

# AVIDA: Alternating method for Visualizing and Integrating Data

Kathryn Dover<sup>1</sup>, Zixuan Cang<sup>2</sup>, Anna Ma<sup>1</sup>, Qing Nie<sup>1,3</sup>, and Roman Vershynin<sup>1</sup>

<sup>1</sup>Department of Mathematics, University of California, Irvine

<sup>2</sup>Department of Mathematics, North Carolina State University

<sup>3</sup>Department of Developmental & Cell Biology, Center for Multiscale Cell Fate Research, University of California, Irvine

## Abstract

High-dimensional multimodal data arises in many scientific fields. The integration of multimodal data becomes challenging when there is no known correspondence between the samples and the features of different datasets. To tackle this challenge, we introduce *AVIDA*, a framework for simultaneously performing data alignment and dimension reduction. In the numerical experiments, Gromov-Wasserstein optimal transport and t-distributed stochastic neighbor embedding are used as the alignment and dimension reduction modules respectively. We show that AVIDA correctly aligns high-dimensional datasets without common features with four synthesized datasets and two real multimodal single-cell datasets. Compared to several existing methods, we demonstrate that AVIDA better preserves structures of individual datasets, especially distinct local structures in the joint low-dimensional visualization, while achieving comparable alignment performance. Such a property is important in multimodal single-cell data analysis as some biological processes are uniquely captured by one of the datasets. In general applications, other methods can be used for the alignment and dimension reduction modules.

## 1 Introduction

Databases are expanding not only in size but also with increasing complexity. In many applications, multiple measurements of a system are taken across different samples or in different feature spaces which produce *multimodal data* such as texts attached to images [18]. Multimodality allows more comprehensive investigation of a system. Establishing connections among the modalities is the foundation of coherent analysis. Recently, the emerging multimodal single-cell omics has become a powerful tool to analyze different aspects of a biological system at the

---

Corresponding authors emails: qnie@uci.edu; rvershyn@uci.edu

same time [43]. Fusing multimodal single-cell data is especially challenging when there are no direct correspondence between the measurements and the samples.

Single-cell RNA sequencing (scRNA-seq) is a recent technology that measures RNA abundance at transcriptomics level with single-cell resolution [34]. The maturation of the technology allows analysis with scRNA-seq assays across many samples that, for example, represent different ages or healthy and diseased individuals [33, 39]. On the other hand, the emerging single-cell assays provide more comprehensive examination of a system, such as single-cell ATAC-seq (scATAC-seq) [4] that measures chromatin accessibility and single-cell Hi-C [25] that explores chromosome architecture.

Integrating the various single-cell assays across different samples provides a comprehensive characterization of a biological system. Many computational methods have been developed to integrate the same single-cell assays of multiple samples [14, 32, 40] or different single-cell assays [12, 13]. In the integration of multiple single-cell omics assays, most current methods rely on the known correspondence between features, for example by mapping chromatin loci to genes and assuming the similarity between the samples. The multi-omics integration becomes a harder problem when no prior correspondence is assumed, for example a gene actually corresponds to multiple loci and accessible loci does not directly indicate gene expression. This leads to a general problem of integrating datasets without known correspondence between features.

When no feature correspondence is given, the structures of the individual datasets can be exploited and matched to integrate the datasets. For example, canonical correlation analysis examines covariances between the datasets but is limited to deriving linear correspondence between the features. When the datasets are represented as graphs with edges annotating pairs of similar data points within each dataset, the integration problem can be addressed using various graph alignment methods [27, 42]. Among the graph alignment methods, *Gromov-Wasserstein optimal transport (GW-OT)* can align graphs based only on the graph structures [23]. It finds a coupling of the distributions representing the graphs that best preserves the intra-dataset distances between the nodes.

Optimal transport (OT) compares and finds connections between measures. It seeks the coupling between distributions with the minimum total coupling cost based on predefined costs between locations [17, 24, 38]. OT has been a versatile tool widely used in practical problems, such as generative deep learning models [3], domain adaptation [8], and image sciences [10]. It has been used to find correspondence between data points in single-cell gene expression data with common features [5, 26, 31]. The aforementioned GW-OT has been used in this field to exploit the structural information within individual datasets. SpaOTsc [5] uses fused Wasserstein-Gromov-Wasserstein optimal transport to improve the integration of spatial data and scRNA-seq data with few shared genes by matching the spatial structure and the structure in scRNA-seq data based on gene expression similarity. SCOT [9] uses Gromov-Wasserstein optimal transport to align scRNA-seq and scATAC-seq data by matching the structures represented by intra-dataset similarity among cells. Pamona [6] uses partial Gromov-Wasserstein optimal transport to partially align scRNA-seq and scATAC-seq data to address the partially overlapping cell populations among different samples.

In addition to studying shared structures revealed by the overlapping part of integrated data, it is equivalently important to examine the structures of non-overlapping part which may depict a biological process uniquely captured by a certain as-

say [41]. Since most integration methods depend on similarities between samples, the dissimilar parts are often overlooked. Efforts have been made to keep the variation among samples examined with the same single-cell assay [41].

In the analysis of high-dimensional multimodal datasets, another crucial step is *dimensionality reduction*. Dimensionality reduction is the process of taking high-dimensional data and finding a representation in lower dimensions that is still meaningful. It has many important applications because dimensionality reduction helps address the curse of dimensionality and other challenges that come with working with high-dimensional data [16]. Principal Component Analysis (PCA) [28] is the most traditional linear technique used in dimensionality reduction but there are many popular non-linear techniques, such as Local Linear Embedding [30], Isomap [35], UMAP [22], and t-SNE [36].

t-SNE is a popular dimensionality reduction and visualization technique that was introduced in 2008 by van der Maaten and Hinton [36]. It has been applied towards a variety of high-dimensional data, including deep learning [21], physics [37], and medicine [1]. Given a high-dimensional dataset, t-SNE outputs a low-dimensional representation. t-SNE works by making pairwise affinities between points in high-dimensions and pairwise affinities between points in low-dimensions. It then uses gradient descent to find the set of points (in low dimensions) that minimize the KL divergence between the two sets of joint probabilities.

In the analysis of multimodal single-cell data, the dimensionality reduction and the integration steps are often performed separately or sequentially, including the existing methods that integrate datasets without known feature correspondences [6,9]. However, these two steps are closely related in that they both preserve the structures from high dimension to low dimension or from the original spaces to the joint space. The benefit of combining these two steps has been shown in a recent work that uses shared features between datasets [15]. In this work, we present a workflow called AVIDA (alternating method for visualizing and integrating data) that alternatively performs dimension reduction and data integration for integrating datasets without common features such that they regulate each other. To demonstrate this workflow, we use t-SNE for the dimension reduction module and Gromov-Wasserstein optimal transport for the integration module. In four synthetic datasets and two real biological datasets with ground truth, we show that AVIDA better preserves the structures of the individual datasets while achieving comparable integration quality compared to other methods.

## 2 Results

### 2.1 Overview of AVIDA

The proposed method is called the *alternating method for visualizing and integrating data*, or AVIDA. AVIDA alternates between improving the low dimensional representation through a dimensionality reduction technique and the alignment of data points in low dimensions across different datasets. The purpose of alternating between dimensionality reduction and alignment is to find a balance between a good representation while still accurately aligning the datasets. We denote AVIDA as a function, taking as input the datasets  $X_1, \dots, X_k$  and is parameterized by the choice of dimensionality reduction and alignment techniques:  $\text{AVIDA}(X_1, X_2, \dots, X_k; \text{DR}, \text{ALIGN})$ . A simplified schematic of the method is shown




Figure 1: A visual schematic of AVIDA.

in Figure 1. As shown in Figure 1, AVIDA can take as input two datasets and organizes the data as a pairwise distance matrix. Next, dimensionality reduction using the given pairwise distance matrix is performed on both datasets independently. An alignment method is used to “align” the datasets in the lower dimensional space and using the aligned data points, a new pairwise distance matrix is formed for each dataset and the process iterates. This framework is flexible in its choice of dimensionality reduction technique (in fact, different dimension reduction algorithms can be used on different datasets if one so chooses) and alignment method.

Suppose one is given two datasets  $X^{(1)}$  and  $X^{(2)}$  and the goal is to create a joint representation of the datasets in a common lower dimensional space. Using some technique DR for dimensionality reduction (e.g., PCA, t-SNE, Random Forests, etc.) and GW-OT for alignment, we can formulate the objective function for AVIDA as  $\text{AVIDA}(X_1, X_2; \text{DR}, \text{GW})$ . The GW-OT objective is defined with respect to the low-dimensional representation of points:

$$\text{GW}(Y^{(1)}, Y^{(2)}) = \sum_{i, j, i', j'} L_{i, j, i', j'} \mathbf{T}_{i, i'} \mathbf{T}_{j, j'} - \epsilon(H(\mathbf{T})), \quad (1)$$

where  $H(\mathbf{T}) = \sum_{i, j} T_{ij} \log(T_{ij})$  is the Entropic regularization term and

$L_{i, j, i', j'} = \|d(y_i^{(1)}, y_j^{(1)}) - d(y_{i'}^{(2)}, y_{j'}^{(2)})\|^2$  with a chosen distance metric  $d(\cdot, \cdot)$ . This objective is minimized by using the projected gradient descent method with KL metric based projections [29],  $T \leftarrow \text{Proj}_{U(\mathbf{a}, \mathbf{b})}^{\text{KL}}(T \odot e^{-\tau(L \otimes T + \epsilon \log(T))})$  where

$U(\mathbf{a}, \mathbf{b}) = \{T \in \mathbb{R}_+^{n_1 \times n_2} : T \mathbf{1} = \mathbf{a}, T^T \mathbf{1} = \mathbf{b}\}$  and  $\tau$  is the step size. The implementation in Python Optimal Transport [11] package is used. The representation for  $Y^{(1)}$  will subsequently be mapped to  $Y^{(2)}$  using the mapping found by minimizing (1) with respect to  $T$ , i.e., by setting  $Y^{(1)} = TY^{(2)}$ . Our combined loss function can be represented as

$$\text{AVIDA}(X^{(1)}, X^{(2)}; \text{DR}, \text{GW}) = \min_{Y^{(1)}, Y^{(2)}} \text{DR}(X^{(1)}, Y^{(1)}) + \text{DR}(X^{(2)}, Y^{(2)}) + \text{GW}(Y^{(1)}, Y^{(2)}), \quad (2)$$

where  $DR(X^{(i)}, Y^{(i)})$  represents the objective loss associate with the dimensionality reduction technique DR. For example, if t-SNE is uses for the DR step, the objective can be represented as the KL loss between probability distributions on the points in high and low dimensions. See 4 for more details.

## 2.2 AVIDA accurately reproduces the intra-dataset structures in integration of synthetic data

We compared AVIDA( $X_1, X_2$ ; TSNE, GW) to both Pamona and SCOT across four simulated datasets and two real-world single-cell multi-omics datasets. We chose Pamona and SCOT as a comparison because they are both advanced integration methods. Table 1 contains the metrics for AVIDA( $X_1, X_2$ ; TSNE, GW), SCOT and Pamona on both the simulated and real-life datasets. We used five different metrics to assess the performance of these methods: the fraction of samples closer than the true match (FOSCTTM), alignment, integration, accuracy and representation loss. The accuracy metric is only included on the datasets where the ground truth is known and an empty cell in the table implies the dataset did not meet that requirement. Details on the metrics are included in Section 4.2.

| Dataset          | Method | FOSCTTM       | Integration       | Accuracy          | Alignment     | Representation Loss |
|------------------|--------|---------------|-------------------|-------------------|---------------|---------------------|
| Bifurcated Tree  | AVIDA  | 0.1202        | 1.0820            | 4.3863            | <b>0.5157</b> | <b>0.3275</b>       |
|                  | Pamona | <b>0.1001</b> | <b>0.0004</b>     | <b>0.008</b>      | 0.9817        | 2.2340              |
|                  | SCOT   | 0.2103        | 1.0016            | 12.2095           | 0.75          | 2.1466              |
| Circular Frustum | AVIDA  | 0.1187        | 0.9699            | 2.9377            | <b>0.4267</b> | <b>0.3955</b>       |
|                  | Pamona | <b>0.0184</b> | <b>0.0003</b>     | <b>0.0012</b>     | 0.9433        | 2.2311              |
|                  | SCOT   | 0.0515        | 1.0032            | 4.3857            | 0.9727        | 1.7083              |
| Dumbbell         | AVIDA  | 0.5228        | 0.5568            | 25.1281           | 0.6385        | <b>0.1220</b>       |
|                  | Pamona | 0.5238        | <b>0.0010</b>     | <b>0.0731</b>     | 0.76          | 1.8240              |
|                  | SCOT   | <b>0.4754</b> | 2.565             | 11.2244           | <b>0.2070</b> | 3.6008              |
| Distant Rings    | AVIDA  | 0.3138        | 0.6847            | 5.3429            | <b>0.639</b>  | <b>0.1916</b>       |
|                  | Pamona | 0.1770        | <b>6.8279e-17</b> | <b>2.0019e-15</b> | 1.0           | 1.8662              |
|                  | SCOT   | <b>0.0056</b> | 0.0791            | 0.2759            | 0.9125        | 0.9261              |
| sc-GEM           | AVIDA  | 0.2070        | 0.4700            | 2.4996            | 0.8994        | <b>0.4879</b>       |
|                  | Pamona | 0.2084        | <b>0.0005</b>     | <b>0.0106</b>     | 0.7164        | 1.6516              |
|                  | SCOT   | <b>0.1818</b> | 2.3164            | 6.9267            | <b>0.5616</b> | 0.8790              |
| scNMT-seq        | AVIDA  | 0.2745        | 0.3631            | 4.5787            | <b>0.6619</b> | <b>1.0489</b>       |
|                  | Pamona | 0.2908        | <b>5.337e-05</b>  | <b>0.0060</b>     | 0.9785        | 2.8408              |
|                  | SCOT   | <b>0.2675</b> | 2.4333            | 28.6287           | 0.7522        | 1.1979              |

Table 1: Metrics for AVIDA( $X_1, X_2$ ; TSNE, GW) (labeled as AVIDA above), Pamona and SCOT experiments.

Our four simulated datasets include a bifurcated tree, a circular frustum (from [20]), a dumbbell and distant rings. The dumbbell and distant rings datasets are introduced in order to highlight the difference between AVIDA and SCOT and Pamona. The dumbbell dataset consists of one dataset with two rings connected by a line and the other containing just a line. The rings dataset consists of two rings that are far apart from each other in high dimensions and the sizes of their radii are much smaller compared with the distance between the two centers. The specific parameters used to generated these datasets are given in Section 4. The evaluations of these methods on the various metrics are given by Table 1.

By looking at Figures 2 and 3, it is clear why we want to introduce these

datasets. In Figure 2, AVIDA clearly preserved the local structure of both datasets




Figure 2: Pamona, AVIDA, SCOT, and t-SNE representation of the dumbbell dataset.

while Pamona and SCOT highlight the linear structure found in both datasets. AVIDA is the only method that is able to successfully integrate the two representations generated by t-SNE’s visualization. Figure 3 shows that Pamona’s method collapses both rings to a single point, destroying the local structure of the data. SCOT is able to integrate the datasets while still preserving some linear structure but compared to t-SNE’s actual visualization, AVIDA has the best representation. Since AVIDA allows t-SNE to construct the local structure of the line before mapping, that structure is preserved in the final visualization.

However, if we were to look at the FOSCTM and accuracy scores in Table 1 for Figure 2 and Figure 3, Pamona scores best because all the points are correctly mapped close together. The datasets illustrate our need for a representation metric since the traditional metrics do not penalize for poor visualizations. We use t-SNE’s loss function as our representation score since it is a popular visualization method, however it could easily be replaced by a loss function from other visualization methods (e.g. UMAP).

### 2.3 AVIDA achieves a balance between structure representation and multimodal dataset alignment

We also compare the outputs from two real-world single-cell multi-omics datasets. The first is sc-GEM, a dataset from [7] which contains both gene expression and




Figure 3: t-SNE, AVIDA, SCOT and Pamona representation of the distant rings dataset.

DNA methylation at multiple loci on human somatic cell samples under conversion to induced pluripotent stem cells. The second is scNMT-seq, a dataset of chromatin accessibility, DNA methylation, and gene expression on mouse gastrulation samples collected at four different time states from [2]. The evaluations of AVIDA, SCOT and Pamona on these datasets are also given in Table 1. In Figure 4, we can see the different visualizations for sc-GEM. The left column of the figure shows the integration between the two datasets and the right column has the datapoints colored by cell. From the visualizations we can see that AVIDA is able to fully integrate the two different datasets where there is some noticeable separation in the SCOT representation. Since this dataset contains the conversion from somatic cells to stem cells, we hope to see a gradient of colors from one end of the visualization to the other which AVIDA is able to achieve. This is a good example of a real-life dataset where AVIDA is able to comparably integrate the datasets while achieving a desirable visualization.

We can also confirm this observation in Table 1. AVIDA is able to achieve FOSCTTM and alignment scores that are comparable to SCOT and Pamona while simultaneously having the best representation loss. The same holds true for scNMT-seq as well. These examples illustrate that AVIDA is comparable to both Pamona and SCOT on real-life datasets while also performing well on the adversarial datasets: the dumbbell and distant rings datasets.

While we did not visualize every dataset here, Figure 5 plots the different visu-




Figure 4: AVIDA, SCOT and Pamona representation of sc-GEM.

alizations by their FOSCTTM and representation scores. The shapes represent the dataset being visualized and the different colors represent the method used for the visualization. We can see that across the different datasets, all three methods have comparable FOSCTTM scores, indicating that the integration of the datasets are similar. However, we can also see that AVIDA by far has the best representation score, indicating a better visualization.

### 3 Discussion

Motivated by the similar fundamental assumptions in dimension reduction and data integration that they both try to preserve the structures of datasets, we developed an alternating method, AVIDA, which combines these two processes for joint visualization of datasets without shared features. Comparing with the methods that perform integration first and then dimension reduction, AVIDA better preserves the detailed structures of the datasets being integrated especially the structures present only one of the datasets. This property allows the identification of mechanisms that can only be revealed with one of the technologies.

In this work, we demonstrate the method using t-SNE for dimension reduction and Gromov-Wasserstein optimal transport for data integration. In general, other dimension reduction methods and integration methods could be used. The




Figure 5: A comparison of methods using integration and visualization.

representation loss used in the comparison can also be used as a control metric about how well the structures of individual datasets are preserved in the joint representation. This metric can be used to find a balance between integration and representation when other methods are used for the dimension reduction and integration modules. The comparison indicates that a method could do a perfect job in integration while missing structures presented in the individual datasets. It is thus important to also evaluate the quality of structure representation of individual datasets when developing joint dimension reduction methods for high-dimensional multimodal datasets.

Despite the improvements on performing the two processes separately, the quality of the joint visualization still heavily depends on the performance of the specific dimension reduction method and integration method. While the quality of dimension reduction can be checked by comparing to the structures present in the original high-dimensional datasets, it is hard to evaluate the integration quality without ground truth. It is thus also important to further validate the result with prior knowledge or assess the robustness of the integration with for example, subsampling.

Upon the joint visualization of multimodal datasets, one major downstream task is to find the correspondence between the non-overlapping features across the datasets. A potential method for this is to track the contributions of original features to the common low-dimensional representations and subsequently find the correspondence between them.

## 4 Methods

AVIDA is a framework that takes input data sets  $\{X^{(\ell)}\}_{i=1}^N$  where the data sets  $X^{(\ell)} \in \mathbb{R}^{n_\ell \times d_\ell}$  need not be in the same feature space. The output of AVIDA is a low dimensional representation of all data sets simultaneously in a single feature space. This is accomplished by alternating between dimensionality reduction and alignment. The AVIDA framework is presented in Algorithm 1. The choice of dimensionality reduction technique and alignment method is up to the user and can be chosen based on the use case. In Section 4.1, we present a detailed im-

plementation of AVIDA using t-SNE for dimensionality reduction and GW-OT for alignment.

---

**Algorithm 1** AVIDA

---

**Input:**  $N$  datasets  $X^{(\ell)} = \{x_i^{(\ell)}\}_{i=1}^{n_\ell} \subset \mathbb{R}^{d_\ell}$ , target dimension  $d$ , Dimensionality Reduction Method  $DR(\cdot)$ , Alignment Method  $ALIGN(\cdot)$ .

**Output:** Low-dimensional representations  $Y^{(\ell)} = \{y_i^{(\ell)}\}_{i=1}^{n_\ell} \subset \mathbb{R}^d$ .

Initialize  $Y_0^{(\ell)}$  for  $\ell \in [N]$  and set  $t = 0$ .

**do**

- Dimensionality reduction step:  
 $\hat{Y}_t^{(\ell)} = DR(X^{(\ell)}, Y_t^{(\ell)})$  for  $\ell \in [N]$ .  $\triangleright$  Input dataset  $X^{(\ell)}$  and initialization  $Y_t^{(\ell)}$
- Alignment step:  
 $[Y_{t+1}^{(1)}, \dots, Y_{t+1}^{(N)}] = ALIGN(\hat{Y}_t^{(1)}, \dots, \hat{Y}_t^{(N)})$ .
- Increment iteration count:  $t = t + 1$ .

**while** stopping criteria not satisfied

Return  $Y^{(\ell)} = Y_t^{(\ell)}$  for  $\ell \in [N]$ .

---

## 4.1 AVIDA with t-SNE and GW-OT

In this section, we present our implementation of the AVIDA framework using t-SNE for dimensionality reduction and GW-OT for alignment, i.e.,  $AVIDA(X_1, X_2; \text{TSNE}, \text{GW})$ .

For simplicity, we assume there are two input data sets  $X^{(1)} = \{x_i^{(1)}\}_{i=1}^{n_1} \subset \mathbb{R}^{d_1}$  and  $X^{(2)} = \{x_i^{(2)}\}_{i=1}^{n_2} \subset \mathbb{R}^{d_2}$  and that the low-dimensional output feature space has dimension  $d = 2$ , i.e.,  $Y^{(1)} = \{y_i^{(1)}\}_{i=1}^{n_1} \subset \mathbb{R}^2$  and  $Y^{(2)} = \{y_i^{(2)}\}_{i=1}^{n_2} \subset \mathbb{R}^2$ .

In the dimensionality reduction step, t-SNE generates pairwise affinity values  $\{p_{ij}^{(\ell)}\}$  for each of the dataset  $X^{(\ell)}$ , as given by

$$p_{j|i}^{(\ell)} = \frac{\exp(-\|x_i^{(\ell)} - x_j^{(\ell)}\|^2/2\sigma_i^{(\ell)})}{\sum_{k \neq i} \exp(-\|x_k^{(\ell)} - x_i^{(\ell)}\|^2/2\sigma_i^{(\ell)})} \quad (3)$$

$$p_{ij}^{(\ell)} = \frac{p_{j|i}^{(\ell)} + p_{i|j}^{(\ell)}}{2n_\ell}, \quad (4)$$

where the  $\sigma_i^{(\ell)}$ 's satisfy

$$\rho = 2^{-\sum_{j \neq i} p_{j|i}^{(\ell)} \log(p_{j|i}^{(\ell)})}, \quad (5)$$

for a perplexity value  $\rho$  chosen by the user. The target probabilities  $q_{ij}^{(\ell)}$  for the representation are defined as:

$$q_{ij}^{(\ell)} = \frac{(1 + \|y_i^{(\ell)} - y_j^{(\ell)}\|^2)^{-1}}{\sum_{i',j'} (1 + \|y_{i'}^{(\ell)} - y_{j'}^{(\ell)}\|^2)^{-1}}. \quad (6)$$

To obtain  $y_i^{(\ell)}$ , t-SNE minimizes the Kullback-Leibler divergence between  $\{p_{ij}^{(\ell)}\}_{j \neq i}$  and  $\{q_{ij}^{(\ell)}\}_{j \neq i}$  using gradient descent:

$$KL(P_\ell || Q_\ell) = \sum_{i,j=1}^{n_\ell} p_{ij}^{(\ell)} \log \left( \frac{p_{ij}^{(\ell)}}{q_{ij}^{(\ell)}} \right), \quad (7)$$

The t-SNE method utilizes a “early exaggeration” phase to artificially highlights the attractions between points in similar neighborhoods, promoting clusters. This period is a very important tool that allows t-SNE to develop local structures in its visualization. The early exaggeration phase occurs in the first 200 iterations of gradient decent in which  $p_{ij}^{(\ell)}$  values scaled by a factor of 4. It has been show that the early exaggeration in t-SNE promotes clustering of similar points [19]. After the first 200 iterations, the  $p_{ij}^{(\ell)}$  values are returned to their original value and t-SNE continues to perform gradient descent.

In the alignment step of AVIDA, GW-OT is used to align data points across data sets. Given the current low dimensional representations outputs from t-SNE,  $Y^{(1)}$  and  $Y^{(2)}$ , the following optimization problem is solved to compute the transport matrix  $\mathbf{T}$ :

$$GW(Y^{(1)}, Y^{(2)}) = \min_{\mathbf{T}} \sum_{i,j,i',j'} \|d(y_i^{(1)}, y_j^{(1)}) - d(y_{i'}^{(2)}, y_{j'}^{(2)})\|^2 \mathbf{T}_{i,i'} \mathbf{T}_{j,j'} - \epsilon(H(\mathbf{T})), \quad (8)$$

where  $H(\mathbf{T}) = \sum_{i,j} \mathbf{T}_{ij} \log(\mathbf{T}_{ij})$  is an Entropic regularization term and  $d(\cdot, \cdot)$  is a chosen distance metric. The representation for  $Y^{(1)}$  is mapped to  $Y^{(2)}$  using the mapping found by minimizing (8), or by computing  $Y^{(1)} = T Y^{(2)}$ . AVIDA( $X^{(1)}, X^{(2)}$ ; TSNE, GW) continues alternating between minimizing the KL loss in t-SNE and using optimal transport to align points until a stopping criteria is reached. In this implementation, we choose to limit the number of iterations to 1000 and perform alignment every 100 iterations after the early exaggeration phase (i.e., after the first 200 iterations) of t-SNE. The pseudo-code for AVIDA( $X^{(1)}, X^{(2)}$ ; TSNE, GW) is provided in Algorithm 2.

## 4.2 Metrics, parameters, hardware

The metrics used in Section 2 are described in detail in this section. For reproducibility, we also include the hardware settings under which these experiments were run and the user selected parameters employed to obtain our numerical results.

### 4.2.1 Metrics

To compare AVIDA( $X_1, X_2$ ; TSNE, GW), Pamona, and SCOT five different metrics are employed: fraction of samples closer than the true match (FOSCTTM), alignment, integration, accuracy, and representation score. The FOSCTTM and alignment are metrics proposed in previous works. FOSCTTM was originally proposed by Liu et al. [20] and was used to validate the performance of SCOT. The alignment score was used in [6] to compare Pamona and SCOT. In addition to the metrics used in previous works, we also introduce a few others to capture various aspects of the output representation. The additional metrics we measure are:

---

**Algorithm 2** AVIDA( $X_1, X_2$ ; TSNE, GW)

---

**Input:** datasets  $X^{(1)} = \{x_1^{(1)}, \dots, x_{n_1}^{(1)}\}$ ,  $X^{(2)} = \{x_1^{(2)}, \dots, x_{n_2}^{(2)}\}$ , perplexity  $\rho$ , and regularization parameter  $\varepsilon$

**Output:** low-dimensional representations:  $Y_0^{(1)} = \{y_1^{(1)}, \dots, y_{n_1}^{(1)}\}$ ,  $Y_0^{(2)} = \{y_1^{(2)}, \dots, y_{n_2}^{(2)}\}$

Compute pairwise affinities  $p_{ij}^{(1)}, p_{ij}^{(2)}$  with perplexity  $\rho$  (using Eq. (3) and Eq. (4))

Initialize solutions  $Y_0^{(1)}, Y_0^{(2)}$  with points drawn i.i.d. from  $\mathcal{N}(0, 10^{-4}I)$

**while**  $t < 1000$  **do**

**if**  $\text{mod}(t, 100) \neq 0$  **then**

**for**  $\ell = 1, 2$  **do**

            Compute pairwise affinities  $q_{ij}^{(\ell)}$  (using Eq. 6)

            Compute gradients  $\Delta_t^{(\ell)} = \frac{\delta}{\delta Y_t^{(\ell)}} \text{TSNE}(X^{(\ell)}, Y_t^{(\ell)})$  (using Eq. 7)

            Set  $Y_t^{(\ell)} = Y_t^{(\ell)} + \Delta_t^{(\ell)}$

**end for**

**else**

        Compute the GW-OT mapping,  $\mathbf{T}$ , between  $Y_t^{(1)}$  and  $Y_t^{(2)}$  (using Eq. 1)

        Set  $Y_{t+1}^{(\ell)} = \mathbf{T}Y_t^{(\ell)}$

**end if**

$t \leftarrow t + 1$

**end while**

---

integration, accuracy, and objective loss. In this section, we define each and the conditions under which these metrics are meaningful. For notational simplicity,  $D \in \mathbb{R}^{n_1 \times n_2}$  such that  $D_{ij} = d(y_i^{(1)}, y_j^{(2)})$  denote the pairwise distance matrix between points in  $Y^{(1)}$  and points in  $Y^{(2)}$ .

The FOSCTTM captures roughly the accuracy of the representation. FOSCTTM operates under the assumption that every point has a “true match” and that the “true matches” should be close together in the lower-dimensional representation. This is formalized as follows. Assume, for simplicity, and  $n_1 = n_2 = n$  and without loss of generality that the true match of  $x_i^{(1)}$  is  $x_i^{(2)}$  for all  $i \in [n]$ . The FOSCTTM is defined as:

$$\text{FOSCTTM} = \sum_{i=1}^n \frac{|\{j : D_{ij} < D_{ii}\}|}{n-1} + \sum_{j=1}^n \frac{|\{i : D_{ij} < D_{jj}\}|}{n-1}. \quad (9)$$

In other words, for each point  $Y^{(1)}$ , determine the fraction of the points  $y_i^{(2)}$  that are closer to  $y_i^{(1)}$  than  $y_i^{(2)}$ . Then, repeat the process for points in  $Y^{(2)}$ . Smaller values of FOSCTTM indicate better performance.

Under these same assumptions (that every point has a true match), we can also define an accuracy score. The idea is that points that are true matches should appear close together in the lower dimensional representation. This is measured

by taking a simple trace of the matrix  $D$ :

$$\text{Accuracy} = \sum_{i=1}^n D_{ii} = \text{tr}(D)$$

The Alignment score used in this work was also used in [6]. The alignment score measures how well aligned the two datasets being integrated are in low dimensions. For the alignment score, we assume that each data set has class labels and that those class labels can be shared across data sets. The points in each data set is split into “shared” and “dataset specific”. “Shared” data points have representation in both  $Y^{(1)}$  and  $Y^{(2)}$  whereas “dataset specific” data points only appear in one of the datasets. The alignment score is computed as follows. Let  $S^{(1)} \cup P^{(1)} = Y^{(1)}$  and  $S^{(2)} \cup V^{(2)} = Y^{(2)}$  where sets  $S^{(\ell)}$  denote the set of all points corresponding to “shared” data points and  $V^{(\ell)}$  denote the set indices of all dataset specific points in  $Y^{(\ell)}$ . The alignment score is defined as:

$$\text{Alignment} = 1 - \frac{|\bar{x}_s - k/(\ell+1)|}{k - k/(\ell+1)},$$

where  $\bar{x}_s$  is the average number of nearest neighbors that are shared points from the same dataset.

The aforementioned metrics have been utilized in previous works. We also propose to use the following for evaluating the representation of the low dimensional data. First, we employ a symmetrized Kullback-Leibler loss with a student t-distribution kernel to evaluate how well the visualization represents the high-dimensional data in an integrated fashion. We refer to this as the Representation Loss:

$$\begin{aligned} \text{Representation Loss} &= \frac{1}{2} \left( \text{KL}(X^{(1)}, Y^{(1)}) + \text{KL}(Y^{(1)}, X^{(1)}) \right) \\ &+ \frac{1}{2} \left( \text{KL}(X^{(2)}, Y^{(2)}) + \text{KL}(Y^{(2)}, X^{(2)}) \right). \end{aligned}$$

Lastly, we want to evaluate how well integrated the two data sets are in low dimensions. We say that integration is the average, minimum distance between a data point in  $Y_1$  and any data point in  $Y_2$ . The integration is defined as:

$$\text{Integration} = \frac{1}{n_1} \sum_{i=1}^{n_1} \min_j D_{ij} + \frac{1}{n_2} \sum_{j=1}^{n_2} \min_i D_{ij}.$$

#### 4.2.2 Parameters

The default perplexity value in most standard implementations of t-SNE is 30. However, depending on the dataset, the perplexity value may need to be adjusted. Table 2 shows the perplexity value choices for each experiment presented in Section 2. In addition to perplexity, another important parameter is  $\epsilon$  in Equation 1.

| Dataset          | Bifurcated Tree | Circular Frustrum | Dumbbell | Distant Rings | sc-GEM | scNMT-seq |
|------------------|-----------------|-------------------|----------|---------------|--------|-----------|
| Perplexity Value | 30              | 60                | 30       | 30            | 50     | 100       |

Table 2: Perplexity choices for each dataset.

For all of our experiments,  $\epsilon$  was set to be  $5 \times 10^{-3}$  but depending on the dataset could be adjusted.

### 4.2.3 Hardware

We ran the experiments on an Intel i7-10750H CPU (base frequency 2.60GHz) with 8GB memory.

## 4.3 Datasets

For our analysis, we introduced two synthetic datasets: the dumbbell dataset and distant rings dataset. The dumbbell dataset consists of two subdatasets,  $X^{(d,1)}, X^{(d,2)} \subset \mathbb{R}^2$  with 200 datapoints each. For all  $0 \leq i \leq 200$ ,

$$\begin{aligned} X_{i,1}^{(d,1)} &\sim 50U(0, 1) \\ X_{i,2}^{(d,1)} &\sim N(0, 1) \end{aligned}$$

where  $U(0, 1)$  is the uniform distribution and  $N(0, 1)$  is the normal distribution. This essentially constructs  $X^{(d,1)}$  as a line in 2D with a little bit of noise. To construct the two rings in  $X^{(d,2)}$ , we consider  $\theta \sim U(0, 2\pi)$  and  $r \sim N(3, 0.5)$ , then use it in our construction.

$$\begin{aligned} X_{i,1}^{(d,2)} &\sim r \cos(\theta), \quad 1 \leq i \leq 50 \\ X_{i,2}^{(d,2)} &\sim r \sin(\theta), \quad 1 \leq i \leq 50 \\ X_{i,1}^{(d,2)} &\sim r \cos(\theta) + 14, \quad 50 < i \leq 100 \\ X_{i,2}^{(d,2)} &\sim r \sin(\theta), \quad 50 < i \leq 100 \end{aligned}$$

The first 50 points in  $X^{(2)}$  are a slightly noisy circle centered at 0, where the next 50 points in the dataset are the same slightly noisy circle centered instead at 14. These two rings are then connected by a line.

$$\begin{aligned} X_{i,1}^{(d,2)} &\sim U(3, 10), \quad 100 < i \leq 200 \\ X_{i,2}^{(d,2)} &\sim N(0, 0.2), \quad 100 < i \leq 200 \end{aligned}$$

This line is the last 100 points and also has small noise across one dimension.

The distant rings dataset also contains two subdatasets,  $X^{(c,1)}, X^{(c,2)} \subset \mathbb{R}$ . Again, we let  $\theta \sim U(0, 2\pi)$  and now we define  $r_1 \sim N(5, 1)$  and  $r_2 \sim N(5, 0.1)$  and define two different rings.

$$\begin{aligned} X_{:,1}^{(c,1)} &\sim r_1 \cos(\theta) \\ X_{:,2}^{(c,1)} &\sim r_1 \sin(\theta) \\ X_{:,1}^{(c,2)} &\sim r_2 \cos(\theta) + 100 \\ X_{:,2}^{(c,2)} &\sim r_2 \sin(\theta) + 100 \end{aligned}$$

Essentially for each dataset, we construct two rings where the distance between them dwarfs the radius of each ring. To make these two rings distinct, we constructed one ring to have much less noise than the other.

## 5 Data Availability

The synthetic data distant rings and dumbbell dataset are available at <https://github.com/kat-dover/AVIDA/tree/main/data> and the bifurcated tree and circular frustrum were downloaded from the SCOT repository <https://rsinghlab.github.io/SCOT/data/>. The sc-GEM data from [7] was downloaded from the SCOT repository given at <https://rsinghlab.github.io/SCOT/data/>. The scNMT-seq data from [2] were downloaded from the Pamona repository given at <https://github.com/caokai1073/Pamona>.

## 6 Code Availability

The AVIDA implementation with t-SNE as the dimension reduction module and Gromov-Wasserstein optimal transport as the alignment module is available at <https://github.com/kat-dover/AVIDA> which will be made publically available on Github upon publication.

## 7 Acknowledgement

R.V. acknowledges support from NSF DMS-1954233, NSF DMS-2027299, U.S. Army 76649-CS, and NSF+Simons Research Collaborations on the Mathematical and Scientific Foundations of Deep Learning. Q.N. was partly supported by a NSF grant DMS1763272 and a grant from the Simons Foundation (594598, QN). Z. C. was partly supported by a start-up fund of North Carolina State University and NSF DMS-2151934.

## References

- [1] W. M. Abdelmoula, B. Balluff, S. Englert, J. Dijkstra, M. J. T. Reinders, A. Walch, L. A. McDonnell, and B. P. F. Lelieveldt. Data-driven identification of prognostic tumor subpopulations using spatially mapped t-sne of mass spectrometry imaging data. *Proceedings of the National Academy of Sciences*, 113(43):12244–12249, 2016.
- [2] R. Argelaguet, S. J. Clark, H. Mohammed, L. C. Stapel, C. Krueger, C.-A. Kapourani, I. Imaz-Rosshandler, T. Lohoff, Y. Xiang, C. W. Hanna, et al. Multi-omics profiling of mouse gastrulation at single-cell resolution. *Nature*, 576(7787):487–491, 2019.
- [3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- [4] J. D. Buenrostro, B. Wu, U. M. Litzenburger, D. Ruff, M. L. Gonzales, M. P. Snyder, H. Y. Chang, and W. J. Greenleaf. Single-cell chromatin accessibility reveals principles of regulatory variation. *Nature*, 523(7561):486–490, 2015.
- [5] Z. Cang and Q. Nie. Inferring spatial and signaling relationships between cells from single cell transcriptomic data. *Nature communications*, 11(1):1–13, 2020.

- [6] K. Cao, Y. Hong, and L. Wan. Manifold alignment for heterogeneous single-cell multi-omics data integration using pamona. *Bioinformatics*, 38(1):211–219, 2022.
- [7] L. F. Cheow, E. T. Courtois, Y. Tan, R. Viswanathan, Q. Xing, R. Z. Tan, D. S. Tan, P. Robson, Y.-H. Loh, S. R. Quake, et al. Single-cell multimodal profiling reveals cellular epigenetic heterogeneity. *Nature methods*, 13(10):833–836, 2016.
- [8] N. Courty, R. Flamary, and D. Tuia. Domain adaptation with regularized optimal transport. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 274–289. Springer, 2014.
- [9] P. Demetsci, R. Santorella, B. Sandstede, W. S. Noble, and R. Singh. Scot: Single-cell multi-omics alignment with optimal transport. *Journal of Computational Biology*, 29(1):3–18, 2022.
- [10] S. Ferradans, N. Papadakis, G. Peyré, and J.-F. Aujol. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences*, 7(3):1853–1882, 2014.
- [11] R. Flamary, N. Courty, A. Gramfort, M. Z. Alaya, A. Boisbunon, S. Chambon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, L. Gautheron, N. T. Gayraud, H. Janati, A. Rakotomamonjy, I. Redko, A. Rolet, A. Schutz, V. Seguy, D. J. Sutherland, R. Tavenard, A. Tong, and T. Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- [12] M. Forcato, O. Romano, and S. Bicciato. Computational methods for the integrative analysis of single-cell data. *Briefings in bioinformatics*, 22(3):bbaa042, 2021.
- [13] Y. Hao, S. Hao, E. Andersen-Nissen, W. M. Mauck III, S. Zheng, A. Butler, M. J. Lee, A. J. Wilk, C. Darby, M. Zager, et al. Integrated analysis of multimodal single-cell data. *Cell*, 184(13):3573–3587, 2021.
- [14] B. Hie, B. Bryson, and B. Berger. Efficient integration of heterogeneous single-cell transcriptomes using scanorama. *Nature biotechnology*, 37(6):685–691, 2019.
- [15] M. S. Jain, K. Polanski, C. D. Conde, X. Chen, J. Park, L. Mamanova, A. Knights, R. A. Botting, E. Stephenson, M. Haniffa, et al. Multimap: dimensionality reduction and integration of multimodal data. *Genome biology*, 22(1):1–26, 2021.
- [16] L. Jimenez and D. Landgrebe. Supervised classification in high-dimensional space: geometrical, statistical, and asymptotical properties of multivariate data. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 28(1):39–54, 1998.
- [17] L. Kantorovitch. On the translocation of masses. *C. R. (Doklady) Acad. Sci. URSS (N.S.)*, 37(199–201), 1942.
- [18] D. Lahat, T. Adali, and C. Jutten. Multimodal data fusion: an overview of methods, challenges, and prospects. *Proceedings of the IEEE*, 103(9):1449–1477, 2015.
- [19] G. C. Linderman and S. Steinerberger. Clustering with t-sne, provably. *SIAM Journal on Mathematics of Data Science*, 1(2):313–332, 2019.

- [20] J. Liu, Y. Huang, R. Singh, J.-P. Vert, and W. S. Noble. Jointly embedding multiple single-cell omics measurements. In *Algorithms in bioinformatics:... International Workshop, WABI..., proceedings. WABI (Workshop)*, volume 143. NIH Public Access, 2019.
- [21] B. Lorincz, A. Stan, and M. Giurgiu. An objective evaluation of the effects of recording conditions and speaker characteristics in multi-speaker deep neural speech synthesis. *ArXiv*, abs/2106.01812, 2021.
- [22] L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [23] F. Mémoli. Gromov–wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11(4):417–487, 2011.
- [24] G. Monge. Memoire sur la theorie des deblais et des remblais. *Histoire de l'Academie Royale des Sciences*, (666-704), 1781.
- [25] T. Nagano, Y. Lubling, T. J. Stevens, S. Schoenfelder, E. Yaffe, W. Dean, E. D. Laue, A. Tanay, and P. Fraser. Single-cell hi-c reveals cell-to-cell variability in chromosome structure. *Nature*, 502(7469):59–64, 2013.
- [26] M. Nitzan, N. Karaiskos, N. Friedman, and N. Rajewsky. Gene expression cartography. *Nature*, 576(7785):132–137, 2019.
- [27] K. O’Connor, B. Yi, K. McGoff, and A. B. Nobel. Graph optimal transport with transition couplings of random walks. *arXiv preprint arXiv:2106.07106*, 2021.
- [28] K. Pearson. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11):559–572, 1901.
- [29] G. Peyré, M. Cuturi, and J. Solomon. Gromov-wasserstein averaging of kernel and distance matrices. In *International Conference on Machine Learning*, pages 2664–2672. PMLR, 2016.
- [30] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.
- [31] G. Schiebinger, J. Shu, M. Tabaka, B. Cleary, V. Subramanian, A. Solomon, J. Gould, S. Liu, S. Lin, P. Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- [32] T. Stuart and R. Satija. Integrative single-cell analysis. *Nature reviews genetics*, 20(5):257–272, 2019.
- [33] M. J. Stubbington, O. Rozenblatt-Rosen, A. Regev, and S. A. Teichmann. Single-cell transcriptomics to explore the immune system in health and disease. *Science*, 358(6359):58–63, 2017.
- [34] V. Svensson, R. Vento-Tormo, and S. A. Teichmann. Exponential scaling of single-cell rna-seq in the past decade. *Nature protocols*, 13(4):599–604, 2018.
- [35] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

- [36] L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [37] M. Verma, G. Matijević, C. Denker, A. Diercke, E. Dineva, H. Balthasar, R. Kamlah, I. Kontogiannis, C. Kuckein, and P. S. Pal. Classification of high-resolution solar  $h\alpha$  spectra using t-distributed stochastic neighbor embedding. *The Astrophysical Journal*, 907(1):54, 2021.
- [38] C. Villani. *Topics in optimal transportation*. American Mathematical Society, Providence, RI, 2003.
- [39] D. E. Wagner, C. Weinreb, Z. M. Collins, J. A. Briggs, S. G. Megason, and A. M. Klein. Single-cell mapping of gene expression landscapes and lineage in the zebrafish embryo. *Science*, 360(6392):981–987, 2018.
- [40] J. D. Welch, V. Kozareva, A. Ferreira, C. Vanderburg, C. Martin, and E. Z. Macosko. Single-cell multi-omic integration compares and contrasts features of brain cell identity. *Cell*, 177(7):1873–1887, 2019.
- [41] L. Zhang and Q. Nie. scmc learns biological variation through the alignment of multiple single-cell genomics datasets. *Genome biology*, 22(1):1–28, 2021.
- [42] S. Zhang and H. Tong. Network alignment: recent advances and future directions. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 3521–3522, 2020.
- [43] C. Zhu, S. Preissl, and B. Ren. Single-cell multimodal omics: the power of many. *Nature methods*, 17(1):11–14, 2020.