



Published in final edited form as:

Proc SIAM Int Conf Data Min. 2020 ; 2020: 316–324. doi:10.1137/1.9781611976236.36.

Harmonic Alignment

Jay S. Stanley III^{†,‡}, Scott Gigante^{†,§}, Guy Wolf^{*,¶}, Smita Krishnaswamy^{*,||}

[‡]Yale University, Appl. Math. Prog.

[§]Yale University, Comp. Bio. & Bioinf. Prog.

[¶]Université de Montréal, Dept. of Math. & Stat.; Mila – Quebec AI Institute

^{||}Yale University, Depts. of Gene. & Comp. Sci.

Abstract

We propose a novel framework for combining datasets via alignment of their intrinsic geometry. This alignment can be used to fuse data originating from disparate modalities, or to correct batch effects while preserving intrinsic data structure. Importantly, we do not assume any pointwise correspondence between datasets, but instead rely on correspondence between a (possibly unknown) subset of data features. We leverage this assumption to construct an isometric alignment between the data. This alignment is obtained by relating the expansion of data features in harmonics derived from diffusion operators defined over each dataset. These expansions encode each feature as a function of the data geometry. We use this to relate the diffusion coordinates of each dataset through our assumption of partial feature correspondence. Then, a unified diffusion geometry is constructed over the aligned data, which can also be used to correct the original data measurements. We demonstrate our method on several datasets, showing in particular its effectiveness in biological applications including fusion of single-cell RNA sequencing (scRNA-seq) and single-cell ATAC sequencing (scATAC-seq) data measured on the same population of cells, and removal of batch effect between biological samples.

1 Introduction

High dimensional data have become increasingly common in many fields of science and technology, and with them the need for robust representations of intrinsic structure and geometry in data, typically inferred via manifold learning methods. Furthermore, modern data collection technologies often produce multisample data, which contain multiple datasets (or data batches) that aim to capture the same phenomena but originate from different equipment, different calibration, or different experimental environments. These introduce new challenges in manifold learning, as naïve treatment in such cases produces data geometry that largely separates different data batches into separate manifolds. Therefore, special processing is required to align and integrate the data manifolds in such cases in order to allow for the study and exploration of relations between and across multiple datasets.

Corr. author, smita.krishnaswamy@yale.edu.

[†]Equal contribution

^{*}Equal contribution

As a particular application field, we focus here on single cell data analysis, which has gained importance with the advent of new sequencing technologies, such as scRNA-seq and scATAC-seq. While numerous works have shown that manifold learning approaches are particularly effective on such data [1], a common challenge in their analysis is a set of technical artifacts termed *batch effects* (caused by data collection from separate experimental runs) that tend to dominate downstream analysis, unless explicitly corrected. For instance, it is often the case that naïve data manifold construction groups such data into clusters that correspond to measurement time or equipment used, rather than by meaningful biological variations. Under such circumstances, it is necessary that batch artifacts be eliminated while actual biological differences between the samples be retained. Further, it can also be the case that different variables of data are measured on the same biological system. For example, cells from the same tissue can be measured with transcriptomic and proteomic technologies. However, the cells themselves are destroyed in each measurement, even though they are sampled from the same underlying cellular manifold. Therefore, there is no correspondence that can be established between the two sets of measurements directly.

Recent manifold learning methods focusing on multisample data often treat each dataset as a different “view” of the same system or latent manifold, and construct a multiview geometry (e.g., based on the popular diffusion maps framework [2]) to represent them (e.g., [3–8]). Importantly, these methods often require at least partial, if not full, bijection between views (e.g., both sets of measurements conducted on the same cells), which is often impossible to obtain in experimental scenarios where data is collected asynchronously or independently. In particular, as mentioned above, genomic and proteomic data (especially at the single-cell resolution) often originate on destructive collection technologies, and thus data point (i.e., cell) correspondence becomes an impractical (if not impossible) assumption to impose on their analysis. Other works attempt to directly match data points, either in the ambient space [9,10] or by local data geometry [11]. These approaches can be very sensitive to differences in sampling density rather than data geometry, as discussed in §2-3. Furthermore, a complete matching is often not feasible as certain datasets may contain distinct local phenomena (e.g., rare subpopulation only captured in one dataset but not present in the other).

In this paper, we formulate the processing of multisample data with no data point correspondence in terms of manifold alignment, and present an approach towards such alignment by bridging the geometric-harmonic framework provided by diffusion geometry [2] (§3.1) together with data feature filtering enabled by graph signal processing [12] (§3.2). Our alignment approach relies on correspondence between underlying features quantified by data collection or measurement systems, phrased here as feature correspondence, which is often more realistic than data point correspondence. Indeed, related systems often observe similar “entities” (e.g., cells, patients) and aim to capture related properties in them. As explained in §2 and §3.2, we treat measured data features as manifold signals (i.e., over the data manifold) and relate them to intrinsic coordinates of a diffusion geometry [2] of each dataset, which also serve as intrinsic data harmonics. Then, as explained in §4, we leverage feature correspondence to capture pairwise relations between the intrinsic diffusion coordinates of the separate data manifolds (i.e., of each dataset). Finally, we use these relations to compute an isometric transformation that aligns the data manifolds on top of each other without distorting their internal structure.

We demonstrate the results of our method in §5 on artificial manifolds and single-cell biological data for both batch effect removal and multimodal data fusion. In each case, our method successfully aligns data manifolds such that they have appropriate neighbors both within and across the two datasets. Further, we show an application of our approach in transfer learning by applying a k-NN classifier to one unlabeled dataset based on labels provided by another dataset (with batch effects between them), and compare the classification accuracy before and after alignment. Finally, comparisons with recently developed methods such as the MNNbased method from [9] show significant improvements in performance and denoising by our harmonic alignment methods.

2 Problem setup

Let $X = \{x_1, \dots, x_M\}$ and $Y = \{y_1, \dots, y_N\}$ be two finite datasets that aim to measure the same phenomena. For simplicity, we assume that both datasets have the same number of features, i.e., $X, Y \subseteq \mathbb{R}^n$ for some sufficiently high dimension n . We consider here a setting where these datasets are collected via different instruments or environment, but are expected to capture some equivalent information which can be used to align the two datasets. To leverage a manifold learning approach in such settings, we consider the common latent geometry of the data as an unknown manifold \mathcal{M} , which is mapped to the two feature spaces via functions $\mathbf{f}, \mathbf{g}: \mathcal{M} \rightarrow \mathbb{R}^n$ that represent the two data spaces. Namely, each data point $x \in X$ is considered as a result $x = \mathbf{f}(z) = (f_1(z), \dots, f_n(z)) \in \mathbb{R}^n$ for some $z \in \mathcal{M}$ and similarly each $y \in Y$ as a result of $y = \mathbf{g}(z) = (g_1(z), \dots, g_n(z)) \in \mathbb{R}^n$. Therefore, we aim to provide a common data representation of both datasets, which captures the geometry of \mathcal{M} while allowing data fusion of X, Y and integrated processing or analysis of their data features.

We note that while we clearly do not have access to the points $z \in \mathcal{M}$ on the underlying manifold, we do have access to a finite sampling of the feature functions $f_s, g_s: \mathcal{M} \rightarrow \mathbb{R}$, $s = 1, \dots, n$, by considering X, Y as points-by-features matrices (i.e., rewriting them, by slight abuse of notation, as $X \in \mathbb{R}^{M \times n}$ and $Y \in \mathbb{R}^{N \times n}$, rather than finite subsets of \mathbb{R}^n) and taking their corresponding columns. Further, previous multiview manifold learning methods typically consider aligned datasets, i.e., assuming that all feature functions are sampled over the same manifold points. Instead, here we remove this assumption, thus allowing independently sampled datasets, and replace it with a feature correspondence assumption. Namely, we assume the feature functions f_s, g_s (for given $1 \leq s \leq n$) aim to capture similar structures in the data and should therefore share some common information, although each may also contain sensor-specific or dataset-specific bias. While this is an informal notion, it fits well with many experimental data collection settings.

3 Preliminaries and background

3.1 Diffusion maps

To learn a manifold geometry from collected data we use the diffusion maps (DM) construction [2], which we briefly describe here for one of the data spaces X , but is

equivalently constructed on Y . This construction starts by considering local similarities, which we quantify via an anisotropic kernel

$$\mathcal{K}(x_i, x_j) = \frac{\mathcal{G}(x_i, x_j)}{\|\mathcal{G}(x_i, \cdot)\|_1 \|\mathcal{G}(x_j, \cdot)\|_1}, \quad (3.1)$$

where $\mathcal{G}(x_i, x_j) = \exp(-\|x_i - x_j\|^2/\sigma)$ is the Gaussian kernel with neighborhood radius $\sigma > 0$. As shown in [2], this kernel provides neighborhood construction that is robust to sampling density variations and enables separation of data geometry from its distribution. Next, the kernel \mathcal{K} is normalized to define transition probabilities $p(x_i, y_j) = \mathcal{K}(x_i, x_j)/\|\mathcal{K}(x_i, \cdot)\|_1$ that define a Markovian diffusion process over the data. Finally, a DM is defined by organizing these probabilities in a row stochastic matrix \mathbf{P} (typically referred to as the diffusion operator) as $\mathbf{P}_{ij} = p(x_i, x_j)$, and using its eigenvalues $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$ and (corresponding) eigenvectors $\{\phi_j\}_{j=1}^N$ to map each $x_i \in X$ to diffusion coordinates $\Phi_t(x_i) = [\lambda_1^t \phi_1(x_i), \dots, \lambda_N^t \phi_N(x_i)]^T$. The parameter t in this construction represents a diffusion time or the number of transitions considered in the diffusion process. To simplify notations, we also use $\Phi_t = \{\Phi_t(x_i): x_i \in X\}$ to denote the DM of the entire dataset X . We note that in general, as t increases, most of the eigenvalue weights $\lambda_j^t, j = 1, \dots, N$, become numerically negligible, and thus truncated DM coordinates (i.e., using only non-negligible weights) can be used for dimensionality reduction purposes, as discussed in [2].

3.2 Graph Fourier transform

A classic result in spectral graph theory (see, e.g., [13]) shows that the discrete Fourier basis (i.e., pure harmonics, such as sines and cosines, organized by their frequencies) can be derived as Laplacian eigenvectors of the ring graphs. This result was recently used in graph signal processing [12] to define a *graph Fourier transform* (GFT) by treating eigenvectors of the graph Laplacian as generalized Fourier harmonics (i.e., intrinsic sines and cosines over a graph). Further, as discussed in [2,14], diffusion coordinates are closely related to these Laplacian eigenvectors, and can essentially serve as geometric harmonics over data manifolds.

In our case, we regard the kernel \mathcal{K} from §3.1 as a weighted adjacency matrix of a graph whose vertices are the data point in X . Then, the resulting normalized graph Laplacian is given by $\mathcal{L} = \mathbf{I} - \mathbf{D}^{1/2} \mathbf{P} \mathbf{D}^{-1/2}$, where \mathbf{D} is a diagonal matrix with $\mathbf{D}_{ii} = \|\mathcal{K}(x_i, \cdot)\|_1$.

Therefore, the eigenvectors of \mathcal{L} can be written as $\psi_j = \mathbf{D}^{1/2} \phi_j$ with corresponding eigenvalues $\omega_j = 1 - \lambda_j$. The resulting GFT of a signal (or function) f over X can thus be written as $\hat{f}[j] = \langle f, \psi_j \rangle = \langle f, \mathbf{D}^{1/2} \phi_j \rangle$. We note that here we treat either ω_j or λ_j as providing a “frequency” organization of their corresponding eigenvectors ψ_j or ϕ_j (treated as intrinsic harmonics). In the latter case, eigenvectors with higher eigenvalues correspond to lower frequencies on the data manifold, and vice versa. As noted before, the same construction of GFT here and DM in §3.1 can be equivalently constructed for Y as well. This frequency-based organization of diffusion coordinates derived from X and Y , and their treatment as geometric harmonics, will be leveraged in §4 to provide an isometric alignment between the

intrinsic data manifolds represented by the DMs of the two datasets by also leveraging the (partial) feature correspondence assumption from §2.

3.3 Related work on manifold alignment

Algorithms for semi-supervised and unsupervised manifold alignment exist in classical statistics [15,16], deep learning [10,17,18] and manifold learning [4,9,11]. As mentioned in §1, much work has been done on finding common manifolds between data based on known (partial) bijection between data points [5–8]. In a sense, these methods can be regarded as nonlinear successors of the classic canonical correlation analysis (CCA) [16], in the same way as many manifold learning methods can be regarded as generalizing PCA. Indeed, similar to PCA, the CCA method finds a common linear projection, but on directions that maximize covariance or correlation (typically estimated empirically via known pointwise correspondence) between datasets rather than just variance within one of them. However, in this work we mainly focus on settings where no data point correspondence is available, and therefore we focus our discussion in this section on related work that operate in such settings.

One of the earliest attempts at manifold alignment (in particular, with no point correspondence), was presented in [11], which proposes a linear method based on embedding a joint graph built over both datasets to preserve local structure in both manifolds. This method provides a mapping from both original features spaces to a new feature space defined by the joint graph, which is shared by both datasets with no assumption of feature correspondence. More recently, in biomedical data analysis, mutual nearest neighbors (MNN) batch correction [9] focuses on families of manifold deformations that are often encountered in biomedical data. There, locally linear manifold alignment is provided by calculating a correction vector for each point in the data, as defined by the distances from the point to all points for which it is a mutual k -nearest neighbor. This correction vector is then smoothed by taking a weighted average over a Gaussian kernel.

Beyond manifold learning settings, deep learning methods have been proposed to provide alignment and transfer learning between datasets. For example, cycle GANs [17] are a class of deep neural network in which a generative adversarial network (GAN) is used to learn a nonlinear mapping from one domain to another, and then a second GAN is used to map back to the original domain. These networks are then optimized to (approximately) satisfy cycle consistency constraints such that the result of applying the full cycle to a data point reproduces the original point. MAGAN [10] is a particular cycle GAN that adds a supervised partial feature correspondence to enforce alignment of two data manifolds over the mapping provided by the trained network. However, this correspondence can be disturbed by noise or sparsity in the data.

Additionally, a similar problem exists in isometric shape matching, albeit limited to low dimensional data (i.e., shapes in at most three dimensions). For example, the method in [19] takes shapes with a known Laplace-Beltrami operator and aligns them using a representation of the corresponding eigenfunctions. This work was extended further in [20] to settings where ambient functions are defined intrinsically by the shape in order to learn region-region correspondences from an unknown bijection. Recent work [21] has relaxed the requirement

for shapes to be isometric and the need for prior knowledge of the Laplace-Beltrami operator, instead estimating the manifold with a kernel density estimate over the shape boundary. However, the application of these methods is limited to shapes, rather than a regime of point clouds as seen in high-dimensional data analysis.

In contrast, in this work we consider more general settings of aligning intrinsic data manifolds in arbitrary dimensions, while being robust to noise, data collection artifacts, and density variations. We provide a nonlinear method for aligning two datasets using their diffusion maps [2] (§3.1) under the assumption of a partial feature correspondence. Unlike MAGAN, we do not need to know in advance which features should correspond, and our results show that even with correspondence as low as 15% we achieve good alignment between data. Further, unlike shape matching methods, we are not limited to datasets describing the boundary of a shape or dominated by density distribution. Our formulation allows us to obtain more information from datasets with partial feature correspondence than methods that assume no correspondence, but without the burden of determining in advance which or how many features correspond. To evaluate our method, in §5 we focus on comparison with MAGAN, as a leading representative of deep learning approaches, and MNN, as a leading representative of manifold learning approaches. We note that to the best of our knowledge, the method in [11] is not provided with standard implementation, and our attempts at implementing the algorithm have significantly under performed other methods. For completeness, partial comparison to this method is demonstrated in Appendix C.

4 Harmonic alignment

Given datasets X, Y , as described in §2, we aim to construct a unified DM over both of them, which represents the global intrinsic structure of their common manifold \mathcal{M} while still retaining local differences between the datasets (e.g., due to distributional differences or local patterns only available in one of the dataset). As mentioned before, global shifts and batch effects often make direct construction of such DM (or even the construction of local neighborhood kernels) over the union of both datasets unreliable and impractical. Instead, we propose here to first construct two separate DMs $\Phi^{(X)}, \Phi^{(Y)}$ (based on eigenpairs $(\lambda_i^{(X)}, \phi_i^{(X)}), i = 1, \dots, M$, and $(\lambda_j^{(Y)}, \phi_j^{(Y)}), j = 1, \dots, N$, correspondingly), which capture the intrinsic geometry of each dataset. We then align their coordinates via an orthogonal transformation that preserves the rigid structure of the data in each DM, which is computed by orthogonalizing a correlation matrix computed between the diffusion coordinates of the two DMs.

However, since the diffusion coordinates are associated with intrinsic notions of frequency on data manifolds (as explained in §3.2, there is no need to compute the correlation between every pair $\phi_i^{(X)}, \phi_j^{(Y)}, i = 1, \dots, M, j = 1, \dots, N$. Indeed, leveraging the interpretation of such coordinate functions as intrinsic *diffusion harmonics*, we can determine that they should not be aligned between the geometry of X and Y if their corresponding frequencies (i.e., captured via the eigenvalues $\lambda_i^{(X)}, \lambda_j^{(Y)}$) are sufficiently far from each other. Therefore, in §4.1 we describe the construction of a bandlimited correlation matrix, and then use it in §4.2

to align the two DMs. Since our alignment method is based on the treatment of DM coordinates as manifold harmonics, we call this method *harmonic alignment*.

4.1 Bandlimited correlation

In order to partition the diffusion harmonics into local frequency bands, we consider the following window functions, which are inspired by the itersine filter bank construction [22]:

$$w_{\xi}(\lambda) = \begin{cases} \sin\left(\frac{\pi}{2}\cos^2\left(\frac{\pi}{2}(\ell\lambda - \xi)\right)\right) \frac{\xi - 1}{\ell} \leq \lambda \leq \frac{\xi + 1}{\ell} \\ 0 & \text{otherwise,} \end{cases}$$

where $\xi = 0, \dots, \ell$ and ℓ considered as a meta-parameter of the construction. We note that experimental evidence indicate that fine tuning ℓ does not significantly affect alignment quality. Each window $w_{\xi}(\cdot)$ is supported on an interval of length 2ℓ around ξ/ℓ , while decaying smoothly from $w_{\xi}(\xi/\ell) = 1$ to zero. Two consecutive windows (i.e., $w_{\xi}(\cdot)$, $w_{\xi+1}(\cdot)$) share an overlap of half their support; otherwise (i.e., $w_{\xi}(\cdot)$, $w_{\xi'}(\cdot)$ with $|\xi - \xi'| \geq 2$) they have disjoint supports. Finally, we recall the spectra (i.e., eigenvalues) of $P^{(X)}$, $P^{(Y)}$ are contained in the interval $[0, 1]$, which is entirely covered by $\ell + 1$ window functions $w_{\xi}(\cdot)$, $\xi = 0, \dots, \ell$, as illustrated in Fig. 1(c). Notice that only half the support of $w_0(\cdot)$ and $w_{\ell}(\cdot)$ are shown in there, since half of their support is below zero or above one, correspondingly.

Using the soft partition defined by $w_{\xi}(\cdot)$, $\xi = 0, \dots, \ell$, we now define bandlimiting weights

$$w_{ij}^{(X,Y)} = \sum_{\xi=1}^{\ell} w_{\xi}(\lambda_i^{(X)}) w_{\xi}(\lambda_j^{(Y)}), \quad (4.2)$$

for $i = 1, \dots, M$ and $j = 1, \dots, M$, between diffusion harmonics of X and Y . As shown in the following lemma, whose proof appears in Appendix B, these weights enable us to quantitatively identify diffusion harmonics that correspond similar frequencies and ignore relations between ones that have significantly different ones.

LEMMA 4.1.—*The bandlimiting weights from Eqn. 4.2 satisfy the following properties:*

$w_{ij}^{(X,Y)}$ is continuous and differentiable

$\lambda_i^{(X)}$ and $\lambda_j^{(Y)}$; if $\lambda_i^{(X)} = \lambda_j^{(Y)}$ then $w_{ij} = 1$; if $|\lambda_i^{(X)} - \lambda_j^{(Y)}| > \frac{2}{\ell}$ then $w_{ij} = 0$; and the rate of change of $w_{ij}^{(X,Y)}$ w.r.t. $|\lambda_i^{(X)} - \lambda_j^{(Y)}|$ is bounded by $O(\ell)$.

Next, we use the weights from Eqn. 4.2 to construct a $M \times N$ bandlimited correlation matrix C defined as

$$[C]_{ij} = w_{ij}^{(X,Y)} \text{corr}(\phi_i^{(X)}, \phi_j^{(Y)}) \quad (4.3)$$

for $i = 1, \dots, M$ and $j = 1, \dots, N$, which only considers correlations between diffusion harmonics within similar frequency bands.

Finally, for each i, j with nonzero weight $w_{ij}^{(X, Y)}$, we now need to compute a correlation between the diffusion harmonics $\phi_i^{(X)}, \phi_j^{(Y)}$. If we had partial data point correspondence, as is assumed in many previous work (e.g., [5,7]), we could estimate such correlation directly from matching parts of the two datasets. However, in our case we do not assume any a priori matching between data points. Instead, we rely on the assumed feature correspondence and leverage the GFT from §3.2 to express the harmonics $\phi_i^{(X)}, \phi_j^{(Y)}$ in terms of the data features via their Fourier coefficients. Namely, we take the GFT of the data features $f_s, g_s, s = 1, \dots, n$ (i.e., the ‘‘columns’’ of the points-by-features representation of X, Y as $M \times n, N \times n$ data matrices, correspondingly), and use them to represent $\phi_i^{(X)}, \phi_j^{(Y)}$ by the n dimensional vectors $\hat{\mathbf{x}}_i = (\hat{f}_1[i], \dots, \hat{f}_n[i])^T$ and $\hat{\mathbf{y}}_j = (\hat{g}_1[j], \dots, \hat{g}_n[j])^T$, correspondingly. Then, we compute a correlation between the harmonics $\phi_i^{(X)}, \phi_j^{(Y)}$ indirectly via a correlation between $\hat{\mathbf{x}}_i, \hat{\mathbf{y}}_j$. For simplicity, and by slight abuse of terminology, we use an inner product in lieu of the latter, to define

$$\text{corr}(\phi_i^{(X)}, \phi_j^{(Y)}) = \langle \hat{\mathbf{x}}_i, \hat{\mathbf{y}}_j \rangle.$$

Therefore, together with Eqn. 4.3, our bandlimited correlation matrix is given by

$$[C]_{ij} = w_{ij}^{(X, Y)} \langle \hat{\mathbf{x}}_i, \hat{\mathbf{y}}_j \rangle = \sum_{\xi=1}^{\ell} w_{\xi} \left(\lambda_i^{(X)} \hat{\mathbf{x}}_i, w_{\xi} \left(\lambda_j^{(Y)} \hat{\mathbf{y}}_j \right) \right).$$

4.2 Rigid alignment

Given the bandlimited correlation matrix C , we use its SVD given by $C = U\Sigma V^T$ to obtain its nearest orthogonal approximation $\mathbf{T} = UV^T$ (e.g., as shown in [23]) that defines an isometric transformation between the diffusion maps of the two samples, which we refer to as harmonic alignment. Finally, we can now compute a unified diffusion map, which can be written in (block) matrix form as

$$\Phi_t^{(X, Y)} = \begin{bmatrix} \Phi_0^{(X)} & \Phi_0^{(X)}\mathbf{T} \\ \Phi_0^{(Y)}\mathbf{T}^T & \Phi_0^{(Y)} \end{bmatrix} \begin{bmatrix} \Lambda^{(X)} & 0 \\ 0 & \Lambda^{(Y)} \end{bmatrix}^t, \quad (4.4)$$

where $\Lambda^{(X)}, \Lambda^{(Y)}$ are diagonal matrices with the diffusion eigenvalues $\{\lambda_i^{(X)}\}_{i=1}^N, \{\lambda_j^{(Y)}\}_{j=1}^M$ (correspondingly) as their main diagonal, and t is an integer diffusion time parameter as in §3.1. A summary of the described steps is presented in Appendix A. While this construction is presented here in terms of two datasets for simplicity, it can naturally be generalized to multiple datasets by considering multiple blocks (rather than the two-by-two block structure in Eqn. 4.4), based on orthogonalizing pairwise bandlimited correlations between datasets. This generalization is discussed in detail in Appendix A.2.

Finally, given aligned DMs in $\Phi_t^{(X, Y)}$, we can construct a new neighborhood kernel over their coordinates (i.e., in terms of a combined diffusion distance) and build a robust unified diffusion geometry over the entire data in $X \cup Y$ that is invariant to batch effects and

also enables denoising of data collection artifacts that depend on environment or technology rather than the underlying measured phenomena. This diffusion geometry can naturally be incorporated in diffusion-based methods for several data processing tasks, such as dimensionality reduction & visualization [24], denoising & imputation [25], latent variable inference [7], and data generation [26]. In particular, in §5.3 we demonstrate the application of harmonic alignment to batch effect removal and multimodal data fusion with various single-cell genomic technologies.

5 Numerical results

5.1 Artificial feature corruption

To demonstrate the accuracy of harmonic alignment, we assess its ability to recover k -nearest neighborhoods after random feature corruption, and compare it to MNN [9] and MAGAN [10], which are leading manifold- and deep-learning methods respectively, as discussed in §3.3. To this end, we drew two random samples X and Y of 1000 MNIST digit images, each of which is a 784-dimensional vector. For each trial, we generate a random orthogonal 784×784 corruption matrix \mathbf{O}_0 . To vary the amount of feature corruption, we produce partial corruption matrices \mathbf{O}_p by randomly substituting $p\%$ of the columns in \mathbf{O}_0 with columns of the identity matrix. Right multiplication of Y by these matrices yields corrupted images with only $p\%$ preserved pixels (Fig. 2(b), ‘Corrupted’).

To assess the alignment of the corrupted images $Y\mathbf{O}_p$ to the uncorrupted images X , we perform lazy classification on digits (i.e., rows) in $Y\mathbf{O}_p$ by using the labels of each aligned image’s k nearest neighbors in X . The results of this experiment, performed for $p = \{0, 5, 10, \dots, 95, 100\}$, are reported in Fig. 2(a). For robustness, at each p we sampled three different non-overlapping pairs X, Y , and for each pair we sampled three random \mathbf{O}_p matrices. It should be noted that while we report results in terms of mean classification accuracy, we do not aim to provide an optimal classifier here. Our evaluation merely aims to provide a quantitative assessment of neighborhood quality before and after alignment. We regard a lazy learner as ideal for such evaluation since it directly exposes the quality of data neighborhoods, rather than obfuscate it via a trained model. For comparison, results for harmonic alignment with a SVM classifier are shown in §5.2 and Fig. 2(d).

In general, none of the methods recovers k -nearest neighborhoods under total corruption, showing 10% accuracy for very small p , essentially giving random chance accuracy. Note that this case clearly violates our (partial) feature correspondence assumption. However, when using sufficiently many band-limited filters, harmonic alignment quickly recovers over 80% accuracy and consistently outperforms both MNN and MAGAN, except under very high correspondence (i.e., when $\mathbf{O}_p \approx \mathbf{I}$). The method proposed by [11] was excluded since it did not show improvement over unaligned classification, but is discussed in supplemental materials for completeness. We note that the performance of harmonic alignment is relatively invariant to the choice of ℓ , with the exclusion of extremely high values. For the remainder of the experiments, we fix $\ell = 8$. All experiments use the default parameter $t = 1$.

Next, we examined the ability of harmonic alignment to reconstruct the corrupted data (Fig. 2(b)). We performed the same corruption procedure with $p = 25\%$ and selected one example of each MNIST digit. Ground truth from Y and corrupted result $Y\mathbf{O}_{25}$ are shown in Fig. 2(b). Then, reconstruction was performed by setting each pixel in a new image to the dominant class average of the 10 nearest neighbors from X . In the unaligned case, we see that most examples give smeared fives or ones; this is likely a random intersection formed by X and $Y\mathbf{O}_{25}$. On the other hand, reconstructions produced by harmonic alignment resemble the original input examples.

Finally, in Fig. 2(c), we consider the effect of data size on obtained alignment. To this end, we fix $p = 35\%$ and vary the size of the two aligned datasets. We compare harmonic alignment, MNN, and MAGAN on input sizes ranging from 200 to 1600 MNIST digits, while again using lazy classification accuracy to measure neighborhood preservation and quantify alignment quality. The results in Fig. 2(d) show that both MNN and MAGAN are not significantly affected by dataset size, and in particular do not improve with additional data. Harmonic alignment, on the other hand, not only outperforms them significantly – its alignment quality increases monotonically with input size.

5.2 Transfer learning

An interesting use of manifold alignment algorithms is transfer learning. In this setting, an algorithm is trained to perform well on a small (e.g., pilot) dataset, and the goal is to extend the algorithm to a new larger dataset (e.g., as more data is being collected) after alignment. In this experiment, we first randomly selected 1,000 uncorrupted examples of MNIST digits, and constructed their DM to use as our training set. Next, we took 65%-corrupted unlabeled points (§5.1) in batches of 1,000, 2,000, 4,000, and 8,000, as a test set for classification using the labels from the uncorrupted examples. As shown in 2(d), with a 5-nearest neighbor lazy classifier, harmonic alignment consistently improves as the dataset gets larger, even with up to *eight* test samples for every one training sample. When the same experiment is performed with a linear SVM, harmonic alignment consistently outperforms other methods with performance being independent of test set size (or train-to-test ratio). This is due to the increased robustness and generalization capabilities of trained SVM. Further discussion of transfer learning is given in the supplementary materials. In addition to showing the use of manifold alignment in transfer learning, this example also demonstrates the robustness of our algorithm to imbalance between samples.

5.3 Biological data

5.3.1 Batch effect correction—To illustrate the need for robust manifold alignment in computational biology, we turn to a simple real-world example from [27] (Fig. 3). This dataset was collected by mass cytometry (CyTOF) of peripheral blood mononuclear cells from patients who contracted dengue fever [27].

The canonical response to dengue infection is upregulation of interferon gamma ($\text{IFN}\gamma$). [28] During early immune response, $\text{IFN}\gamma$ works in tandem with acute phase cytokines such as tumor necrosis factor alpha ($\text{TNF}\alpha$) to induce febrile response and inhibit viral replication [29]. We thus expect to see upregulation of these two cytokines together.

In Fig. 3(a), we show the relationship between IFN γ and TNF α without denoising. Note that there is a substantial difference between the IFN γ distributions of the two samples (Earth Mover's Distance [EMD] = 2.699). In order to identify meaningful relationships in CyTOF data, it is common to denoise it first [1]. We used a graph low-pass filter proposed in [25] to denoise the cytokine data. The results of this denoising are shown in Fig. 3(b). This procedure introduced more technical artifacts by enhancing differences between batches, as seen by the increased EMD (3.127) between IFN γ distributions of both patients. This is likely due to substantial connectivity differences between the two batch submanifold in combined data manifold.

Next, we performed harmonic alignment of the two patient profiles (Fig. 3(c)). Harmonic alignment corrected the difference between IFN γ distributions and restored the canonical correlation of IFN γ and TNF α (EMD=0.135). This example illustrates the utility of harmonic alignment for biological data, where it can be used for integrated analysis of data collected across different experiments, patients, and time points.

5.3.2 Multimodal Data Fusion—Since cells contain numerous types of components that are informative of their state (genes, proteins, epigenetics), modern experimental technologies are starting to measure of each of these components separately at the single cell level. Since most single-cell assays are destructive, it is challenging or impossible to obtain all desired measurements in the same cells. It is therefore desirable to perform each assay on a subset of cells from a single sample, and align these datasets *in silico* to obtain a pseudo-joint profile of the multiple data types.

To demonstrate the utility of harmonic alignment in this setting, we use a dataset obtained from [30] of 11,296 cells from adult mouse kidney collected by a joint measurement technique named sci-CAR, which measures *both* gene expression (scRNA-seq) and chromatin accessibility (scATAC-seq) in the same cells simultaneously. The datasets are normalized separately as in [25], using a square root transformation for the scRNA-seq and a log transformation with a pseudocount of 1 for the scATAC-seq data, and finally the dimensionality of each dataset is reduced to 100 using truncated SVD. After randomly permuting the datasets to scramble the correspondence between them, we align the two manifolds in order to recover the known bijection between data modalities. Let $f(i) \in F$ be the scRNA-seq measurement of cell i , and $g(i) \in G$ be the scATAC-seq measurement of cell i . Fig. 3(d) shows the average percentage overlap of neighborhoods of $f(i)$ in F with neighborhoods of $g(i)$ in G , before and after alignment with: MAGAN, MNN and Harmonic Alignment. Harmonic Alignment most accurately recovers cell neighborhoods, thereby allowing the generation of *in silico* joint profiles across data types and obviating the need for expensive or infeasible *in vitro* joint profiling.

6 Conclusion

We presented a novel method for processing multisample data, which contains multiple sampled datasets that differ by global shifts or *batch effects*. To perform data fusion and provide a single stable representation of the entire data, we proposed to learn an intrinsic diffusion geometry of each individual datasets and then align them together using the duality

between diffusion coordinates used in manifold learning and manifold harmonics used in graph signal processing. While previous methods for data manifold alignment relied on known bijective correspondence between data points, our method replaces such strict requirement by considering feature correspondence in the sense that corresponding features across samples or datasets should have similar intrinsic regularity (or “frequency” composition) on the diffusion geometry of each sampled dataset in the data. Our *harmonic alignment* leverages this understanding to compute cross-dataset similarity between manifold harmonics, which is then used to construct an isometric transformation that aligns the data manifolds. Results show that our method is effective in resolving both artificial misalignment and biological batch effects, thus allowing data fusion and transfer learning. We expect future applications of harmonic alignment to include, for example, the use of multimodal data fusion to understand complex molecular processes through three or more different data modalities.

Acknowledgments

This work was partially funded by: the Gruber Foundation [S.G.]; IVADO (l’institut de valorisation des données) [G.W.]; Chan-Zuckerberg Initiative grants 182702 & CZF2019-002440 [S.K.]; and NIH grants R01GM135929 & R01GM130847 [G.W.,S.K.].

A: Algorithm

A.1 Standard implementation

Consider the standard setting of harmonic alignment as described in Section 4. We provide here a detailed description of the Harmonic Alignment algorithm in pseudocode.

Let X and Y be collections of data points $\vec{x} \in \mathbb{R}^d$. Let the diffusion time t and band count ℓ be positive integers. Let $\sigma^{(X)}, \sigma^{(Y)}$ be bandwidth functions (generally either positive constants or a function of the distance from a point to its k th nearest neighbor) and the anisotropy $q^{(X)}, q^{(Y)} \in [0, 1]$. Let the kernel parameters $\mathcal{K}^{(i)} = \{\sigma^{(i)}, q^{(i)}\}$.

Then the aligned data is given by $\Phi_t^{(X,Y)} = \text{HarmonicAlignment}(X, Y, \{\mathcal{K}^{(X)}, \mathcal{K}^{(Y)}\}, t, \ell)$, where the first $\#|X|$ points in $\Phi_t^{(X,Y)}$ are the aligned coordinates of X and the last $\#|Y|$ points in $\Phi_t^{(X,Y)}$ are the aligned coordinates of Y .

Algorithm A.1

function HarmonicAlignment ($X, Y, \mathcal{K}, t, \ell$) = $\Phi_t^{(X,Y)}$

Input:	Data sets	(X, Y)
	Kernel parameters	$\mathcal{K} = \{\mathcal{K}^{(X)}, \mathcal{K}^{(Y)}\}$
	Alignment diffusion time	t
	Alignment band count	ℓ
Output:	Aligned diffusion map	$\Phi_t^{(X,Y)}$
	1: for $Z=X, Y$ do	

```

2:  $\mathcal{L}^{(Z)} \leftarrow \{W^{(Z)}, \mathcal{L}^{(Z)}, D^{(Z)}\} \leftarrow \text{GaussKernelGraph}(Z, \mathcal{X}^{(Z)}) \quad \triangleright \text{Alg. A.2}$ 
3:  $\Psi^{(Z)}, \Lambda^{(Z)} \leftarrow \text{SVD}(I - \mathcal{L}^{(Z)}) \quad \triangleright \text{Graph Fourier basis}$ 
4: end for
5:  $\Phi_t^{(X,Y)} \leftarrow \text{Align}(\{\Psi^{(X)}, \Lambda^{(X)}, D^{(X)}\}, \{\Psi^{(Y)}, \Lambda^{(Y)}, D^{(Y)}\}, t, \ell) \quad \triangleright \text{Alg. A.3}$ 
6: return  $\Phi_t^{(X,Y)}$ 

```

Algorithm A.2

function GaussKernelGraph (X, ϵ) = $\{W, D, \mathcal{L}\}$

Apply a Gaussian kernel to a dataset and compute the corresponding graph matrices.

```

Input:   Dataset            $X = \{\vec{x}_1, \dots, \vec{x}_N : x \in \mathcal{X}\} \subseteq \mathbb{R}^d$ 
           Bandwidth function  $\sigma: X \mapsto \mathbb{R}$ 
Output: Kernel matrix      $W$ 
           Degree matrix       $D$ 
           Normalized Laplacian  $\mathcal{L}$ 

1: for  $i = 1, N$  do
2:    $D(i, i) \leftarrow 0$ 
3:   for  $j = 1, N$  do
4:      $W(i, j) \leftarrow \frac{1}{2} \left( \exp\left(\frac{-\|\vec{x}_i - \vec{x}_j\|_2^2}{2\epsilon(x_i)}\right) + \exp\left(\frac{-\|\vec{x}_i - \vec{x}_j\|_2^2}{2\epsilon(x_j)}\right) \right) \quad \triangleright \text{Symmetric kernel}$ 
5:      $D(i, i) \leftarrow D(i, i) + W(i, j) \quad \triangleright \text{Degrees}$ 
6:   end for
7: end for
8:  $\mathcal{L} \leftarrow I - D^{-1/2} W D^{-1/2} \quad \triangleright \text{Normalized graph Laplacian}$ 
9: return  $\{W, D, \mathcal{L}\}$ 

```

A.2 Multiple dataset alignment

While the previous construction was presented in terms of two datasets for simplicity, we can naturally generalize harmonic alignment to n datasets by considering multiple blocks (rather than the two-by-two block structure in 4.4), based on orthogonalizing pairwise bandlimited correlations between datasets. We briefly elaborate this approach.

Consider $\mathcal{X} = \{\mathcal{X}^{(i)} \subset \mathbb{R}^d : \#\mathcal{X}^{(i)} = N_i\}_i^n$. As in the case with $n = 2$ datasets, we apply a Gaussian kernel to each $\mathcal{X}^{(i)} \in \mathcal{X}$, which yields a Laplacian $\mathcal{L}^{(i)}$ and degree matrix $D^{(i)}$. Diagonalization of $I - \mathcal{L}^{(i)}$ produces a Fourier basis $\{\Psi^{(i)}, \Lambda^{(i)}\}$. As before, we consider the eigenspace $\Psi^{(i)} = \Psi^{(i)} \setminus \psi_1^{(i)}, \Lambda^{(i)} = \Lambda^{(i)} \setminus \lambda_1^{(i)}$, from which

Algorithm A.3

function Align ($\mathcal{G}^{(X)}, \mathcal{G}^{(Y)}, t, \ell$) = $\Phi_t(X, Y)$

Compute and apply an alignment matrix to two diffusion maps.

Input: Laplacian eigensystems and degrees $\mathcal{G}^{(X)}, \mathcal{G}^{(Y)}; \mathcal{G}^{(Z)} = \{\Psi^{(Z)}, \Lambda^{(Z)}, D^{(Z)}\}$
 Alignment diffusion t
 Alignment band count ℓ

Output: Aligned diffusion map $\Phi_t(X, Y)$

- 1: **for** $Z=X, Y$ **do**
- 2: $\Psi^{(Z)} \leftarrow \Psi^{(Z)} \setminus \psi_1^{(Z)}; \Lambda^{(Z)} \leftarrow \Lambda^{(Z)} \setminus \lambda_1^{(Z)}$
- 3: $\Phi^{0(Z)} \leftarrow D^{(Z)^{1/2}} \Psi^{(Z)}$ ▷ Graph Fourier transform
- 4: $\hat{Z} \leftarrow \Psi^{(Z)^T} Z$
- 5: **end for**
- 6: $w^{(X, Y)} \leftarrow$ Bandlimiting Weights ($\Lambda^{(X)}, \Lambda^{(Y)}, \ell$) ▷ Alg. A.4
- 7: **for** $i=2, N_1$ **do**
- 8: **for** $j=2, N_2$ **do**
- 9: $C(i-1, j-1) \leftarrow w_{ij}(X, Y) \langle \hat{X}(i-1, :), \hat{Y}(j-1, :) \rangle$ ▷ Bandlimited correlatios
- 10: **end for**
- 11: **end for**
- 12: $U, S, V \leftarrow$ SVD(C)
- 13: $T \leftarrow UV^T$ ▷ Orthogonalization Sec. 4.2
- 14: $\Phi_t^{(X, Y)} \leftarrow \begin{bmatrix} \Phi^{0(X)} & \Phi^{0(X)} \mathbf{T} \\ \Phi^{0(Y)} \mathbf{T}^T & \Phi^{0(Y)} \end{bmatrix} \begin{bmatrix} \Lambda^{(X)} & 0 \\ 0 & \Lambda^{(Y)} \end{bmatrix}^t$
- 15: **return** $\Phi_t(\bar{X}, Y)$

Algorithm A.4

function BandlimitingWeights ($\Lambda^{(X)}, \Lambda^{(Y)}, \ell$) = $w^{(X, Y)}$

Compute the joint bandlimiting weights for two graphs.

Input: Normalized Laplacian eigenvalues $\Lambda = \left\{ \Lambda^{(X)}, \Lambda^{(Y)}; \Lambda^{(Z)} = \left\{ \lambda_j^{(Z)} \right\}_{j=2}^{\#|Z|} \right\}$
 Alignment band count ℓ

Output: Pairwise frequency weights $w^{(X, Y)}: \Lambda^{(X)} \times \Lambda^{(Y)} \mapsto [0, 1]$

- 1: $f \leftarrow f: \lambda, \ell, \xi \mapsto I(\xi - 1 \leq \lambda \ell \leq \xi + 1) \sin\left(\frac{\pi}{2} \cos^2\left(\frac{\pi}{Y}(\ell \lambda - \xi)\right)\right)$ ▷ Itersine wavele 4.1
- 2: **for** $i=2, N_2$ **do**
- 3: **for** $j=2, N_2$ **do**

```

4:      $w_{ij}(X, Y) \leftarrow 0$ 
5:     for  $\xi = 1, \ell$  do
6:          $w_{ij}(X, Y) \leftarrow w_{ij}(X, Y) + f(\lambda_i(X), \ell, \xi)f(\lambda_j(Y), \ell, \xi)$ 
7:     end for
8: end for
9: end for
10: return  $w^{(X, Y)}$ 

```

we compute a graph Fourier transform $\widehat{X}^{(i)} = \Psi^{(i)T} X^{(i)}$ and diffusion map $\Phi_0^{(i)} = D^{(i)1/2} \Psi^{(i)}$ for each dataset. This is the same initialization that one performs for the case when $n = 2$ as in Algs. A.1, A.3.

Next, for every pair $i, j \in \{1, \dots, n\} \times \{1, \dots, n\}$ we generate bandlimiting weights $w^{(i,j)} = \text{BandlimitingWeights}(\Lambda^{(X)}, \Lambda^{(Y)}, \ell)$ (Alg. A.4). Then the correlation between each diffusion map pair is

$$C^{(i,j)}(h-1, k-1) = w^{(i,j)}(\lambda_h^{(i)}, \lambda_k^{(j)}) \left\langle \widehat{X}^{(i)}(h-1, :), \widehat{X}^{(j)}(k-1, :) \right\rangle_{k=2, \dots, N_j}^{h=2, \dots, N_i}$$

Factoring $C^{(i,j)} = U^{(i,j)} S^{(i,j)} V^{(i,j)T}$, we have the rigid alignment operator

$$T^{(i \rightarrow j)} = U^{(i,j)} V^{(i,j)T}$$

and its adjoint $T^{(j \rightarrow i)} = T^{(i \rightarrow j)T}$.

Next, let $B^{(1,2,\dots,n)}(i,j)$ be an $N_i - 1 \times N_j - 1$ matrix such that

$$B^{(1,2,\dots,n)}(i,j) = \begin{cases} \Phi_0^{(i)} & i = j \\ \Phi_0^{(i)} T^{(i \rightarrow j)} & j > i \\ \Phi_0^{(i)} T^{(j \rightarrow i)} & j < i. \end{cases}$$

Then the i, j block $\Phi_i^{(1,2,\dots,n)}(i,j)$ is the $N_i - 1 \times N_j - 1$ submatrix of the diffusion coordinates of $\mathcal{X}^{(i)}$ aligned into the diffusion space of $\mathcal{X}^{(i)}$,

$$\Phi_i^{(1,2,\dots,n)}(i,j) = B^{(1,2,\dots,n)}(i,j) \Lambda^{(j)T}.$$

This matrix is

$$\Phi_t^{(1,2,\dots,n)} = \begin{bmatrix} \Phi_0^{(1)} & \Phi_0^{(1)T(1 \rightarrow 2)} & \dots & \Phi_0^{(t)T(1 \rightarrow n)} \\ \Phi_0^{(2)T(2 \rightarrow 1)} & \Phi_0^{(2)} & \dots & \Phi_0^{(2)T(2 \rightarrow n)} \\ \vdots & \vdots & \ddots & \vdots \\ \Phi_0^{(2)T(n \rightarrow 1)} & \Phi_0^{(2)T(n \rightarrow 2)} & \dots & \Phi^{(n)0} \end{bmatrix} \begin{bmatrix} \Lambda^{(X)} \\ \Lambda^{(2)} \\ \vdots \\ \Lambda^{(n)} \end{bmatrix}^t.$$

Alg. A.5 summarizes this process. When $n = 2$, it simplifies to Alg. A.1.

A.3 Runtime analysis and implementation

Here we provide an informal analysis of algorithmic runtime and suggest a collection of possible improvements that could be made in order to increase the efficiency of our algorithm. We show that the algorithm with cubic sample complexity and linear feature complexity in its naïve implementation, and can be reduced to quadratic sample complexity with relatively straightforward modifications. Breaking the problem down via a divide-and-conquer approach could yield further improvements to give an algorithm with linear sample complexity. We note that in all experiments in this paper, the naïve approach was used.

The runtime complexity of a naïve implementation of Alg. A.1 is

$$\mathcal{O} \left(\underbrace{N_1^3 + N_2^3}_{\text{(a) embedding}} + \underbrace{(N_1^2 + N_2^2)d}_{\text{(b) GFT}} + \underbrace{N_1 N_2 (d + N_1)}_{\text{(c) correlation \& SVD}} + \underbrace{N_1 N_2 (N_1 + N_2)}_{\text{(d) alignment}} \right) \subset \mathcal{O}(N_2^3 + N_1 N_2 d),$$

where $N_1 < N_2$ are the size of two data sets to align and d is the number of dimensions. The major costs of the proposed algorithm are partitioned according to their step in the algorithm (see underbraces).

Some simple observations about the rank of each system will pave the way to reducing alignment runtime. Our primary tool will be randomized truncated SVD, e.g. [31,32], which computes the first k singular vectors of an $m \times n$ system in $\mathcal{O}(mn \log k)$.

First we reduce the size of the input data. Assuming that the features of $X^{(j)}$ are independent, the simplest setting for dimensionality and rank reduction occurs when $N_1 < d$ (recalling that $N_1 < N_2$). It is clear by construction that the rank of the correlation matrix is at most

$$r_{\max} = \min\{d, N_1\}.$$

Algorithm A.5

function MultiAlignment ($\mathcal{X}, \mathcal{X}, t, \ell$) = $\Phi_t^{(1,\dots,n)}$

Align the diffusion maps of multiple datasets.

Input: Data sets $\mathcal{X} = \{\mathcal{X}^{(1)}, \dots, \mathcal{X}^{(n)}\}$

	Kernel parameters	$\mathcal{K} = \{\mathcal{K}^{(X)}, \dots, \mathcal{K}^{(n)}\}$	
	Alignment diffusion time	t	
	Alignment band count	ℓ	
Output:	Unified diffusion map	$\Phi_t^{(1, \dots, n)}$	


```

1:  for i=1,n do
2:     $\mathcal{G}^{(i)} \leftarrow \{W^{(i)}, \mathcal{L}^{(i)}, D^{(i)}\} \leftarrow \text{GaussKernelGraph}(\mathcal{X}^{(i)}, \mathcal{K}^{(i)})$  ▷ Alg. A.2
3:     $\Psi^{(i)}, \Lambda^{(i)} \leftarrow \text{SVD}(I - \mathcal{L}^{(i)})$  ▷ Graph Fourier basis
4:     $\Psi^{(i)} \leftarrow \Psi^{(i)} \setminus \psi_1^{(i)}; \Lambda^{(i)} \leftarrow \Lambda^{(i)} \setminus \lambda_1^{(i)}$ 
5:     $\widehat{\mathcal{X}}^{(i)} \leftarrow \Psi^{(i)T} \mathcal{X}^{(i)}$  ▷ Graph Fourier transform
6:     $\Phi_0^{(i)} \leftarrow D^{(i)1/2} \Psi^{(i)}$ 
7:  end for
8:  for i = 1, n do
9:     $\Phi_t^{(i, i)}(1, \dots, n) \leftarrow \Phi_0^{(i)} \Lambda^{(i)t}$ 
10:   for j = i + 1, n do
11:      $w^{(i, j)} \leftarrow \text{BandlimitingWeights}(\Lambda^{(i)}, \Lambda^{(j)}, \ell)$  ▷ Alg. A.4
12:     for  $\ell = 2, N_j$  do
13:       for k = 2,  $N_j$  do.
14:          $C^{(i, j)}(\ell - 1, k - 1) \leftarrow w^{(i, j)}(\lambda_\ell^{(i)}, \lambda_k^{(j)}) \left( \widehat{\mathcal{X}}^{(i)}(\ell - 1, :), \widehat{\mathcal{X}}^{(j)}(k - 1, :) \right)$  ▷ Bandlimited correlation
15:       end for
16:     end for
17:      $U^{(i, j)}, V^{(i, j)} \leftarrow C^{(i, j)}$ 
18:      $T^{(i \rightarrow j)} \leftarrow U^{(i, j)} V^{(i, j)T}; T^{(j \rightarrow i)} \leftarrow V^{(i, j)} U^{(i, j)T}$ 
19:      $\Phi_t^{(i, j)}(1, \dots, n) \leftarrow \Phi_0^{(i)} T^{(i \rightarrow j)} \Lambda^{(j)t}$ 
20:      $\Phi_t^{(j, i)}(1, \dots, n) \leftarrow \Phi_0^{(j)} T^{(j \rightarrow i)} \Lambda^{(i)t}$ 
21:   end for
22: end for
23: return  $\Phi_t^{(1, \dots, n)}$ 

```

Thus, if the input data $X^{(i)} \subset \mathbb{R}^d$ for $d > r_{\max}$, then the most efficient algorithm will use truncated SVD to only consider the first r_{\max} principal components of each $X^{(i)}$ as alignment features.

Applying rank reduction to only the input data introduces an additional $\log r_{\max}$ term to the embedding (a) through a truncated SVD

$$\mathbf{O}\left(\left(N_1^3 + N_2^3\right) + \left(N_1^2 + N_2^2\right)\log r_{\max}\right).$$

However, the GFT **(b)** now runs in

$$\mathbf{O}\left(\left(N_1^2 + N_2^2\right)r_{\max}\right).$$

Subsequently, correlation and orthogonalization **(c)** runs in

$$\mathbf{O}\left(N_1 N_2 (r_{\max} + N_1)\right).$$

where the r_{\max} term is due to the product of a $N_1 \times r_{\max}$ matrix with a $r_{\max} \times N_2$ matrix and the second is the full SVD of the $N_1 \times N_2$ correlation matrix.

The same argument can be applied to reduce the number of diffusion coordinates such that r_{\max} components are taken. The correlation **(c)** then collapses to two r_{\max}^3 operations: one is the matrix product of two square $r_{\max} \times r_{\max}$ matrices and the other is the full SVD of this product. This reduces the total complexity of harmonic alignment to

$$\begin{aligned} & \mathbf{O}\left(\underbrace{\left(N_1^2 + N_2^2\right)3\log r_{\max}}_{\text{(a) embedding}} + \underbrace{\left(N_1 + N_2\right)r_{\max}^2}_{\text{(b) GFT}} + \underbrace{2r_{\max}^3}_{\text{(c) correlation \& SVD}} + \underbrace{N_1 N_2 (2r_{\max})}_{\text{(d) alignment}}\right) \\ & \subset \mathbf{O}\left(N_2^2 \log r_{\max} + N_1 N_2 r_{\max}\right). \end{aligned}$$

It is often the case that one is only interested in $k \ll r_{\max}$ PCA components or diffusion components. For example, [33] proves an optimal singular value cutoff for denoising of $m \times n$ data matrices. One could select a different rank for each PCA and diffusion maps operation; we will denote these as $\{k_{PCA}^{(1)}, k_{PCA}^{(2)}\}$ and $\{k_{DM}^{(1)}, k_{DM}^{(2)}\}$, simplifying to $k_{PCA} = \max\{k_{PCA}^{(1)}, k_{PCA}^{(2)}\}$, $k_{DM} = \max\{k_{DM}^{(1)}, k_{DM}^{(2)}\}$ and $k = \max\{k_{PCA}, k_{DM}\}$.

The total complexity of harmonic alignment is now

$$\mathbf{O}\left(\underbrace{\left(N_1^2 + N_2^2\right)3\log k_{PCA}}_{\text{(a) embedding}} + \underbrace{\left(N_1 + N_2\right)k_{PCA}^2}_{\text{(b) GFT}} + \underbrace{2k_{DM}^3}_{\text{(c) corr. \& SVD}} + \underbrace{2N_1 N_2 k_{DM}}_{\text{(d) alignment}}\right) \subset \mathbf{O}\left(N_2^2 \log k + N_1 N_2 k\right).$$

Finally, we can use the multiple alignment algorithm presented in Alg. A.5 to ‘chunk’ very large datasets to reduce runtime. The general scheme would be to break the input into many smaller datasets of N_c points. The result is a set of $\sum_i \lfloor N_i / N_c \rfloor \approx \frac{N_1 + N_2}{N_c}$ smaller problems that run in

$$\mathbf{O} \left(\frac{N_1 + N_2}{N_c} \left(\underbrace{6N_c \log k}_{\text{(a) embedding } PCA} + \underbrace{2N_c k^2}_{\text{(b) GFT } PCA} + \underbrace{2k^3}_{\text{(c) corr SVD } DM} + \underbrace{2N_c^2 k}_{\text{(d) alignment } DM} \right) \right) \subset \mathbf{O}(N_2 N_c k),$$

Further analysis must be done to examine the effect on the accuracy of the output when one divides-andconquers in this way. It is clear that in practice it is important to accordingly adjust k as the rank structure will vary depending on the sizes of the submatrices chosen to align.

B: Proof of Lemma 4.1

The bandlimiting weights from (4.2) satisfy the following properties:

CLAIM B.1.

$w_\xi(\lambda)$ is continuous in λ .

Proof. As a piecewise function of continuous functions, it suffices to check that $w_\xi(\lambda)$ is continuous at $\lambda = \frac{\epsilon - 1}{\ell}$ and $\lambda = \frac{\xi + 1}{\ell}$.

$$\begin{aligned} \lim_{h \rightarrow 0^+} w_\xi \left(\frac{\xi \pm 1}{\ell} \mp h \right) &= \lim_{h \rightarrow 0^+} \left(\sin \left(\frac{\pi}{Y} \cos^2 \left(\frac{\pi}{Y} \left(\ell \left(\frac{\xi \pm 1}{\ell} \mp h \right) - \xi \right) \right) \right) \right) \\ &= \sin \left(\frac{\pi}{Y} \cos^2 \left(\frac{\pi}{Y} \left(\ell \left(\frac{\xi \pm 1}{\ell} \right) - \xi \right) \right) \right) \\ &= \sin \left(\frac{\pi}{Y} \cos^2 \left(\pm \frac{\pi}{Y} \right) \right) = 0 = \lim_{h \rightarrow 0^-} w_\xi \left(\frac{\xi \pm 1}{\ell} \mp h \right). \end{aligned}$$

COROLLARY B.1.

As a sum of continuous functions in $\lambda_i^{(X)}$ and $\lambda_j^{(Y)}$, $w_{ij}^{(X,Y)}$ is continuous in $\lambda_i^{(X)}$ and $\lambda_j^{(Y)}$.

CLAIM B.2.

$w_\xi(\lambda)$ is differentiable in λ .

Proof. Note first that

$$\frac{dw_\xi}{d\lambda} = \begin{cases} -\frac{\pi^2 \ell}{Y} \cos \left(\frac{\pi}{Y} (\ell \lambda - \xi) \right) \cos \left(\frac{\pi}{Y} \cos^2 \left(\frac{\pi}{Y} (\ell \lambda - \xi) \right) \right) \sin \left(\frac{\pi}{Y} (\ell \lambda - \xi) \right) & \frac{\xi - 1}{\ell} \leq \lambda \leq \frac{\xi + 1}{\ell}; \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.1})$$

$$= \begin{cases} -\frac{\pi^2 \ell}{4} \sin(\pi(\ell \lambda - \xi)) \cos\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}(\ell \lambda - \xi)\right)\right) \frac{\xi - 1}{\ell} \leq \lambda \leq \frac{\xi + 1}{\ell}; \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.2})$$

Then, as a piecewise function of continuous functions, it suffices to check that $\frac{dw_\xi}{d\lambda}$ is continuous at $\lambda = \frac{\xi - 1}{\ell}$ and $\lambda = \frac{\xi + 1}{\ell}$.

$$\begin{aligned} \lim_{h \rightarrow 0^+} \frac{dw_\xi\left(\frac{\xi \pm 1}{\ell} \mp h\right)}{d\lambda} &= -\frac{\pi^2 \ell}{4} \sin(\pm \pi) \cos\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}\right)\right) \\ &= 0 = \lim_{h \rightarrow 0^-} \frac{dw_\xi\left(\frac{\xi \pm 1}{\ell} \mp h\right)}{d\lambda}. \end{aligned}$$

□

COROLLARY B.2.

As a sum of differentiable functions in $\lambda_i^{(X)}$ and $\lambda_j^{(Y)}$, $w_{ij}^{(X,Y)}$ is differentiable in $\lambda_i^{(X)}$ and $\lambda_j^{(Y)}$.

CLAIM B.3.

If $k \in \mathbb{Z}$ such that $0 \leq k \leq \ell$ and $\frac{k}{\ell} \leq \lambda \leq \frac{k+1}{\ell}$ then $w_\xi(\lambda) = 0$ for all $\xi \notin \{k, k+1\}$.

Proof. Assume $k \in \mathbb{Z}$ such that

$$0 \leq k \leq \ell \text{ and } \frac{k}{\ell} \leq \lambda \leq \frac{k+1}{\ell}. \text{ Let } \xi \in \mathbb{Z} \text{ such that } \xi \notin \{k, k+1\}.$$

For the case where $\xi > k+1$, then $\xi \geq k+2$ and so $\lambda \leq \frac{\xi-1}{\ell}$.

For the case where $\xi < k$, then $\xi \leq k-1$ and so $\lambda \geq \frac{\xi+1}{\ell}$.

In each case, this implies that $w_\xi(\lambda) = 0$. □

CLAIM B.4.

If $\lambda_i^{(X)} = \lambda_j^{(Y)}$ then $w_{ij} = 1$.

Proof. Assume $\lambda_i^{(X)} = \lambda_j^{(Y)} = \lambda$ and let $k \in \mathbb{Z}$ such that $\frac{k}{\ell} \leq \lambda \leq \frac{k+1}{\ell}$.

Then

$$\begin{aligned}
w_{ij}^{(X,Y)} &= \sum_{\xi=1}^{\ell} w_{\xi}(\lambda) w_{\xi}(\lambda) \\
&= w_k(\lambda) w_k(\lambda) + w_{k+1}(\lambda) w_{k+1}(\lambda) \text{ [by Claim [B.3]} \\
&= \sin^2\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}(\ell\lambda - k)\right)\right) + \sin^2\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}(\ell\lambda - k - 1)\right)\right) \\
&= \sin^2\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}(\ell\lambda - k)\right)\right) + \sin^2\left(\frac{\pi}{Y} \sin^2\left(\frac{\pi}{Y}(\ell\lambda - k)\right)\right) \left[\text{since } \cos\left(x - \frac{\pi}{Y}\right) = \sin(x)\right] \\
&= \sin^2\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}(\ell\lambda - k)\right)\right) + \sin^2\left(\frac{\pi}{Y} \left(1 - \cos^2\left(\frac{\pi}{Y}(\ell\lambda - k)\right)\right)\right) \left[\text{since } \sin^2(x) + \cos^2(x) = 1\right] \\
&= \sin^2\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}(\ell\lambda - k)\right)\right) + \cos^2\left(\frac{\pi}{Y} \cos^2\left(\frac{\pi}{Y}(\ell\lambda - k)\right)\right) \left[\text{since } \sin\left(\frac{\pi}{Y} - x\right) = \cos(x)\right] \\
&= 1.
\end{aligned}$$

□

CLAIM B.5.

If $|\lambda_i^{(X)} - \lambda_j^{(Y)}| \geq \frac{2}{\ell}$ then $w_{ij}^{(X,Y)} = 0$.

Proof. Assume $|\lambda_i^{(X)} - \lambda_j^{(Y)}| \geq \frac{2}{\ell}$ and let $k_i^{(X)}, k_j^{(Y)} \in \mathbb{Z}$ such that

$$\frac{k_i^{(X)}}{\ell} \leq \lambda_i^{(X)} \leq \frac{k_i^{(X)} + 1}{\ell} \text{ and } \frac{k_j^{(Y)}}{\ell} \leq \lambda_j^{(Y)} \leq \frac{k_j^{(Y)} + 1}{\ell}.$$

Assume without loss of generality that $\lambda_i^{(X)} < \lambda_j^{(Y)}$. Then by Claim B.3,

$$\lambda_i^{(X)} \leq \lambda_j^{(Y)} - \frac{2}{\ell} \leq \frac{k_j^{(Y)} + 1}{\ell} - \frac{2}{\ell} = \frac{k_j^{(Y)} - 1}{\ell} \Rightarrow w_{\xi}(\lambda_i^{(X)}) = 0 \text{ for all } \xi \geq k_j^{(Y)}$$

and since $w_{\xi}(\lambda_j^{(Y)}) = 0$ for all $\xi \notin \{k_j^{(Y)}, k_j^{(Y)} + 1\}$, then

$$w_{\xi}(\lambda_i^{(X)}) w_{\xi}(\lambda_j^{(Y)}) = 0 \text{ for all } \xi.$$

□

CLAIM B.6.

The rate of change of $w_{ij}^{(X,Y)}$ w.r.t. both $\lambda_i^{(X)}$ and $\lambda_j^{(Y)}$ is bounded by $O(\ell)$.

Proof. From Equation B.2,

$$\left| \frac{dw_{\xi}}{d\lambda} \right| \leq \frac{\pi^2 \ell}{4} \text{ for all } \lambda.$$

Without loss of generality we consider only $\frac{dw_{ij}^{(X,Y)}}{d\lambda_i^{(X)}}$. Let $k \in \mathbb{Z}$ such that $\frac{k}{\ell} \leq \lambda_i^{(X)} \leq \frac{k+1}{\ell}$.

$$\begin{aligned}
\frac{dw_{ij}^{(X,Y)}}{d\lambda_i^{(X)}} &= \sum_{\xi=1}^{\ell} \frac{dw_{\xi}}{d\lambda_i^{(X)}} w_{\xi}(\lambda_j^{(Y)}) \\
&= \frac{dw_k}{d\lambda_i^{(X)}} w_k(\lambda_j^{(Y)}) + \frac{dw_{k+1}}{d\lambda_i^{(X)}} w_{k+1}(\lambda_j^{(Y)}) \\
&\leq \frac{dw_k}{d\lambda_i^{(X)}} + \frac{dw_{k+1}}{d\lambda_i^{(X)}}.
\end{aligned}$$

So

$$\left| \frac{dw_{ij}^{(X,Y)}}{d\lambda_i^{(X)}} \right| \leq \left| \frac{dw_k}{d\lambda_i^{(X)}} + \frac{dw_{k+1}}{d\lambda_i^{(X)}} \right| \leq \left| \frac{dw_k}{d\lambda_i^{(X)}} \right| + \left| \frac{dw_{k+1}}{d\lambda_i^{(X)}} \right| \leq \frac{\pi^2 \ell}{Y}.$$

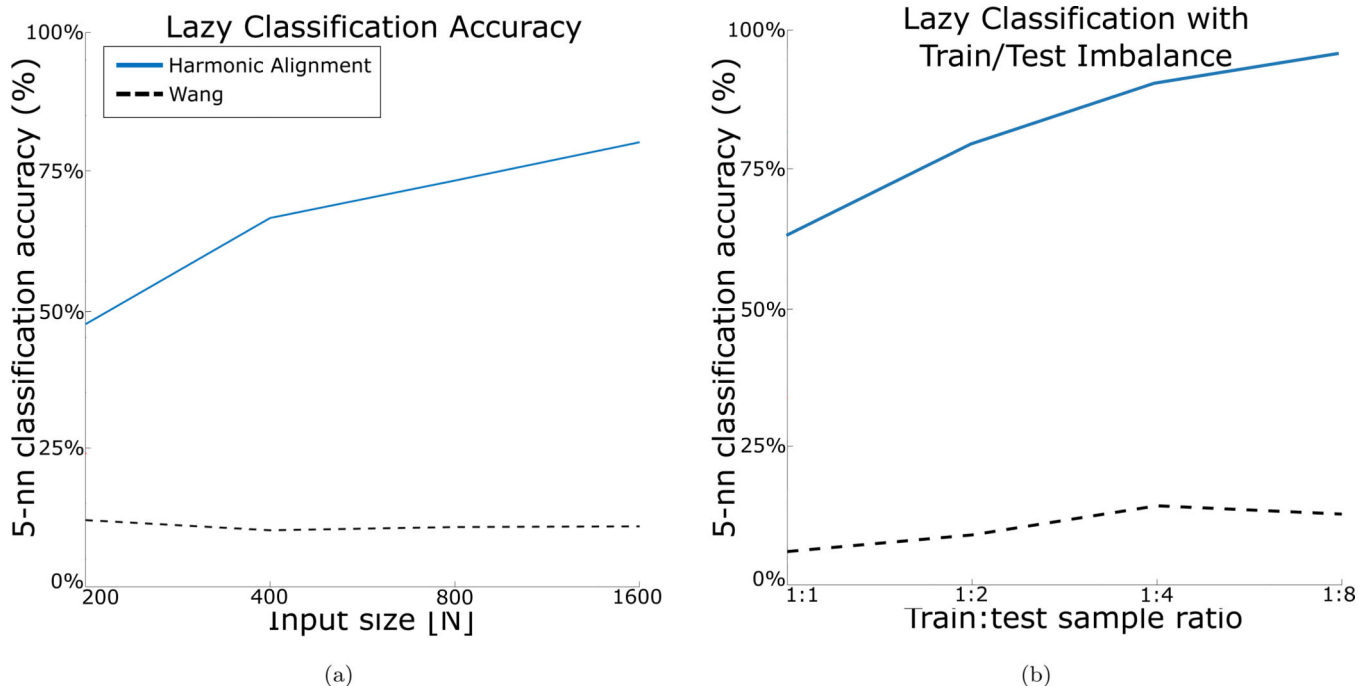
□

COROLLARY B.3.

The rate of change of $w_{ij}^{(X,Y)}$ w.r.t. $|\lambda_i^{(X)} - \lambda_j^{(Y)}|$ is bounded by $O(\ell)$.

C: Comparison to Wang and Mahadevan

Despite being a natural candidate for comparison to our method, unfortunately no standard implementation of the method proposed by [11] is available. Our implementation of their method performed extremely poorly (worse than random) on the comparisons and is extremely computationally intensive. The method is therefore not shown in the main comparisons; however, for completeness, the results are shown in Figure 4.

**Figure 4:**

Recovery of k-neighborhoods under feature corruption. Mean over 3 iterations is reported for each method. (a) Lazy classification accuracy relative to input size with unlabeled randomly corrupted digits with 35% preserved pixels. (b) Transfer learning performance. For each ratio, 1K uncorrupted, labeled digits were sampled from MNIST, and then 1K, 2K, 4K, and 8K (x-axis) unlabeled points were sampled and corrupted with 35% column identity.

References

- [1]. Moon KR, Stanley J, Burkhardt D, van Dijk D, Wolf G, and Krishnaswamy S, “Manifold learning-based methods for analyzing single-cell rna-sequencing data,” *Curr. Opin. in Syst. Biol.*, 2017.
- [2]. Coifman RR and Lafon S, “Diffusion maps,” *Appl. & Comp. Harm. Anal.*, vol. 21, no. 1, pp. 5–30, 2006.
- [3]. Ham J, Lee DD, and Saul LK, “Semisupervised alignment of manifolds.,” in *AISTATS*, pp. 120–127, 2005.
- [4]. Wang C and Mahadevan S, “Manifold alignment using procrustes analysis,” in *Proc. of the 25th ICML*, pp. 1120–1127, 2008.
- [5]. Coifman RR and Hirn MJ, “Diffusion maps for changing data,” *Appl. & Comp. Harm. Anal.*, vol. 36, no. 1, pp. 79–107, 2014.
- [6]. Tuia D and Camps-Valls G, “Kernel manifold alignment for domain adaptation,” *PloS ONE*, vol. 11, no. 2, p. e0148655, 2016.
- [7]. Lederman RR and Talmon R, “Learning the geometry of common latent variables using alternating-diffusion,” *Appl. & Comp. Harm. Anal.*, vol. 44, no. 3, pp. 509–536, 2018.
- [8]. Boumal N, Bendory T, Lederman RR, and Singer A, “Heterogeneous multireference alignment: A single pass approach,” in *Proc. of the 52nd CISS*, pp. 1–6, 2018.
- [9]. Haghverdi L, Lun AT, Morgan MD, and Marioni JC, “Batch effects in single-cell RNA-sequencing data are corrected by matching mutual nearest neighbors,” *Nat. biotechnology*, vol. 36, no. 5, p. 421, 2018.

- [10]. Amodio M and Krishnaswamy S, “MAGAN: Aligning biological manifolds,” in Proc. of the *35th ICML*, pp. 215–223, 2018.
- [11]. Wang C and Mahadevan S, “Manifold alignment without correspondence,” in *IJCAI*, vol. 2, p. 3, 2009.
- [12]. Shuman DI, Narang SK, Frossard P, Ortega A, and Vandergheynst P, “The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains,” *IEEE Sig. Proc. Mag.*, vol. 30, no. 3, pp. 83–98, 2013.
- [13]. Brooks R, Gordon C, and Perry PA, eds., *Geometry of the Spectrum*, vol. 173 of *Contemp. Math* AMS, 1994.
- [14]. Nadler B, Lafon S, Kevrekidis I, and Coifman RR, “Diffusion maps, spectral clustering and eigenfunctions of fokker-planck operators,” in *Adv. in NeurIPS*, pp. 955–962, 2006.
- [15]. Gower JC, “Generalized procrustes analysis,” *Psychometrika*, vol. 40, no. 1, pp. 33–51, 1975.
- [16]. Thompson B, *Canonical correlation analysis: Uses and interpretation*. No. 47 in *Quantitative applications in the social sciences*, Sage, 1984.
- [17]. Zhu J, Park T, Isola P, and Efros AA, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in Proc. of *ICCV*, pp. 2242–2251, 2017.
- [18]. Kim T, Cha M, Kim H, Lee JK, and Kim J, “Learning to discover cross-domain relations with generative adversarial networks,” in Proc. of the *34th ICML*, pp. 1857–1865, 2017.
- [19]. Ovsjanikov M, Ben-Chen M, Solomon J, Butscher A, and Guibas L, “Functional maps: A flexible representation of maps between shapes,” *ACM Trans. Graph.*, vol. 31, pp. 30:1–30:11, 7 2012.
- [20]. Pokrass J, Bronstein AM, Bronstein MM, Sprechmann P, and Sapiro G, “Sparse models for intrinsic shape correspondence,” in *Perspectives in Shape Analysis*, pp. 211–230, Springer, 2016.
- [21]. Vestner M, Litman R, Rodolà E, Bronstein AM, and Cremers D, “Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space,” *CoRR*, vol. abs/1701.00669, 2017.
- [22]. Perraudin N, Holighaus N, Søndergaard PL, and Balazs P, “Designing gabor windows using convex optimization,” arXiv:1401.6033, 2014.
- [23]. Schönemann PH, “A generalized solution of the orthogonal procrustes problem,” *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966.
- [24]. Moon KR, van Dijk D, Wang Z, Burkhardt D, Chen W, van den Elzen A, Hirn MJ, Coifman RR, Ivanova NB, Wolf G, and Krishnaswamy S, “Visualizing transitions and structure for high dimensional data exploration,” bioRxiv, p. 120378, 2017.
- [25]. van Dijk D, Sharma R, Nainys J, Yim K, Kathail P, Carr A, Burdziak C, Moon KR, Chaffer CL, Pattabiraman D, Bieri B, Mazutis L, Wolf G, Smita K, and Pe’er D, “Recovering gene interactions from single-cell data using data diffusion,” *Cell*, vol. 174, no. 3, pp. 716–729.e27, 2018. [PubMed: 29961576]
- [26]. Lindenbaum O, Stanley J, Wolf G, and Krishnaswamy S, “Geometry based data generation,” in *Adv. in NeurIPS*, pp. 1400–1411, 2018.
- [27]. Amodio M, van Dijk D, Srinivasan K, Chen WS, Mohsen H, Moon KR, Campbell A, Zhao Y, Wang X, Venkataswamy M, Desai A, Kumar RV,P, Montgomery R, Wolf G, and Krishnaswamy S, “Exploring single-cell data with deep multitasking neural networks,” *Nat. Methods*, 2019. DOI: 10.1038/s41592-019-0576-7.
- [28]. Chakravarti A and Kumaria R, “Circulating levels of tumour necrosis factor-alpha & interferon-gamma in patients with dengue & dengue haemorrhagic fever during an outbreak,” *Indian J. Med. Res.*, vol. 123, no. 1, p. 25, 2006. [PubMed: 16567864]
- [29]. Ohmori Y, Schreiber RD, and Hamilton TA, “Synergy between interferon- γ and tumor necrosis factor- α in transcriptional activation is mediated by cooperation between signal transducer and activator of transcription 1 and nuclear factor κ b,” *J. Biol. Chem.*, vol. 272, no. 23, pp. 14899–14907, 1997.
- [30]. Cao J, Cusanovich DA, Ramani V, Aghamirzaie D, Pliner HA, Hill AJ, Daza RM, McFaline-Figueroa JL, Packer JS, Christiansen L, Steemers FJ, Adey AC, Trapnell C, and Shendure J, “Joint profiling of chromatin accessibility and gene expression in thousands of single cells,” *Science*, vol. 361, no. 6409, pp. 1380–1385, 2018. [PubMed: 30166440]

- [31]. Rokhlin V, Szlam A, and Tygert M, “A randomized algorithm for principal component analysis,” *SIAM Journal on Matrix Analysis and Applications*, vol. 31, no. 3, pp. 1100–1124, 2010.
- [32]. Halko N, Martinsson P-G, and Tropp JA, “Finding structure with randomness: Stochastic algorithms for constructing approximate matrix decompositions,” 2009.
- [33]. Donoho DL, Gavish M, and Johnstone IM, “Optimal shrinkage of eigenvalues in the spiked covariance model,” *arXiv preprint arXiv:1311.0851*, 2013.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

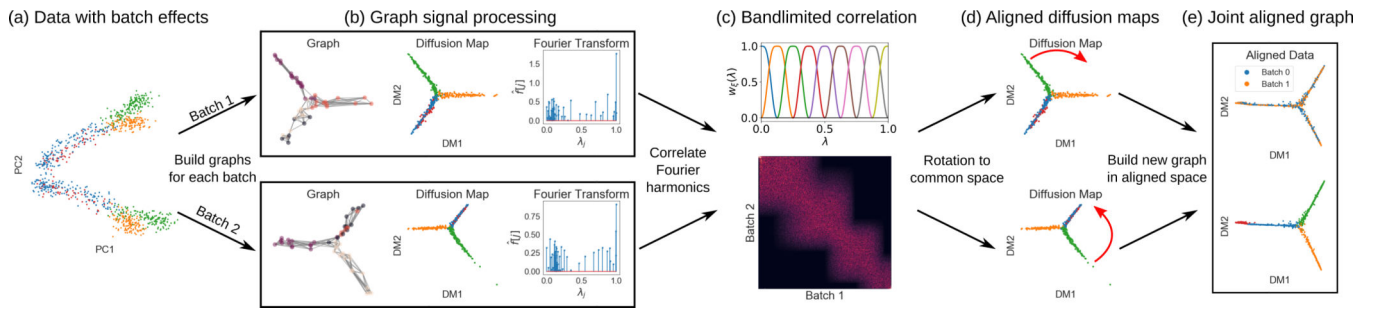

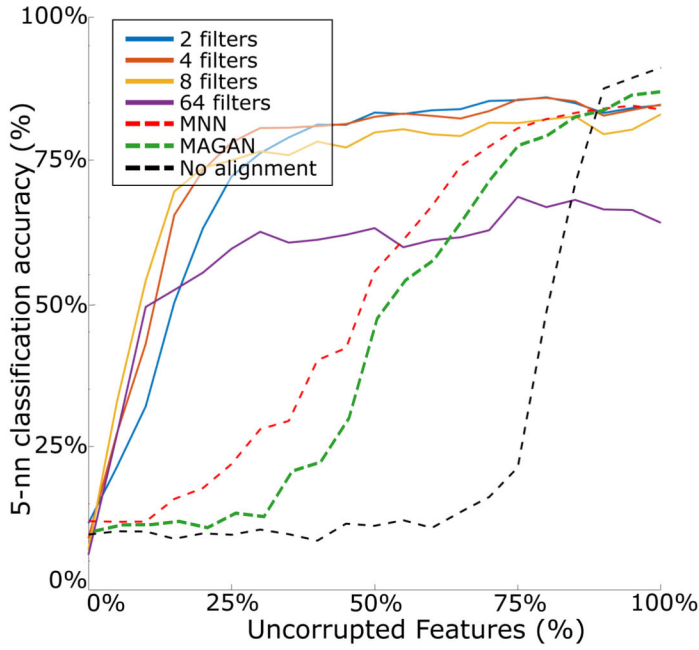
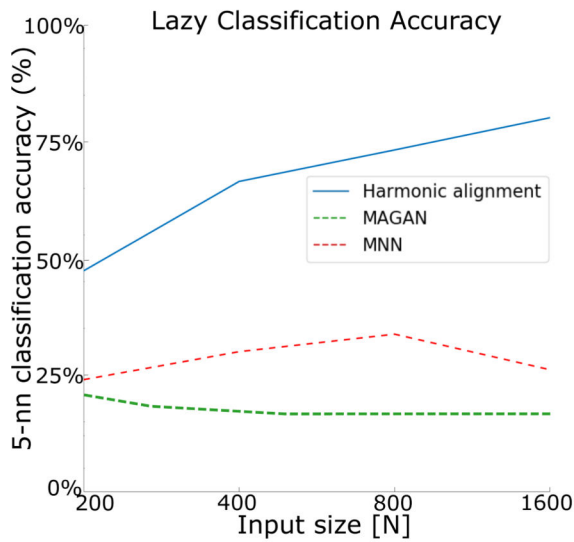


Figure 1:
Schematic representation of the harmonic alignment method.

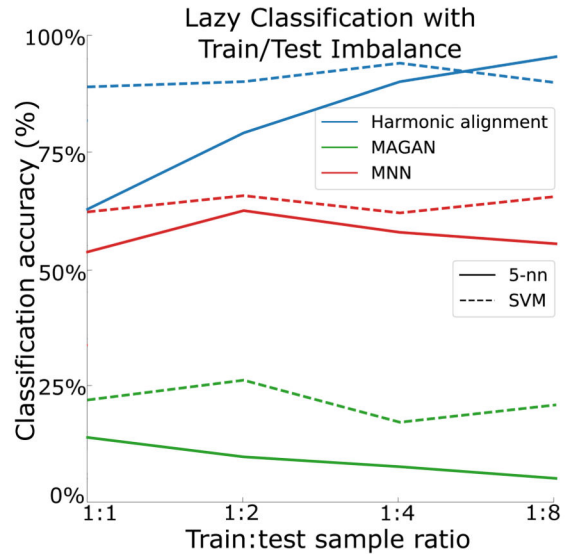


(a)

(b)



(c)

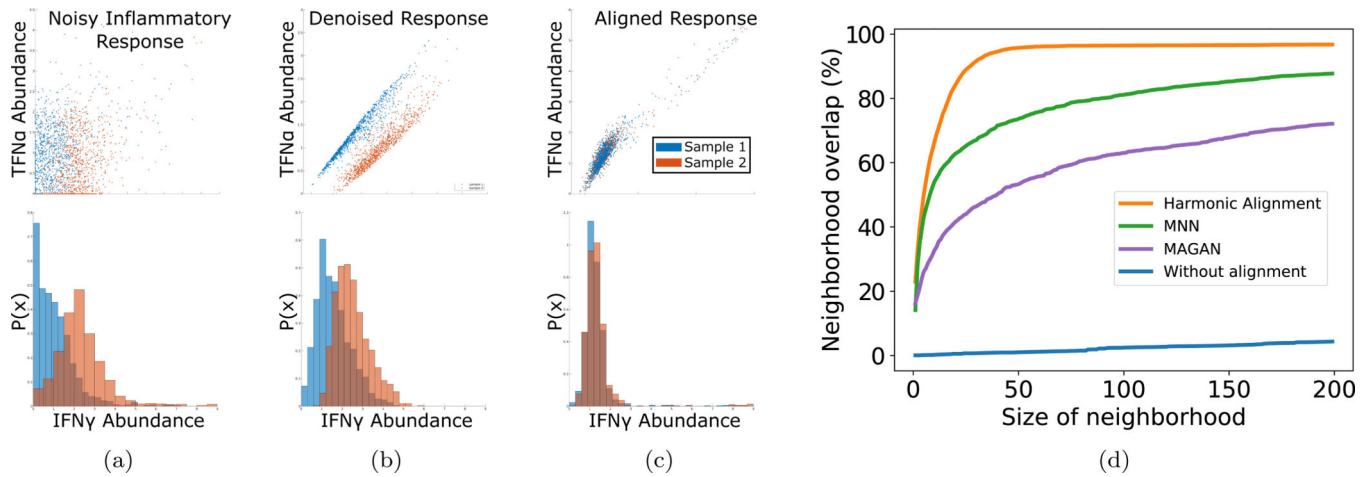


(d)

Figure 2:

Recovery of k -nearest neighborhoods under feature corruption. Mean over 3 iterations is reported for each method. (a) At each iteration, two sets X and Y of 1000 points were sampled from MNIST. Y was then distorted by a 784×784 corruption matrix \mathbf{O}_p for various identity percentages p (§5.1). Subsequently, a lazy classification scheme was used to classify points in $Y\mathbf{O}_p$ using a 5-nearest neighbor vote from X . Results for harmonic alignment with $\ell \in \{2, 4, 8, 64\}$ (§4.1), mutual nearest neighbors (MNN), and classification without alignment are shown. (b) Reconstruction of digits with only 25% uncorrupted features. Left:

Input digits. Left middle: 75% of the pixels in the input are corrupted. Right middle: Reconstruction without harmonic alignment. Right: Reconstruction after harmonic alignment. (c) Lazy classification accuracy relative to input size with unlabeled randomly corrupted digits with 35% preserved pixels. (d) Transfer learning performance. For each ratio, 1K uncorrupted, labeled digits were sampled from MNIST, and then 1K, 2K, 4K, and 8K (x-axis) unlabeled points were sampled and corrupted with 35% column identity.

**Figure 3:**

(a)-(c) *Batch effect removal*. 4K cells were subsampled from two single-cell mass cytometry immune profiles on blood samples of two patients infected with Dengue fever. *Top*: Both patients exhibit heightened IFN γ (x-axis), a pro-inflammatory cytokine associated with TNF α (y-axis) *Bottom*: IFN γ histograms for each batch. (a) Data before denoising. (b) Denoising of unaligned data enhances a technical effect between samples in IFN γ . (c) Harmonic alignment corrects the IFN γ shift. (d) *Multimodal data fusion*. Overlap of cell neighborhoods from joint gene expression and chromatin profiling of single cells. Harmonic alignment most accurately recovers the pointwise relationship between the manifolds.