv.org/abs/2405.00838>.

673 Jules Samaran, Gabriel Peyré, and Laura Cantini. scconfluence: single-cell diagonal integration with
 674 regularized inverse optimal transport on weakly connected features. *Nature Communications*, 15
 675 (1):7762, 2024.

676 Othmane Sebbouh, Marco Cuturi, and Gabriel Peyré. Structured transforms across spaces with
 677 cost-regularized optimal transport. In *International Conference on Artificial Intelligence and
 678 Statistics*, pp. 586–594. PMLR, 2024.

679 Rico Sennrich, Barry Haddow, and Alexandra Birch. Improving neural machine translation models
 680 with monolingual data. *arXiv preprint arXiv:1511.06709*, 2015.

681 Xintian Sun, Benji Peng, Charles Zhang, Fei Jin, Qian Niu, Junyu Liu, Keyu Chen, Ming Li, Pohsun
 682 Feng, Ziqian Bi, et al. From pixels to prose: Advancing multi-modal language models for remote
 683 sensing. *arXiv preprint arXiv:2411.05826*, 2024.

684 Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *European
 685 conference on computer vision*, pp. 776–794. Springer, 2020.

686 Vayer Titouan, Ievgen Redko, Rémi Flamary, and Nicolas Courty. Co-optimal transport. *Advances in
 687 neural information processing systems*, 33:17559–17570, 2020.

688 Hugues Van Assel, Cédric Vincent-Cuaz, Nicolas Courty, Rémi Flamary, Pascal Frossard, and
 689 Titouan Vayer. Distributional reduction: Unifying dimensionality reduction and clustering with
 690 gromov-wasserstein. *arXiv preprint arXiv:2402.02239*, 2024.

691 Hugues Van Assel, Mark Ibrahim, Tommaso Biancalani, Aviv Regev, and Randall Balestrierio. Joint
 692 embedding vs reconstruction: Provable benefits of latent space prediction for self supervised
 693 learning. *Advances in neural information processing systems*, 39, 2025.

702 Titouan Vayer. A contribution to optimal transport on incomparable spaces. *arXiv preprint*
703 [arXiv:2011.04447](https://arxiv.org/abs/2011.04447), 2020.

704

705 Johnny Xi, Jana Osea, Zuheng Xu, and Jason Hartford. Propensity score alignment of unpaired
706 multimodal data, 2024. URL <https://arxiv.org/abs/2404.01595>.

707

708 Karren Dai Yang, Anastasiya Belyaeva, Saradha Venkatachalam, Karthik Damodaran, Abigail
709 Katcoff, Adityanarayanan Radhakrishnan, GV Shivashankar, and Caroline Uhler. Multi-domain
710 translation between single-cell imaging and sequencing data using autoencoders. *Nature communications*, 12(1):31, 2021.

711

712 Zhenguo Yang, Zehang Lin, Lingni Guo, Qing Li, and Wenyin Liu. Mmed: a multi-domain and
713 multi-modality event dataset. *Information Processing & Management*, 57(6):102315, 2020.

714

715 Heming Yao, Phil Hanslovsky, Jan-Christian Huetter, Burkhard Hoeckendorf, and David Richmond.
716 Weakly supervised set-consistency learning improves morphological profiling of single-cell images.
717 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
6978–6987, 2024.

718

719 Shengjia Zhao, Jiaming Song, and Stefano Ermon. Infovae: Balancing learning and inference
720 in variational autoencoders. In *Proceedings of the aaai conference on artificial intelligence*,
volume 33, pp. 5885–5892, 2019.

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A GROOVE EXTENDED DETAILS
757758 A.1 ON-THE-FLY BACKTRANSLATION
759

760 Each modality-specific encoder $g_\theta^{(m)}$ consists of a 3-layer MLP with batch normalization and ReLU
761 activations, while each decoder $d_\theta^{(m)}$ uses a 2-layer MLP with an additional linear output layer (see
762 next subsection). A shared linear projection (a coupling layer) f_θ connects all modality-specific
763 encoders to their respective decoders (i.e. all modality specific embeddings are passed through this),
764 giving the final encoder composition $f_\theta^{(m)} = f_\theta \circ g_\theta^{(m)}$.
765

766 In weakly paired single-cell data, each modality can be viewed as a different language with shared
767 semantic information. We encourage cross-modally entangled representations by ensuring latents
768 enable meaningful cross-modal translation through a self-backtranslation strategy. Let \bar{m} denote the
769 other modality. Given a sample’s latent representation from modality m : $\mathbf{z}_i^{(m)} = f_\theta \circ g_\theta^{(m)}(\mathbf{x}_i^{(m)})$,
770 we perform cross-modal translation in three steps:

- 771 1. **Cross-modal generation:** Using the decoder of modality \bar{m} in inference mode, generate
772 $\mathbf{x}_i^{(m \rightarrow \bar{m})} = d_\theta^{(\bar{m})}(\mathbf{z}_i^{(m)})$
773
- 774 2. **Re-encoding:** Switch to training mode and encode the generated sample:
775 $\mathbf{z}_i^{(m \rightarrow \bar{m})} = f_\theta \circ g_\theta^{(\bar{m})}(\mathbf{x}_i^{(m \rightarrow \bar{m})})$
776
- 777 3. **Backtranslation:** Reconstruct the original modality: $\mathbf{x}_i^{(m \rightarrow \bar{m} \rightarrow m)} = d_\theta^{(m)}(\mathbf{z}_i^{(m \rightarrow \bar{m})})$

778 This *on-the-fly* (*on-the-fly*) process creates synthetic pseudo-paired samples within each mini-batch,
779 contingent on cross-modally informative latent representations.
780

781 A.2 ENCODER ARCHITECTURE AND REGULARIZATION
782

783 The first layer of the encoder projects each data modality from its native dimension $k^{(m)}$ to twice the
784 size of the final embedding dimension ($2 \times d$). This design follows the recommendation of Samaran
785 et al. (2024), who employ a variational-like encoder architecture where, for each d -dimensional
786 embedding, the encoder outputs both a mean and log-variance parameter.
787

788 However, unlike standard variational autoencoders, we do not optimize the variational Evidence
789 Lower BOund (ELBO), which combines reconstruction loss with Kullback-Leibler (KL) divergence.
790 Instead, we minimize only the standard reconstruction loss (like a vanilla autoencoder). This design
791 choice is motivated by prior work (Zhao et al., 2019), which demonstrated that KL divergence
792 can conflict with reconstruction objectives and degrade downstream inference performance—a
793 phenomenon we also observed in our internal analyses.
794

795 Our approach leverages the encoder’s outputted parameters to define a Gaussian posterior distribution
796 with diagonal covariance. Specifically, we interpret the $2d$ -dimensional encoder output as mean μ
797 and log-variance $\log \sigma^2$ parameters for a d -dimensional Gaussian distribution. During training, we
798 sample from this distribution using the reparameterization trick (Kingma & Welling, 2013).
799

800 To ensure numerical stability and provide mild regularization, we add a small fixed constant (10^{-4})
801 to the diagonal covariance matrix during training only. This stochastic sampling mechanism serves
802 three important purposes: (1) it introduces non-determinism in the encoder during training, (2) it
803 prevents overcrowding of samples in the latent space by encouraging distributional rather than point
804 estimates, and (3) it provides implicit regularization for the decoder, helping to prevent overfitting by
805 requiring it to reconstruct from a distribution of latent codes rather than deterministic points.
806

807 During inference, the encoder operates deterministically using only the mean parameters μ .
808

809 A.3 DECODER

810 The decoder takes the d -dimensional output from the shared projection layer f_θ and passes it through
811 a 2-layer MLP with 1D batch normalization and ReLU activation and then a final linear layer top
812 project the embeddings from d -dimensions back to the native $k^{(m)}$ features
813

810
811 A.4 SIMILARITY KERNELS

812 The choice of similarity kernel significantly influences the quality and characteristics of learned
 813 representations. While cosine similarity remains the standard choice in most contrastive approaches,
 814 learning representations on a hypersphere (Radford et al., 2021; Chen et al., 2020), recent work
 815 demonstrates that a heavy tailed t -distribution parameterized can yield more expressive representations
 816 in Euclidean space (Böhm et al., 2022; Hu et al., 2022).

817
818
$$\text{cosine similarity} \quad \quad \quad t\text{-distribution similarity}$$

 819
$$\text{sim}(\mathbf{a}, \mathbf{b}) := \exp\left(\frac{\langle \mathbf{a}, \mathbf{b} \rangle}{\|\mathbf{a}\|_2^2 \cdot \|\mathbf{b}\|_2^2} \cdot \frac{1}{\tau}\right) \quad (7) \quad \text{sim}(\mathbf{a}, \mathbf{b}) := \left[1 + \frac{\|\mathbf{a} - \mathbf{b}\|_2^2}{\tau\eta}\right]^{-\frac{\eta+1}{2}} \quad (8)$$

 820
 821

822 such that τ controls the temperature (bandwidth) parameter and η is the degrees of freedom (which is
 823 set 1 for this work).

824

825 A.5 MINI-BATCH BALANCED UNDERSAMPLING STRATEGY

826

827 To address possible class imbalance in our GroupCLIP framework, we implement a balanced under-
 828 sampling strategy that ensures equal representation across all classes within each batch. The algorithm
 829 first computes the instance count n_l for each label l in the first modality, assuming identical label
 830 distributions across both modalities. We then identify the minority class with the smallest instance
 831 count: $n_{\min} = \min_i(n_l)$.

832 Given an initial batch size B and number of labels L , we compute the effective batch size using:
 833 $B_{\text{eff}} = B - (B \bmod L)$, where each label contributes exactly B_{eff}/L samples per batch. This
 834 ensures the batch size is evenly divisible by the number of labels. Importantly, our sampling strategy
 835 respects the minority labels constraint by never sampling more than $(B_{\text{eff}}/L) \leq n_{\min}$ instances from
 836 any label across the entire training process, preventing oversampling of the minority label.

837 For each batch, we randomly sample B_{eff}/L instances from each label in both modalities indepen-
 838 dently, ensuring balanced representation while maintaining the weakly paired nature of the data. This
 839 approach prevents dominant labels from overwhelming the contrastive learning signal and ensures
 840 that all labels contribute equally to the learned representations, which is particularly important when
 841 dealing with imbalanced multi-modal datasets.

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 A.6 ALGORITHM SKETCH
 865

Algorithm 1 GROOVE Training Procedure

866 **Require:** Modality-specific datasets $\mathcal{D}^{(1)}, \mathcal{D}^{(2)}$
 867 **Require:** Encoders $g_\theta^{(1)}, g_\theta^{(2)}$, Decoders $d_\theta^{(1)}, d_\theta^{(2)}$, Shared Projection f_θ
 870 **Require:** Hyperparameters α (reconstruction), β (GroupCLIP)

871 1: **for** while max training iteration is not reached **do**
 872 2: Sample balanced mini-batches $(\mathbf{x}^{(1)}, t^{(1)}) \sim \mathcal{D}^{(1)}, (\mathbf{x}^{(2)}, t^{(2)}) \sim \mathcal{D}^{(2)}$ \triangleright See Section A.5
 873 **Step 1: Within-Modality Reconstruction and Contrastive Alignment**
 874 3: $\mathbf{z}^{(1)} \leftarrow f_\theta(g_\theta^{(1)}(\mathbf{x}^{(1)}))$
 875 4: $\mathbf{z}^{(2)} \leftarrow f_\theta(g_\theta^{(2)}(\mathbf{x}^{(2)}))$
 876 5: $\hat{\mathbf{x}}^{(1)} \leftarrow d_\theta^{(1)}(\mathbf{z}^{(1)})$
 877 6: $\hat{\mathbf{x}}^{(2)} \leftarrow d_\theta^{(2)}(\mathbf{z}^{(2)})$
 878 7: $\mathcal{L}_{\text{recon}} \leftarrow \sum_{m \in \{1,2\}} \text{MSE}(\mathbf{x}^{(m)}, \hat{\mathbf{x}}^{(m)})$
 879 8: $\mathcal{L}_{\text{GroupCLIP}} \leftarrow \frac{1}{2|\mathcal{D}_z^{(1)}|} \sum_{\mathbf{z}^{(1)} \in \mathcal{D}_z^{(1)}} \ell^{(1)} + \frac{1}{2|\mathcal{D}_z^{(2)}|} \sum_{\mathbf{z}^{(2)} \in \mathcal{D}_z^{(2)}} \ell^{(2)}$ \triangleright See Equation 4
 880 9: $\mathcal{L}^{\text{step1}} \leftarrow \beta \cdot \mathcal{L}_{\text{recon}} + \alpha \cdot \mathcal{L}_{\text{GroupCLIP}}$
 881 10: Update parameters θ using gradients from $\mathcal{L}^{\text{step1}}$
 882 **Step 2: On-the-Fly Backtranslation**
 883 11: Set $g_\theta^{(1)}, g_\theta^{(2)}, d_\theta^{(1)}, d_\theta^{(2)}, f_\theta$ to Eval \triangleright Generate cross-modal samples w/ nograd
 884 12: $\mathbf{x}^{(1 \rightarrow 2)} \leftarrow d_\theta^{(2)} \circ f_\theta \circ g_\theta^{(1)}(\mathbf{x}^{(1)})$
 885 13: $\mathbf{x}^{(2 \rightarrow 1)} \leftarrow d_\theta^{(1)} \circ f_\theta \circ g_\theta^{(2)}(\mathbf{x}^{(2)})$
 886 14: Set $g_\theta^{(1)}, g_\theta^{(2)}, d_\theta^{(1)}, d_\theta^{(2)}, f_\theta$ to Train
 887 15: $\mathbf{z}^{(1 \rightarrow 2)} \leftarrow f_\theta(g_\theta^{(2)}(\mathbf{x}^{(1 \rightarrow 2)}))$ \triangleright Re-encode translated samples
 888 16: $\mathbf{z}^{(2 \rightarrow 1)} \leftarrow f_\theta(g_\theta^{(1)}(\mathbf{x}^{(2 \rightarrow 1)}))$
 889 17: $\hat{\mathbf{x}}^{(1 \rightarrow 2 \rightarrow 1)} \leftarrow d_\theta^{(1)}(\mathbf{z}^{(1 \rightarrow 2)})$ \triangleright Reconstruct original modality
 890 18: $\hat{\mathbf{x}}^{(2 \rightarrow 1 \rightarrow 2)} \leftarrow d_\theta^{(2)}(\mathbf{z}^{(2 \rightarrow 1)})$
 891 19: $\mathcal{L}_{\text{bt}} \leftarrow \frac{1}{2} [\text{MSE}(\mathbf{x}^{(1)}, \hat{\mathbf{x}}^{(1 \rightarrow 2 \rightarrow 1)}) + \text{MSE}(\mathbf{x}^{(2)}, \hat{\mathbf{x}}^{(2 \rightarrow 1 \rightarrow 2)})]$
 892 20: $\mathcal{L}^{\text{step2}} \leftarrow \beta \cdot \mathcal{L}_{\text{bt}}$
 893 21: Update parameters θ using gradients from $\mathcal{L}^{\text{step2}}$
 894 22: **return** trained encoders $f_\theta^{(1)}, f_\theta^{(2)}$

895
 896
 897
 898
 899
 900
 901
 902
 903
 904
 905
 906
 907
 908
 909
 910
 911
 912
 913
 914
 915
 916
 917

918 **B BASELINE DETAILS**
 919

920 Baseline Implementation Details All methods use identical encoder and decoder architectures² with a
 921 consistent latent embedding dimension of $d = 128$ across all experiments. For DVAE, we adopt
 922 the previously reported optimal hyperparameter settings (Ryu et al., 2025). We evaluate GROOVE
 923 with both cosine and t-distribution (tdist) similarity kernels, setting $\alpha = 1.0, \tau = 0.2$, for all
 924 experiments and set $\beta = 0.1$ for single-cell data training. Note that we did not perform any rigorous
 925 hyperparameter exploration, we do not claim the reported results are the best-case performance of
 926 GROOVE for any of the evaluated data sets.

927 We evaluated all representation learners in conjunction with five OT approaches: two standard
 928 methods and three label-constrained variants. The standard OT approaches include EOT and EG-
 929 WOT (Peyré et al., 2016; Kantorovich, 1960). The label constrained approaches include: labeledEOT,
 930 labeledEGWOT and labeledCOOT (co-optimal transport) (Ryu et al., 2025; Titouan et al., 2020). For
 931 all analyses, we use the default entropic regularizer settings from the Perturb-OT package³.
 932

933 **C SIMULATION DETAILS**
 934

935 To simulate complex multi-modal cellular data, we define a probabilistic model that captures shared
 936 and modality-specific latent structures. The model begins by defining a shared latent space and two
 937 unshared, modality-specific latent spaces, which are then combined. The following perturbation types
 938 are used:

939

- 940 1. Shared perturbations: coordinately affect both modalities through the shared latent space
 941 with identical effect sizes and cell-specific penetrance values
- 942 2. Modality-specific perturbations: independently target each modality’s unique dimensions
 943 with separate effect sizes and penetrance parameters

944 We initialize the simulations with the following settings:
 945

- 946 • Latent dimensions: 10
- 947 • Shared variation proportions:
 - 949 – 100% shared: 10 shared, 0 unique dimensions
 - 950 – 80% shared: 8 shared, 2 unique dimensions per modality
 - 951 – 50% shared: 5 shared, 5 unique dimensions per modality
- 952 • Experimental design: 9 perturbations + 1 control condition
- 953 • Sample size: 100 cells per condition per modality
- 954 • Feature size: 1000 and 500 observed features for modalities 1 and 2, respectively

956 **C.1 GENERATIVE MODEL**

957 We begin by defining the latent variables. A shared latent variable Z is sampled coefficient-wise from
 958 a scaled standard normal distribution for n total cells and d_s shared dimensions:
 959

$$960 Z_{ij} \sim \mathcal{N}(0, \text{scale}^2) \quad \text{for } i = 1, \dots, n, j = 1, \dots, d_s$$

963 where scale is a global latent signal strength (default 0.1).

964 Similarly, we define two unshared latent variables, U_X and U_Y , with d_u unshared dimensions:
 965

$$966 U_{X,ij} \sim \mathcal{N}(0, \text{scale}^2), \quad U_{Y,ij} \sim \mathcal{N}(0, \text{scale}^2).$$

967 These are concatenated to form the full latent representations for each modality:
 968

$$969 V_X = [Z \parallel U_X], \quad V_Y = [Z \parallel U_Y],$$

971 ²PS does not require a decoder

³<https://github.com/Genentech/Perturb-OT>

972 with total dimensionality $d = d_s + d_u$.
 973

974 Each modality is associated with a transformation matrix. For modality X with p_X features, the
 975 coefficients of $A_X \in \mathbb{R}^{d \times p_X}$ are sampled as

$$(A_X)_{jk} \sim \mathcal{N}(0, 1).$$

977 Similarly, for modality Y with p_Y features:

$$(A_Y)_{jk} \sim \mathcal{N}(0, 1).$$

980 Bias vectors $b_X \in \mathbb{R}^{p_X}$ and $b_Y \in \mathbb{R}^{p_Y}$ have coefficients

$$(b_X)_j \sim \mathcal{N}(0, 1), \quad (b_Y)_j \sim \mathcal{N}(0, 1).$$

983 Scaling factors $s_X \in \mathbb{R}^{p_X}$ and $s_Y \in \mathbb{R}^{p_Y}$ control feature variability:

$$(s_X)_j \sim \Gamma(1, 1), \quad (s_Y)_j \sim \Gamma(1, 1).$$

986 To account for modality-specific noise, perturbation parameters are defined coefficient-wise:

$$\begin{aligned} \mu_{X,j} &\sim \mathcal{N}(0, 1), & \mu_{Y,j} &\sim \mathcal{N}(0, 1), \\ \text{offsetsd}_{X,j} &\sim \mathcal{N}(0, 1), & \text{offsetsd}_{Y,j} &\sim \mathcal{N}(0, 1), \\ \sigma_{X,j} &= \exp(-3.0 + \text{offsetsd}_{X,j}), & \sigma_{Y,j} &= \exp(-3.0 + \text{offsetsd}_{Y,j}). \end{aligned}$$

991 The noise for each cell i and feature j is then:

$$\xi_{X,ij} \sim \mathcal{N}(\mu_{X,j}, \sigma_{X,j}^2) \cdot \frac{\text{scale}}{\text{snr}}, \quad \xi_{Y,ij} \sim \mathcal{N}(\mu_{Y,j}, \sigma_{Y,j}^2) \cdot \frac{\text{scale}}{\text{snr}},$$

995 where snr is the signal-to-noise ratio (default 0.2).

996 We next incorporate L perturbations across $L + 1$ conditions (including one control). Perturbations
 997 target shared or unshared latent dimensions in a cyclic manner.

998 For **shared perturbations**, the target dimension is

$$t_s(I) = ((I - 1) \bmod d_s).$$

1001 Effect sizes are sampled as

$$\begin{aligned} |e_s(I)| &\sim \max(3, \Gamma(1, 1)), & \text{sign}_s(I) &\sim 2 \cdot \text{Bernoulli}(0.5) - 1, \\ e_s(I) &= \text{sign}_s(I) \cdot |e_s(I)|. \end{aligned}$$

1004 Each cell i has penetrance

$$q_{s,i} \sim \text{Beta}(1, 10).$$

1007 For cells under perturbation I , the targeted latent dimension is shifted in both modalities:

$$v'_{X,its(I)} = v_{X,its(I)} + e_s(I)q_{s,i}, \quad v'_{Y,its(I)} = v_{Y,its(I)} + e_s(I)q_{s,i}.$$

1009 For **modality-specific perturbations**, the target dimension is

$$t_u(I) = ((I - 1) \bmod d_u) + d_s.$$

1012 Effect sizes are defined separately for each modality:

$$\begin{aligned} |e_{uX}(I)| &\sim \max(3, \Gamma(1, 1)), & \text{sign}_{uX}(I) &\sim 2 \cdot \text{Bernoulli}(0.5) - 1, & e_{uX}(I) &= \text{sign}_{uX}(I)|e_{uX}(I)|, \\ |e_{uY}(I)| &\sim \max(3, \Gamma(1, 1)), & \text{sign}_{uY}(I) &\sim 2 \cdot \text{Bernoulli}(0.5) - 1, & e_{uY}(I) &= \text{sign}_{uY}(I)|e_{uY}(I)|. \end{aligned}$$

1016 With penetrance

$$q_{uX,i} \sim \text{Beta}(1, 10), \quad q_{uY,i} \sim \text{Beta}(1, 10),$$

1018 the latent variables are perturbed independently:

$$v''_{X,its(I)} = v'_{X,its(I)} + e_{uX}(I)q_{uX,i}, \quad v''_{Y,its(I)} = v'_{Y,its(I)} + e_{uY}(I)q_{uY,i}.$$

1021 Let V_X^{pert} and V_Y^{pert} denote the final latent spaces. The observed data are generated as

$$X_{ij} = \left((V_X^{\text{pert}} + \xi_X) A_X + b_X \right)_{ij}, \quad Y_{ij} = \left((V_Y^{\text{pert}} + \xi_Y) A_Y + b_Y \right)_{ij},$$

1024 where \odot denotes element-wise multiplication.

1025 This model was implemented and simulations were generated using `Pyro`.

1026 D PERTURB-MULTIOME PRE-PROCESSING

1028 We downloaded the Perturb-Multiome data from (Martin-Rufino et al., 2025). We focused on the data
1029 at day 14 which showed the least cell type heterogeneity. To reduce the dimensionality and sparsity of
1030 the ATAC-seq data, we filtered ATAC-seq peaks to those measured in at least 1% of cells, and mapped
1031 peaks to the closest gene within with 100kb using the Gencode M38 annotation file. We, used the
1032 inverse document frequency and SVD to reduce the ATACseq to 256 components. For the RNA-seq
1033 data, we normalized the counts to 10,000 per cell and applied the log1p transformation. We filtered
1034 genes that are expressed in fewer than 1% of cells. We then subset to the top 512 highly variable
1035 genes using the Seurat approach with the scanpy implementation. We subset each perturbation to 128
1036 cells per perturbation and split the data into RNA-seq and ATAC-seq for cross modal representation
1037 and prediction. We tested approaches accuracy at leveraging the ATAC-seq data to impute into the
1038 RNA-seq data.

1039 E PERTURB-CITE-SEQ PRE-PROCESSING

1040 We downloaded the Perturb-CITE-seq data from Frangieh et al. (Frangieh et al., 2021). We focused
1041 on the IFNy condition and subset the data to that condition. The data was split into RNA-seq and
1042 CITE-seq for separate preprocessing. For RNA-seq, we normalized the total counts per cell to 10,000
1043 and applied the log1p transformation. For both modalities we subtracted out the average expression
1044 of cells with control perturbations so all values in the gene or protein expression matrix are relative
1045 changes to the average control. To reduce the number of perturbations, we computed the energy
1046 distance for each target gene against the controls and subset to perturbations with at least 50 cells
1047 and an energy distance of 0.05. This resulted in 18 perturbations that were used for downstream
1048 tasks. Finally, we subset the gene expression to the top 500 highly variable genes using the Seurat
1049 method implemented in scanpy. We then tested approaches accuracy at leveraging the CITE-seq data
1050 to impute into the RNA-seq data.

1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

1080 **F METRICS EXTENDED DETAILS**
 1081

1082 We evaluate cross-modal matching and prediction performance using eight complementary metrics
 1083 that capture different aspects of alignment quality and distributional similarity.

1084 **Trace Metric** Assuming the sample indices correspond to the true matching, we compute the average
 1085 weight on correct matches, which is the normalized trace of the transport plan \mathbf{T} :

$$1087 \text{Trace}(\mathbf{T}) = \frac{1}{n} \text{Tr}(\mathbf{T}) = \frac{1}{n} \sum_{i=1}^n \mathbf{T}_{ii} \quad (9)$$

1090 The transport plan \mathbf{T} is first row-normalized such that $\sum_j \mathbf{T}_{ij} = 1$ for all i . A uniformly random
 1091 matching assigns $\mathbf{T}_{ij} = 1/n$ for each cell, yielding $\text{Trace}(\mathbf{T}) = 1/n$. Perfect matching yields
 1092 $\text{Trace}(\mathbf{T}) = 1$.

1093 **Barycentric FOSCTTM** We compute the Fraction Of Samples Closer Than the True Match using
 1094 barycentric projection. Given matching matrix \mathbf{T} and target data $\mathbf{X}^{(1)}$, we project to obtain $\hat{\mathbf{X}}^{(1)} =$
 1095 $\mathbf{T}\mathbf{X}^{(1)}$. For each projected sample $\hat{\mathbf{x}}_i^{(1)}$, we compute the Euclidean distance to all samples in $\mathbf{X}^{(1)}$
 1096 and calculate the fraction of samples closer than the true match:

$$1098 \text{FOSCTTM}(\mathbf{T}, \mathbf{X}^{(1)}) = \frac{1}{n} \sum_{i=1}^n \frac{1}{n-1} \sum_{j \neq i} \mathbf{1}\{d(\hat{\mathbf{x}}_i^{(1)}, \mathbf{x}_j^{(1)}) < d(\hat{\mathbf{x}}_i^{(1)}, \mathbf{x}_i^{(1)})\}, \quad (10)$$

1101 where $d(\cdot, \cdot)$ denotes Euclidean distance. The final reported symmetric Barycentric FOSCTTM is an
 1102 average over both both modalities: $0.5 \times (\text{FOSCTTM}(\mathbf{T}, \mathbf{X}^{(1)}) + \text{FOSCTTM}(\mathbf{T}^\top, \mathbf{X}^{(2)}))$

1104 Lower values indicate better matching quality, with random matching expected to yield 0.5.

1105 **Mean Squared Error (MSE)** For direct prediction evaluation, we compute the MSE between true
 1106 samples $\mathbf{X}^{(1)}$ and predicted samples $\hat{\mathbf{X}}^{(1)}$:

$$1108 \text{MSE} = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i^{(1)} - \hat{\mathbf{x}}_i^{(1)}\|_2^2. \quad (11)$$

1110 We report the mean MSE across all features.

1112 **1-Wasserstein Distance (WD)** To assess distributional similarity, we compute the 1-Wasserstein
 1113 distance between true and predicted samples, averaged across features:

$$1115 \text{WD} = \frac{1}{d} \sum_{j=1}^d W_1(\mathbf{X}_{:,j}^{(1)}, \hat{\mathbf{X}}_{:,j}^{(1)}), \quad (12)$$

1117 where W_1 denotes the 1-Wasserstein distance between univariate distributions and $\mathbf{X}_{:,j}$ represents
 1118 the j -th feature column.

1120 **Cosine Similarity** We compute the average cosine similarity between corresponding true and
 1121 predicted feature vectors:

$$1123 \text{Cosine} = \frac{1}{d} \sum_{j=1}^d \frac{\mathbf{X}_{:,j}^{(1)} \cdot \hat{\mathbf{X}}_{:,j}^{(1)}}{\|\mathbf{X}_{:,j}^{(1)}\|_2 \|\hat{\mathbf{X}}_{:,j}^{(1)}\|_2}, \quad (13)$$

1125 where the dot product and norms are computed across samples for each feature.

1126 **KNN-based Metrics** To evaluate neighborhood preservation, we construct k -nearest neighbor graphs
 1127 for both true and predicted data using cosine similarity. Let \mathbf{G}^{true} and \mathbf{G}^{pred} denote the binary
 1128 adjacency matrices of the respective KNN graphs (with $k = 10$). For each sample i , we treat $\mathbf{G}_{i,:}^{\text{true}}$ as
 1129 ground truth labels and $\mathbf{G}_{i,:}^{\text{pred}}$ as predictions, then compute:

1131 **KNN Recall:** The fraction of true neighbors correctly identified:

$$1132 \text{KNN Recall} = \frac{1}{n} \sum_{i=1}^n \frac{\sum_j \mathbf{G}_{i,j}^{\text{true}} \mathbf{G}_{i,j}^{\text{pred}}}{\sum_j \mathbf{G}_{i,j}^{\text{true}}}. \quad (14)$$

1134
 1135 **KNN Average Precision (KNN PR):** The average precision score for each sample's neighborhood
 1136 prediction:

1137
$$\text{KNN PR} = \frac{1}{n} \sum_{i=1}^n \text{AP}(\mathbf{G}_{i,:}^{\text{true}}, \mathbf{G}_{i,:}^{\text{pred}}), \quad (15)$$

 1138

1139 where AP denotes the average precision score.

1140 **KNN ROC-AUC (KNN ROC):** The area under the ROC curve for neighborhood prediction:

1141
 1142
$$\text{KNN ROC} = \frac{1}{n} \sum_{i=1}^n \text{AUC}(\mathbf{G}_{i,:}^{\text{true}}, \mathbf{G}_{i,:}^{\text{pred}}), \quad (16)$$

 1143
 1144

1145 where AUC denotes the area under the receiver operating characteristic curve.

1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

1188 G SIMULATIONS EXTENDED RESULTS
1189
1190
11911192 **Table 6:** Statistical significance tests for cross-modal matching metrics comparing top-ranked
1193 GROOVE (Target) versus top-ranked non-GROOVE (Baseline) methods. One-sided paired t-tests
1194 and win rates computed across replicates, with test direction towards metric improvement. Δ = Target
1195 - Baseline mean difference; n.s. = $p > 0.1$.

1196 Shared 1197 Prop.	1198 Metric	1199 Target Method	1200 Baseline Method	1201 Avg. Δ	1202 Win Rate (%)	1203 p-value
1200 100%	Bary. FOSCTTM Trace	GROOVE (cosine)+LabeledCOOT GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledCOOT DAVAE+LabeledCOOT	-0.039 0.186	100 100	8.50e-05 6.54e-06
1201 80%	Bary. FOSCTTM Trace	GROOVE (cosine)+LabeledCOOT GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledEOT DAVAE+LabeledEOT	0.034 0.072	0 100	n.s. 1.29e-04
1203 50%	Bary. FOSCTTM Trace	GROOVE (cosine)+LabeledEOT GROOVE (cosine)+LabeledEOT	DAVAE+LabeledEOT DAVAE+LabeledEOT	0.000 -0.007	60 30	n.s. n.s.

1206
1207 **Table 7:** Statistical significance tests for cross-modal imputation metrics comparing top-ranked
1208 GROOVE (Target) versus top-ranked non-GROOVE (Baseline) methods. One-sided paired t-tests
1209 and win rates computed across replicates, with test direction towards metric improvement. Δ = Target
1210 - Baseline mean difference; n.s. = $p > 0.1$.

1212 Shared 1213 Prop.	1214 Metric	1215 Target Method	1216 Baseline Method	1217 Avg. Δ	1218 Win Rate (%)	1219 p-value
1215 100%	Cos-sim	GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledCOOT	0.018	100	4.61e-05
	KNN PR	GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledCOOT	0.028	100	1.11e-03
	KNN ROC	GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledCOOT	0.017	100	2.17e-04
	KNN Recall	GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledCOOT	0.033	100	2.17e-04
	MSE	GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledCOOT	-0.022	100	2.46e-06
	WD	GROOVE (cosine)+LabeledCOOT	DAVAE+LabeledCOOT	-0.027	100	3.02e-06
1221 80%	Cos-sim	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.001	50	n.s.
	KNN PR	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.002	50	n.s.
	KNN ROC	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.001	60	n.s.
	KNN Recall	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.002	60	n.s.
	MSE	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.030	10	n.s.
	WD	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.001	40	n.s.
1227 50%	Cos-sim	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.005	90	1.35e-02
	KNN PR	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	-0.001	50	n.s.
	KNN ROC	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	-0.002	30	n.s.
	KNN Recall	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	-0.003	30	n.s.
	MSE	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	-0.003	50	n.s.
	WD	GROOVE (cosine)+LabeledEOT	PS+LabeledEOT	0.011	20	n.s.

1233
1234
1235
1236
1237
1238
1239
1240
1241

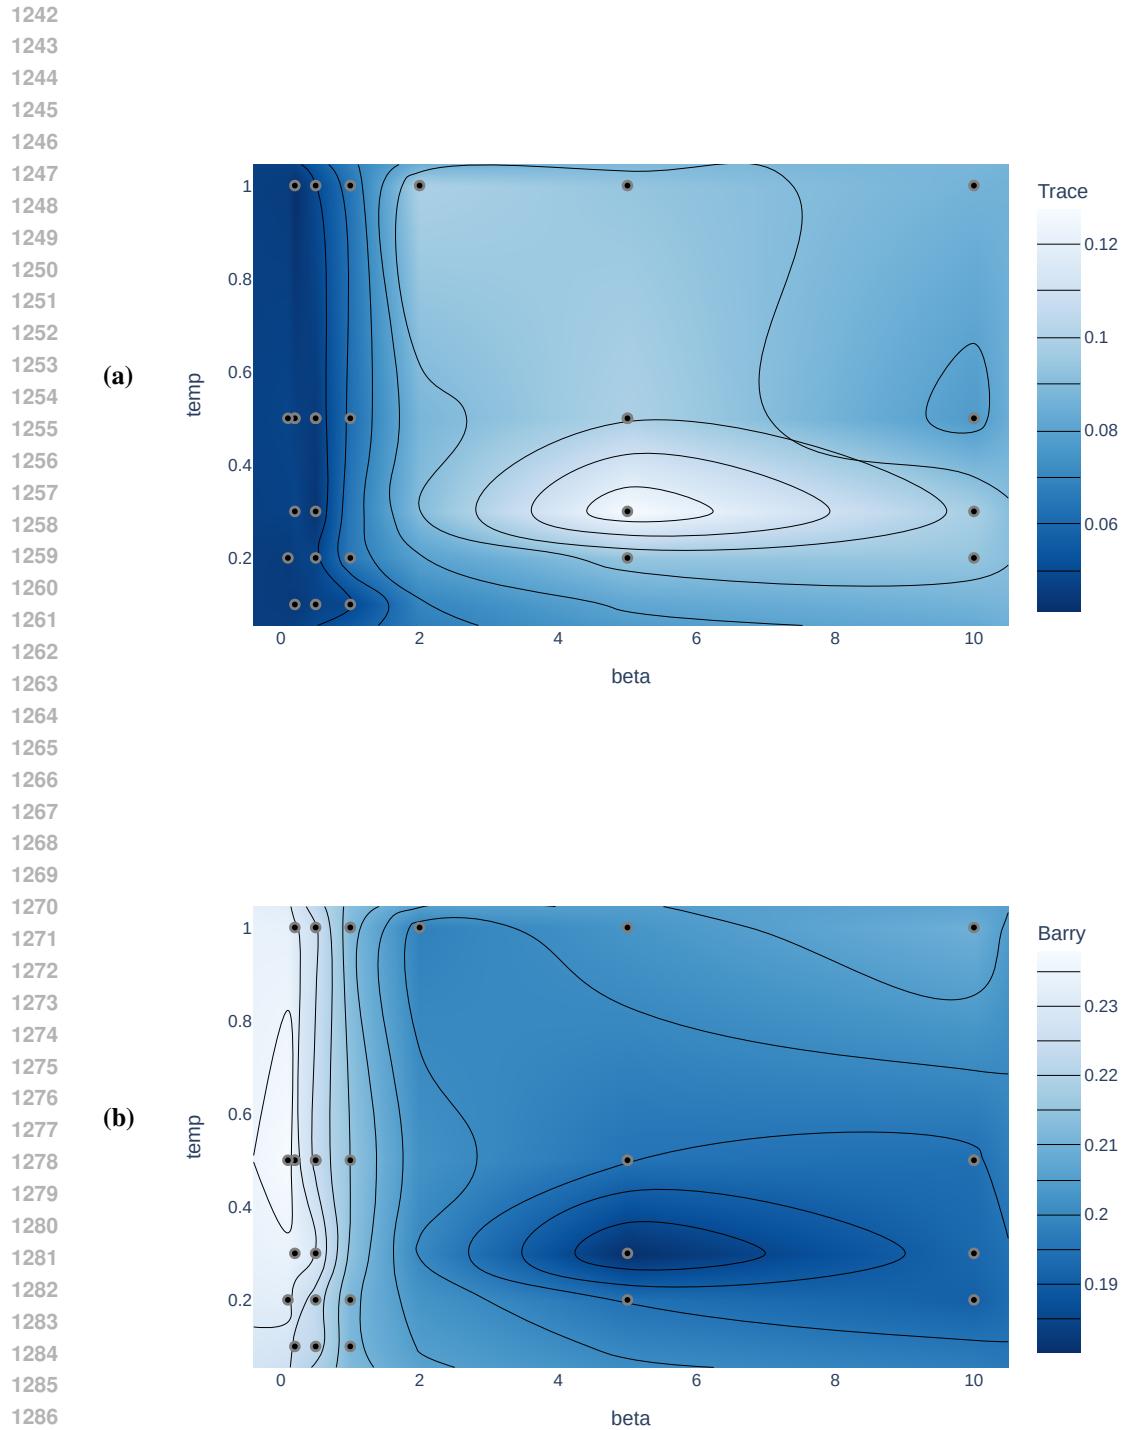


Figure 2: Hyperparameter sensitivity landscape for matching performance. Contour plots show average performance across 100%, 80%, and 50% shared variation settings for each combination of β (x-axis) and τ (y-axis), profiled using Optuna-based hyperparameter search. $\alpha = 1$ for all analysis in this work. (a) Trace-based matching performance (higher is better). (b) Barycentric FOSCTTM (lower is better).

1296 **H PERTURB-CITE-SEQ EXTENDED RESULTS**
 1297
 1298
 1299
 1300

1301 Table 8: Matching performance metrics for top 10 method combinations in Perturb-CITE-seq dataset
 1302 with 5-fold evaluation. SEs follow \pm ; best in bold, second-best underlined.
 1303

Method	Mean Rank	Trace	Bary. FOSCTTM
GROOVE (cosine)+ LabeledEOT	3.0	<u>0.039</u> \pm 0.002	0.381 \pm 0.004
PS+ LabeledEOT	3.0	<u>0.039</u> \pm 0.002	<u>0.383</u> \pm 0.002
GROOVE (tdist)+ LabeledEGWOT	3.5	0.033 \pm 0.001	0.365 \pm 0.001
GROOVE (cosine)+ LabeledCOOT	5.0	0.040 \pm 0.003	0.451 \pm 0.004
GROOVE (cosine)+ LabeledEGWOT	5.0	0.031 \pm 0.001	0.369 \pm 0.001
GROOVE (tdist)+ LabeledEOT	5.5	0.036 \pm 0.001	0.394 \pm 0.001
DAVAE+ LabeledEGWOT	7.5	0.024 \pm 0.000	<u>0.383</u> \pm 0.003
DAVAE+ LabeledEOT	7.5	0.032 \pm 0.001	<u>0.397</u> \pm 0.011
GROOVE (tdist)+ LabeledCOOT	7.5	0.036 \pm 0.002	0.452 \pm 0.004
PS+ LabeledEGWOT	9.0	0.023 \pm 0.000	0.389 \pm 0.001

1319 Table 9: Matching performance metrics for top 10 method combinations in Perturb-CITE-seq dataset
 1320 with leave one perturbation out evaluation. SEs follow \pm ; best in bold.
 1321

Method	Mean Rank	Trace	Bary. FOSCTTM
PS+ LabeledEOT	2.5	0.008 \pm 0.001	0.487 \pm 0.004
GROOVE (tdist)+ LabeledEOT	4.5	0.008 \pm 0.002	0.490 \pm 0.003
GROOVE (cosine)+ EGW	6.0	0.008 \pm 0.001	0.491 \pm 0.003
GROOVE (tdist)+ EOT	6.0	0.007 \pm 0.001	0.485 \pm 0.003
GROOVE (cosine)+ EOT	6.5	0.007 \pm 0.001	0.489 \pm 0.003
GROOVE (cosine)+ LabeledCOOT	6.5	0.010 \pm 0.002	0.495 \pm 0.003
PS+ EOT	6.5	0.007 \pm 0.001	0.488 \pm 0.003
GROOVE (cosine)+ LabeledEOT	7.0	0.007 \pm 0.001	0.487 \pm 0.003
DAVAE+ EOT	7.0	0.008 \pm 0.002	0.491 \pm 0.003
DAVAE+ LabeledCOOT	7.5	0.008 \pm 0.002	0.495 \pm 0.002

1338 Table 10: Imputation performance metrics for top 10 method combinations in Perturb-CITE-seq
 1339 dataset with leave one perturbation out evaluation. SEs follow \pm ; best in bold, second-best underlined,
 1340 homogeneous metrics unannotated.

Method	Mean Rank	MSE	Cos-sim	KNN Recall	KNN PR	KNN ROC	WD
GROOVE (cosine)+ LabeledEOT	7.00	0.262 \pm 0.001	0.057 \pm 0.007	0.075 \pm 0.008	0.072 \pm 0.008	0.505 \pm 0.001	0.352 \pm 0.001
GROOVE (tdist)+ LabeledEGWOT	7.33	0.262 \pm 0.001	0.065 \pm 0.008	0.074 \pm 0.008	0.072 \pm 0.008	0.505 \pm 0.000	0.355 \pm 0.001
DAVAE+ LabeledEOT	7.83	0.262 \pm 0.001	0.082 \pm 0.008	0.074 \pm 0.008	0.073 \pm 0.008	0.504 \pm 0.001	0.354 \pm 0.000
GROOVE (tdist)+ LabeledCOOT	8.00	0.282 \pm 0.001	0.021 \pm 0.003	0.077 \pm 0.008	0.072 \pm 0.008	0.506 \pm 0.001	0.300 \pm 0.001
PS+ LabeledEOT	8.17	0.263 \pm 0.001	0.070 \pm 0.008	0.074 \pm 0.008	0.072 \pm 0.008	0.505 \pm 0.001	0.348 \pm 0.001
GROOVE (cosine)+ EGW	8.67	0.262 \pm 0.001	0.085 \pm 0.010	0.074 \pm 0.008	0.072 \pm 0.008	0.504 \pm 0.001	0.363 \pm 0.000
DAVAE+ LabeledCOOT	9.17	0.290 \pm 0.002	<u>0.031</u> \pm <u>0.003</u>	0.074 \pm 0.007	0.073 \pm 0.008	0.504 \pm 0.001	0.289 \pm 0.001
GROOVE (cosine)+ LabeledEGWOT	9.67	0.262 \pm 0.001	0.066 \pm 0.008	0.073 \pm 0.008	0.072 \pm 0.008	0.504 \pm 0.001	0.355 \pm 0.000
GROOVE (tdist)+ LabeledEOT	9.67	0.262 \pm 0.001	0.054 \pm 0.007	0.074 \pm 0.008	0.072 \pm 0.008	0.505 \pm 0.001	0.352 \pm 0.001
DAVAE+ LabeledEGWOT	10.33	0.262 \pm 0.001	0.088 \pm 0.009	0.071 \pm 0.008	0.073 \pm 0.008	0.502 \pm 0.001	0.354 \pm 0.001

1350
 1351 **Table 11: GROOVE ablation analysis performance metrics in Perturb-CITE-seq dataset. SEs follow**
 1352 **±; best in bold, second-best underlined, homogeneous metrics unannotated.**

1353 1354 Abation Type	1353 Bary. FOSCTTM	1353 MSE	1353 Cos-sim	1353 KNN Recall	1353 KNN PR	1353 KNN ROC	1353 WD
GROOVE (cosine)	0.368±0.005	0.261±0.001	0.047±0.002	0.019±0.001	0.017±0.000	0.502±0.000	0.353±0.000
No GroupCLIP	0.380±0.003	0.261±0.001	0.048±0.001	0.019±0.000	0.017±0.000	0.502±0.000	0.353±0.001
Autoencoder only	0.381±0.002	0.261±0.001	0.048±0.004	0.018±0.000	0.017±0.000	0.502±0.000	0.353±0.001

1357
 1358
 1359
 1360
I PERTURB-MULTIOME EXTENDED RESULTS
 1361
 1362

1363
 1364
 1365
 1366
 1367
 1368
 1369
Table 12: Imputation performance metrics for top 10 method combinations in Perturb-Multiome
 1370 **dataset with 5-fold evaluation. SEs follow ±; best in bold, second-best underlined.**

1372 1373 Method	1372 Mean Rank	1372 MSE	1372 Cos-sim	1372 KNN Recall	1372 KNN PR	1372 KNN ROC	1372 WD
GROOVE (cosine)+ EOT	4.50	0.308±0.003	0.098±0.028	0.052±0.005	0.029±0.001	0.515±0.003	0.428±0.005
GROOVE (cosine)+ LabeledEOT	5.17	0.311±0.001	<u>0.075±0.004</u>	<u>0.066±0.003</u>	<u>0.032±0.001</u>	<u>0.523±0.002</u>	0.433±0.001
PS+ EOT	6.17	0.308±0.003	<u>0.113±0.029</u>	0.043±0.004	0.027±0.001	0.511±0.002	0.426±0.004
GROOVE (tdist)+ LabeledEOT	6.50	0.311±0.001	0.071±0.002	<u>0.066±0.003</u>	<u>0.032±0.001</u>	<u>0.523±0.002</u>	0.433±0.001
GROOVE (tdist)+ EOT	7.33	0.310±0.002	0.090±0.022	0.047±0.007	0.028±0.002	0.513±0.003	0.430±0.003
GROOVE (cosine)+ LabeledCOOT	7.50	0.355±0.023	0.091±0.054	0.045±0.005	0.028±0.001	0.512±0.002	0.323±0.008
PS+ LabeledEOT	7.50	0.311±0.001	0.070±0.002	0.058±0.003	0.030±0.001	0.518±0.001	0.433±0.001
GROOVE (tdist)+ LabeledCOOT	9.33	0.364±0.029	0.047±0.079	<u>0.048±0.005</u>	0.028±0.001	0.514±0.003	0.317±0.011
PS+ LabeledCOOT	9.50	0.334±0.006	0.041±0.025	0.048±0.005	0.028±0.001	0.514±0.002	0.379±0.001
PS+ LabeledEGWOT	10.00	0.311±0.001	0.075±0.003	0.037±0.006	0.026±0.001	0.508±0.003	0.430±0.001

1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
Table 13: Matching performance metrics for top 10 method combinations in Perturb-Multiome dataset
 1390 **with leave one perturbation out evaluation. SEs follow ±; best in bold, second-best underlined.**

1393 1394 Method	1393 Mean Rank	1393 Trace	1393 Bary. FOSCTTM
GROOVE (tdist)+ LabeledCOOT	1.5	0.014±0.003	<u>0.481±0.019</u>
GROOVE (tdist)+ LabeledEOT	3.0	0.008±0.000	0.480±0.004
GROOVE (cosine)+ EOT	4.5	0.008±0.000	0.487±0.004
DAVAE+ LabeledCOOT	4.5	0.010±0.002	0.488±0.009
GROOVE (cosine)+ LabeledEOT	5.0	0.008±0.000	0.486±0.003
GROOVE (tdist)+ EOT	7.0	0.008±0.000	0.487±0.003
GROOVE (tdist)+ LabeledEGWOT	7.0	0.008±0.000	0.496±0.005
GROOVE (cosine)+ LabeledCOOT	7.5	<u>0.011±0.003</u>	0.500±0.018
PS+ EOT	9.5	0.008±0.000	0.497±0.002
DAVAE+ LabeledEOT	10.5	0.008±0.000	0.494±0.002

1404

1405 Table 14: Matching performance metrics for top 10 method combinations in Perturb-Multiome dataset
1406 with 5-fold evaluation. SEs follow \pm ; best in bold, second-best underlined.

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

J EXTENDED DISCUSSION

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

J.1 SYNERGY BETWEEN GROUPCLIP AND LABELED-CONSTRAINED OT

While GroupCLIP already leverages perturbation labels during representation learning, label-constrained OT methods provide complementary benefits at the alignment stage. Specifically, GroupCLIP operates at the group level during training, encouraging same-label samples to cluster together across modalities through *soft constraints*. Now consider the fact that perturbation effects are not always orthogonal. Under this common setting, the latent representations of cells with similar perturbation effects will be closer to each other (or more correlated) in the latent space. This is indeed the type of behavior one would desire from a useful/informative latent. Label-constrained OT on the otherhand sets a *hard constraint* across groups (see Ryu et al. (2025)). That is, it explicitly prevents cross-label matches while finding optimal pairings within each label group. This is a useful constraint when the objective is to match an individual cell within a perturbation. The empirical benefits of this synergy between our soft-constraint and Labelled OT’s hard constraint can be quite strong: in Table 1 (100% shared), GROOVE with LabeledCOOT achieves Trace=0.856 versus markedly worse performance with unlabeled OT variants not even being in the top-5 (trace metrics strictly less than 5th highest combination).

J.2 MATCHING VERSUS IMPUTATION TRADEOFF

Our empirical results reveal an important observation: superior matching performance does not necessarily translate to improved downstream imputation (Section 5). This phenomenon may reflect a tension between these two objectives that warrants further consideration.

Xi et al. (2024) argue that reconstruction losses force models to learn modality-specific noise, which is “counterproductive to matching.” This claim rests on the assumption that all meaningful perturbation-induced variation is perfectly shared across modalities. However, biological reality is often more complex: perturbation effects frequently manifest differently across modalities, with some responses observable in only one modality (private or modality-specific variation) (Argelaguet et al., 2020; Lin & Zhang, 2023).

This can induce a tradeoff. For optimal matching, representations *can* capture only shared perturbation-relevant variation while discarding modality-specific information as “noise.” Since in this task success is defined by recovering true instance-level pairs, i.e., the “real” co-measured cells. However, for imputation and prediction, the most useful training/inference samples are not necessarily the true paired cells, but rather the empirically most similar cells in the biologically relevant latent space (which might also factor in private variation). Fundamentally, perturbations induce similarity structures that supersedes instance-level pairing. Two unpaired cells subjected to the same perturbation

Method	Mean Rank	Trace	Bary. FOSCTTM
GROOVE (cosine)+ LabeledCOOT	2.0	0.048 ± 0.009	0.431 ± 0.028
GROOVE (cosine)+ LabeledEOT	2.5	0.041 ± 0.001	$\underline{\mathbf{0.427}} \pm \mathbf{0.008}$
GROOVE (tdist)+ LabeledEOT	4.5	0.040 ± 0.000	0.441 ± 0.003
DAVAE+ LabeledCOOT	6.0	$\mathbf{0.052} \pm \mathbf{0.010}$	0.470 ± 0.025
GROOVE (tdist)+ LabeledEGWOT	6.0	0.040 ± 0.000	0.448 ± 0.003
GROOVE (cosine)+ LabeledEGWOT	6.5	0.039 ± 0.001	0.446 ± 0.003
GROOVE (tdist)+ LabeledCOOT	6.5	0.047 ± 0.017	0.457 ± 0.050
PS+ LabeledEOT	7.0	0.039 ± 0.000	0.445 ± 0.003
DAVAE+ LabeledEGWOT	8.5	0.039 ± 0.001	0.448 ± 0.003
PS+ LabeledEGWOT	8.5	0.039 ± 0.001	0.447 ± 0.003

1458 may exhibit greater functional similarity, and thus provide more informative training signal for
1459 imputation than a cell’s true paired counterpart, particularly when factoring weak perturbation effects
1460 and technical measurement noise. A purely matching optimized representation, as Xi et al. (2024)
1461 does, discard this type of variation or stucture which would be helpful for downstream imputation.

1462 GROOVE navigates this tradeoff through the inclusion of the reconstruction objective(s). The
1463 reconstruction and backtranslation losses (Equation 5) preserve sample-specific information, including
1464 modality-specific variation. While this may reduce pure matching performance by retaining what
1465 matching frameworks consider “noise,” it provides empirically beneficial information for downstream
1466 imputation tasks. This explains why combinations with superior matching scores do not always
1467 achieve the best imputation performance (like in Tables 1 and 2). The consistency of this matching-
1468 imputation discordance also is need in real real datasets (Sections 5.1 and 5.2) suggests this represents
1469 a genuine phenomenon rather than an artifact of a particular dataset or experimental design. We
1470 believe this tradeoff warrants further theoretical investigation and empirical characterization, which
1471 we defer to future work.

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511