
Diffusion Earth Mover’s Distance and Distribution Embeddings

Alexander Tong^{*1} Guillaume Huguet^{*23} Amine Natik^{*23} Kincaid MacDonald⁴ Manik Kuchroo⁵¹
Ronald R. Coifman⁴ Guy Wolf^{†23} Smita Krishnaswamy^{†51}

Abstract

We propose a new fast method of measuring distances between large numbers of related high dimensional datasets called the Diffusion Earth Mover’s Distance (EMD). We model the datasets as distributions supported on common data graph that is derived from the affinity matrix computed on the combined data. In such cases where the graph is a discretization of an underlying Riemannian closed manifold, we prove that Diffusion EMD is topologically equivalent to the standard EMD with a geodesic ground distance. Diffusion EMD can be computed in $\tilde{O}(n)$ time and is more accurate than similarly fast algorithms such as tree-based EMDs. We also show Diffusion EMD is fully differentiable, making it amenable to future uses in gradient-descent frameworks such as deep neural networks. Finally, we demonstrate an application of Diffusion EMD to single cell data collected from 210 COVID-19 patient samples at Yale New Haven Hospital. Here, Diffusion EMD can derive distances between patients on the manifold of cells at least two orders of magnitude faster than equally accurate methods. This distance matrix between patients can be embedded into a higher level patient manifold which uncovers structure and heterogeneity in patients. More generally, Diffusion EMD is applicable to all datasets that are massively collected in parallel in many medical and biological systems.

1. Introduction

With the profusion of modern high dimensional, high throughput data, the next challenge is the integration and analysis of collections of related datasets. Examples of this are particularly prevalent in single cell measurement modalities where data (such as mass cytometry, or single cell RNA sequencing data) can be collected in a multitude of patients, or in thousands of perturbation conditions (Shifrut et al., 2018). These situations motivate the organization and embedding of datasets, similar to how we now organize data points into low dimensional embeddings, e.g., with PHATE (Moon et al., 2019), tSNE (van der Maaten & Hinton, 2008), or diffusion maps (Coifman & Lafon, 2006)). The advantage of such organization is that we can use the datasets as rich high dimensional features to characterize and group the patients or perturbations themselves. In order to extend embedding techniques to entire datasets, we have to define a distance between datasets, which for our purposes are essentially high dimensional point clouds. For this we propose a new form of Earth Mover’s Distance (EMD), which we call *Diffusion EMD*¹, where we model the datasets as distributions supported on a common data affinity graph. We provide two extremely fast methods for computing Diffusion EMD based on an approximate multiscale kernel density estimation on a graph.

Optimal transport is uniquely suited to the formulation of distances between entire datasets (each of which is a collection of data points) as it generalizes the notion of the shortest path between two points to the shortest set of paths between distributions. Recent works have applied optimal transport in the single-cell domain to interpolate lineages (Schiebinger et al., 2019; Yang & Uhler, 2019; Tong et al., 2020), interpolate patient states (Tong & Krishnaswamy, 2020), integrate multiple domains (Demetci et al., 2020), or similar to this work build a manifold of perturbations (Chen et al., 2020). All of these approaches use the standard *primal* formulation of the Wasserstein distance. Using either entropic regularization approximation and the Sinkhorn algorithm (Cuturi, 2013) to solve the discrete distribution case or a neural network based approach in the continuous formulation (Ar-

^{*}Equal contribution; [†]Equal senior-author contribution. ¹Dept. of Comp. Sci., Yale University, New Haven, CT, USA ²Dept. of Math. & Stat., Université de Montréal, Montréal, QC, Canada ³Mila – Quebec AI Institute, Montréal, QC, Canada ⁴Dept. of Math., Yale University, New Haven, CT, USA ⁵Department of Genetics, Yale University, New Haven, CT, USA. Correspondence to: Smita Krishnaswamy <smita.krishnaswamy@yale.edu>.

¹Python implementation is available at <https://github.com/KrishnaswamyLab/DiffusionEMD>.

jovsky et al., 2017). We will instead use the *dual* formulation through the well-known Kantorovich-Rubinstein dual to efficiently compute optimal transport between many distributions lying on a common low-dimensional manifold in a high-dimensional measurement space. This presents both theoretical and computational challenges, which are the focus of this work.

Specifically, we will first describe a new Diffusion EMD that is an L^1 distance between density estimates computed using multiple scales of diffusion kernels over a graph. Using theory on the Hölder-Lipschitz dual norm on continuous manifolds (Leeb & Coifman, 2016), we show that as the number of samples increases, Diffusion EMD is equivalent to the Wasserstein distance on the manifold. This formulation reduces the computational complexity of computing K -nearest Wasserstein-neighbors between m distributions over n points from $O(m^2n^3)$ for the exact computation to $\tilde{O}(mn)$ with reasonable assumptions on the data. Finally, we will show how this can be applied to embed large sets of distributions that arise from a common graph, for instance single cell datasets collected on large patient cohorts.

Our contributions include: 1. A new method for computing EMD for distributions over graphs called Diffusion EMD. 2. Theoretical analysis of the relationship between Diffusion EMD and the snowflake of a standard EMD. 3. Fast algorithms for approximating Diffusion EMD. 4. Demonstration of the differentiability of this framework. 5. Application of Diffusion EMD to embedding massive multi-sample biomedical datasets.

2. Preliminaries

We now briefly review optimal transport definitions and classic results from diffusion geometry.

Notation. We say that two elements A and B are equivalent if there exist $c, C > 0$ such that $cA \leq B \leq CA$, and we denote $A \simeq B$. The definition of A and B will be clear depending on the context.

Optimal Transport. Let μ, ν be two probability distributions on a measurable space Ω with metric $d(\cdot, \cdot)$, $\Pi(\mu, \nu)$ be the set of joint probability distributions π on the space $\Omega \times \Omega$, where for any subset $\omega \subset \Omega$, $\pi(\omega \times \Omega) = \mu(\omega)$ and $\pi(\Omega \times \omega) = \nu(\omega)$. The 1-Wasserstein distance W_d also known as the earth mover's distance (EMD) is defined as:

$$W_d(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \int_{\Omega \times \Omega} d(x, y) \pi(dx, dy). \quad (1)$$

When μ, ν are discrete distributions with n points, then Eq. 1 is computable in $O(n^3)$ with a network-flow based algorithm (Peyré & Cuturi, 2019).

Let $\|\cdot\|_{L_d}$ denote the Lipschitz norm w.r.t. d , then the dual of Eq. 1 is:

$$W_d(\mu, \nu) = \sup_{\|f\|_{L_d} \leq 1} \int_{\Omega} f(x) \mu(dx) - \int_{\Omega} f(y) \nu(dy). \quad (2)$$

This formulation is known as the Kantorovich dual with f as the witness function. Since it is in general difficult to optimize over the entire space of 1-Lipschitz functions, many works optimize the cost over a modified family of functions such as functions parameterized by clipped neural networks (Arjovsky et al., 2017), functions defined over trees (Le et al., 2019), or functions defined over Haar wavelet bases (Gavish et al., 2010).

Data Diffusion Geometry Let $(\mathcal{M}, d_{\mathcal{M}})$ be a connected Riemannian manifold, we can assign to \mathcal{M} sigma-algebras and look at \mathcal{M} as a ‘‘metric measure space’’. We denote by Δ the Laplace-Beltrami operator on \mathcal{M} . For all $x, y \in \mathcal{M}$ let $h_t(x, y)$ be the heat kernel, which is the minimal solution of the heat equation:

$$\left(\frac{\partial}{\partial t} - \Delta_x \right) h_t = 0, \quad (3)$$

with initial condition $\lim_{t \rightarrow 0} h_t(x, y) = \delta_y(x)$, where $x \mapsto \delta_y(x)$ is the Dirac function centered at y , and Δ_x is taken with respect to the x argument of h_t . Note that as shown in Grigor'yan et al. (2014), the heat kernel captures the local intrinsic geometry of \mathcal{M} in the sense that as $t \rightarrow 0$, $\log h_t(x, y) \simeq -d_{\mathcal{M}}^2(x, y)/4t$.

Here, in Sec. 4.2 (Theorem 1) we discuss another topological equivalent of the geodesic distance with a diffusion distance derived from the heat operator $\mathbf{H}_t := e^{-t\Delta}$ that characterizes the solutions of the heat equation (Eq. 3), and is related to the heat kernel via $\mathbf{H}_t f = \int h_t(\cdot, y) f(y) dy$ (see Lafferty & Lebanon, 2005; Coifman & Lafon, 2006; Grigor'yan et al., 2014, for further details).

It is often useful (particularly in high dimensional data) to consider data as sampled from a lower dimensional manifold embedding in the ambient dimension. This manifold can be characterized by its local structure and in particular, how heat propagates along it. Coifman & Lafon (2006) showed how to build such a propagation structure over discrete data by first building a graph with affinities

$$(\mathbf{K}_{\epsilon})_{ij} := e^{-\|x_i - x_j\|_2^2 / \epsilon} \quad (4)$$

then considering the density normalized operator $\mathbf{M}_{\epsilon} := \mathbf{Q}^{-1} \mathbf{K}_{\epsilon} \mathbf{Q}^{-1}$, where $\mathbf{Q}_{ii} := \sum_j (\mathbf{K}_{\epsilon})_{ij}$. Lastly, a Markov diffusion operator is defined by

$$\mathbf{P}_{\epsilon} := \mathbf{D}^{-1} \mathbf{M}_{\epsilon}, \text{ where } \mathbf{D}_{ii} := \sum_j (\mathbf{M}_{\epsilon})_{ij}. \quad (5)$$

Both D and Q are diagonal matrices. By the law of large numbers, the operator P_ϵ admits a natural continuous equivalent \tilde{P}_ϵ , i.e., for n i.i.d. points, the sums modulo n converge to the integrals. Moreover, in Coifman & Lafon (2006, Prop. 3) it is shown that $\lim_{\epsilon \rightarrow 0} \tilde{P}_\epsilon^{t/\epsilon} = e^{-t\Delta} = H_t$. In conclusion, the operator P_ϵ converges to \tilde{P}_ϵ as the sample size increases and \tilde{P}_ϵ provide an approximation of the Heat kernel on the manifold. Henceforth, we drop the subscript of P_ϵ to lighten the notation, further we will use the notation P_ϵ for the operator as well as for the matrix (it will be clear in the context).

3. EMD through the L^1 metric between multiscale density estimates

The method of efficiently approximating EMD that we will consider here is the approximation of EMD through density estimates at multiple scales. Previous work has considered densities using multiscale histograms over images (Indyk & Thaper, 2003), wavelets over images and trees (Shirdhonkar & Jacobs, 2008; Gavish et al., 2010) and densities over hierarchical clusters (Le et al., 2019). Diffusion EMD also uses a hierarchical set of bins at multiple scales but with smooth bins determined by the heat kernel, which allows us to show equivalence to EMD with a ground distance of the manifold geodesic. These methods are part of a family that we call *multiscale earth mover's distances* that first compute a set of multiscale density estimates or histograms where L^1 differences between density estimates realizes an effective *witness function* and have (varying) equivalence to the Wasserstein distance between the distributions. This class of multiscale EMDs are particularly useful in computing embeddings of distributions where the Wasserstein distance is the ground distance, as they are amenable to fast nearest neighbor queries. We explore this application further in Sec. 4.6 and these related methods in Sec. B.1 of the Appendix.

4. Diffusion Earth Mover's Distance

We now present the Diffusion EMD, a new Earth Mover's distance based on multiscale diffusion kernels as depicted in Fig. 1. We first show how to model multiple datasets as distributions on a common data graph and perform multiscale density estimates on this graph.

4.1. Data Graphs and Density Estimates on Graphs

Let $\mathcal{X} = \{X_1, X_2, \dots, X_m\}$, $\cup_{j=1}^m X_j \subseteq \mathcal{M} \subset \mathbb{R}^d$, be a collection of datasets with $n_i = |X_i|$ and $n = \sum_i n_i$. Assume that the X_i 's are independently sampled from a common underlying manifold $(\mathcal{M}, d_{\mathcal{M}})$ which is a Riemannian closed manifold (compact and without boundary) immersed in a (high dimensional) ambient space \mathbb{R}^d , with

geodesic distance $d_{\mathcal{M}}$. Further, assume that while the underlying manifold is common, each dataset is sampled from a different distribution over it, as discussed below. Such collections of datasets arise from several related samples of data, for instance single cell data collected on a cohort of patients with a similar condition.

Here, we consider the datasets in \mathcal{X} as representing distributions over the common data manifold, which we represent in the finite setting as a common data graph $G_{\mathcal{X}} = (V, E, w)$ with $V = \cup_{j=1}^m X_j$ and edge weights determined by the Gaussian kernel (see Eq. 4), where we identify edge existence with nonzero weights. Then, we associate each X_i with a density measure $\mu_i^{(t)} : V \rightarrow [0, 1]$, over the entire data graph. To compute such measures, we first create indicator vectors for the individual datasets on it, let $\mathbf{1}_{X_i} \in \{0, 1\}^n$ be a vector where for each $v \in V$, $\mathbf{1}_{X_i}(v) = 1$ if and only if $v \in X_i$. We then derive a kernel density estimate by applying the diffusion operator constructed via Eq. 5 over the graph G to these indicator functions to get scale-dependent estimators

$$\mu_i^{(t)} := \frac{1}{n_i} P^t \mathbf{1}_{X_i}, \quad (6)$$

where the scale t is the diffusion time, which can be considered as a meta-parameter (e.g., as used in Burkhardt et al., 2021) but can also be leveraged in multiscale estimation of distances between distributions as discussed here. Indeed, as shown in Burkhardt et al. (2021), at an appropriately tuned single scale, this density construction yields a discrete version of kernel density estimation.

4.2. Diffusion Earth Mover's Distance Formulation

We define the Diffusion Earth Mover's Distance between two datasets $X_i, X_j \in \mathcal{X}$ as

$$W_{\alpha, K}(X_i, X_j) := \sum_{k=0}^K \|T_{\alpha, k}(X_i) - T_{\alpha, k}(X_j)\|_1 \quad (7)$$

where $0 < \alpha < 1/2$ is a meta-parameter used to balance long- and short-range distances, which in practice is set close to $1/2$, K is the maximum scale considered here, and

$$T_{\alpha, k}(X_i) := \begin{cases} 2^{-(K-k-1)\alpha} (\mu_i^{(2^{k+1})} - \mu_i^{(2^k)}) & k < K \\ \mu_i^{(2^K)} & k = K \end{cases} \quad (8)$$

Further, to set K , we note that if the Markov process governed by P converges (i.e., to its stationary steady state) in polynomial time w.r.t. $|V|$, then one can ensure that beyond $K = O(\log |V|)$, all density estimates would be essentially indistinguishable as shown by the following lemma, whose proof appears in the Appendix:




Figure 1. Diffusion EMD first embeds datasets into a common data graph G , then takes multiscale diffusion KDEs for each of the datasets. These multiscale KDEs are then used to compute the Diffusion Earth Mover's Distance between the datasets that can be used in turn to create graphs and embeddings (PHATE (Moon et al., 2018) shown here) of the datasets.

Lemma 1. *There exists a $K = O(\log |V|)$ such that $\mu_i^{(2^K)} \simeq \mathbf{P}^{2^K} \mathbf{1}_{X_i} \simeq \phi_0$ for every $i = 1, \dots, n$, where ϕ_0 is the trivial eigenvector of \mathbf{P} associated with the eigenvalue $\lambda_0 = 1$.*

To compute the Diffusion EMD $W_{\alpha, K}$ in Eq. 7 involves first calculating diffusions of the datasets μ , second calculating differences and weighting them in relation to their scale, this results in a vector per distribution of length $O(nK)$, and finally computing the L^1 norm between them. The most computationally intensive step, computing the diffusions, is discussed in further detail in Sec. 4.4.

4.3. Theoretical Relation to EMD on Manifolds

We now provide a theoretical justification of the Diffusion EMD defined via Eq. 7 by following the relation established in Leeb & Coifman (2016) between heat kernels and the EMD on manifolds. Leeb & Coifman (2016) define the following ground distance for EMD over \mathcal{M} by leveraging the geometric information gathered from the L^1 distances between kernels at different scales.

Definition 1. *The diffusion ground distance between $x, y \in \mathcal{M}$ is defined as*

$$D_\alpha(x, y) := \sum_{k \geq 0} 2^{-k\alpha} \|h_{2^{-k}}(x, \cdot) - h_{2^{-k}}(y, \cdot)\|_1,$$

for $\alpha \in (0, 1/2)$, the scale parameter $K \geq 0$ and $h_t(\cdot, \cdot)$ the heat kernel on \mathcal{M} .

Note that $D_\alpha(\cdot, \cdot)$ is similar to the diffusion distance defined in Coifman & Lafon (2006), which was based on L^2 notions rather than L^1 here. Further, the following result from Leeb & Coifman (2016, see Sec. 3.3) shows that the distance $D_\alpha(\cdot, \cdot)$ is closely related to the intrinsic geodesic one $d_{\mathcal{M}}(\cdot, \cdot)$.

Theorem 1. *Let $(\mathcal{M}, d_{\mathcal{M}})$ be a closed manifold with geodesic distance $d_{\mathcal{M}}$, and let $\alpha \in (0, 1/2)$. The metric $D_\alpha(\cdot, \cdot)$ defined via Def. 1 is equivalent to $d_{\mathcal{M}}(\cdot, \cdot)^{2\alpha}$.*

The previous theorem justifies why in practice we let α close to $1/2$, because we want the snowflake distance $d_{\mathcal{M}}^{2\alpha}(\cdot, \cdot)$ to approximate the geodesic distance of \mathcal{M} . The notion of equivalence established by this result is such that $D_\alpha(x, \cdot) \simeq d(x, \cdot)^{2\alpha}$. It is easy to verify that two equivalent metrics induce the same topology. We note that while here we only consider the Heat kernel, a similar result holds (see Theorem 4 in the Appendix) for a more general family of kernels, as long as they satisfy certain regularity conditions.

For a family of operators $(\mathbf{A}_t)_{t \in \mathbb{R}^+}$ we define the following metric on distributions; let μ and ν be two distributions:

$$\widehat{W}_{\mathbf{A}_t}(\mu, \nu) = \|\mathbf{A}_1(\mu - \nu)\|_1 + \sum_{k \geq 0} 2^{-k\alpha} \|(\mathbf{A}_{2^{-(k+1)}} - \mathbf{A}_{2^{-k}})(\mu - \nu)\|_1. \quad (9)$$

The following result shows that applying this metric for the family of operators \mathbf{H}_t to get $\widehat{W}_{\mathbf{H}_t}$ yields an equivalent of the EMD with respect to the diffusion ground distance $D_\alpha(\cdot, \cdot)$.

Theorem 2. *The EMD between two distributions μ, ν on a closed Riemannian manifold $(\mathcal{M}, d_{\mathcal{M}})$ w.r.t. the diffusion ground distance $D_\alpha(\cdot, \cdot)$, defined via Def. 1, given by*

$$W_{D_\alpha}(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{M} \times \mathcal{M}} D_\alpha(x, y) \pi(dx, dy), \quad (10)$$

is equivalent to $\widehat{W}_{\mathbf{H}_t}$. That is $W_{D_\alpha} \simeq \widehat{W}_{\mathbf{H}_t}$, where \mathbf{H}_t is the Heat operator on \mathcal{M} .

Proof. In Proposition 15 of Leeb & Coifman (2016), it is shown that \mathcal{M} is separable w.r.t. $D_\alpha(\cdot, \cdot)$, hence we can use the Kantorovich-Rubinstein theorem. We let Λ_α , the space of functions that are Lipschitz w.r.t. $D_\alpha(\cdot, \cdot)$ and $\|\cdot\|_{\Lambda_\alpha^*}$, the norm of its dual space Λ_α^* . The norm is defined by

$$\|T\|_{\Lambda_\alpha^*} := \sup_{\|f\|_{\Lambda_\alpha} \leq 1} \int_{\mathcal{M}} f dT.$$

In Theorem 4 of Leeb & Coifman (2016), it is shown that both W_{D_α} and \widehat{W}_{H_t} are equivalent to the norm $\|\cdot\|_{\Lambda_\alpha^*}$. \square

We consider the family of operators $(P_\epsilon^{t/\epsilon})_{t \in \mathbb{R}^+}$, which is related to the continuous equivalent of the stochastic matrix defined in Eq. 5. In practice, we use this family of operators to approximate the heat operator H_t . Indeed, when we take a small value of ϵ , as discussed in Sec. 2, we have from Coifman & Lafon (2006) that this is a valid approximation.

Corollary 2.1. *Let P_ϵ be the continuous equivalent of the stochastic matrix in Eq. 5. For ϵ small enough, we have:*

$$\widehat{W}_{P_\epsilon^{t/\epsilon}} \simeq W_{D_\alpha}. \quad (11)$$

Eq. 11 motivates our use of Eq. 7 to compute the Diffusion EMD. The idea is to take only the first K terms in the infinite sum $\widehat{W}_{P_\epsilon^{t/\epsilon}}$ and then choosing $\epsilon := 2^{-K}$ would give us exactly Eq. 7. We remark that the summation order of Eq. 7 is inverted compared to Eq. 11, but in both cases the largest scale has the largest weight. Finally, we state one last theorem that brings our distance closer to the Wasserstein w.r.t. $d_M(\cdot, \cdot)$; we refer the reader to the Appendix for its proof.

Theorem 3. *Let $\alpha \in (0, 1/2)$ and (M, d_M) be a closed manifold with geodesic d_M . The Wasserstein distance w.r.t. the diffusion ground distance $D_\alpha(\cdot, \cdot)$ is equivalent to the Wasserstein distance w.r.t. the snowflake distance $d_M(\cdot, \cdot)^{2\alpha}$ on M , that is $W_{D_\alpha} \simeq W_{d_M^{2\alpha}}$.*

Corollary 3.1. *For each $1 \leq i, j \leq m$, let $X_i, X_j \in \mathcal{X}$ be two datasets with size n_i and n_j respectively, and let μ_i and μ_j be the continuous distributions corresponding to the ones of X_i and X_j , let K be the largest scale and put $N = \min(K, n_i, n_j)$. Then, for sufficiently big $N \rightarrow \infty$ (implying sufficiently small $\epsilon = 2^{-K} \rightarrow 0$):*

$$W_{\alpha, K}(X_i, X_j) \simeq W_{d_M^{2\alpha}}(\mu_i, \mu_j), \quad (12)$$

for all $\alpha \in (0, 1/2)$.

In fact we can summarize our chain of thought as follows

$$\begin{aligned} W_{\alpha, K}(X_i, X_j) &\stackrel{(a)}{\simeq} \widehat{W}_{P_\epsilon^{t/\epsilon}}(\mu_i, \mu_j) &&\stackrel{(b)}{\simeq} \widehat{W}_{H_t}(\mu_i, \mu_j) \\ &\stackrel{(c)}{\simeq} W_{D_\alpha}(\mu_i, \mu_j) &&\stackrel{(d)}{\simeq} W_{d_M^{2\alpha}}(\mu_i, \mu_j), \end{aligned}$$

where the approximation (a) is due to the fact that the discrete distributions on X_i and X_j converge respectively to μ_i and μ_j when $\min(n_i, n_j) \rightarrow \infty$. Further, $W_{\alpha, K}(X_i, X_j)$ approximate the infinite series $\widehat{W}_{P_\epsilon^{t/\epsilon}}(\mu_i, \mu_j)$ as in Eq. 9 when $K \rightarrow \infty$, note also that we take $\epsilon = 2^{-K}$ so that the largest scale in Eq. 7 is exactly 2^K . The approximation in (b) comes from the approximation of the heat operator as in Coifman & Lafon (2006), (c) comes from Theorem 2 and (d) comes from Theorem 1.

Algorithm 1 Chebyshev embedding

Input: $n \times n$ graph kernel K , $n \times m$ distributions μ , maximum scale K , and snowflake constant α .

Output: $m \times (K + 1)n$ distribution embeddings b

$Q \leftarrow \text{Diag}(\sum_i K_{ij})$

$K^{norm} \leftarrow Q^{-1} K Q^{-1}$

$D \leftarrow \text{Diag}(\sum_i K_{ij}^{norm})$

$M \leftarrow D^{-1/2} K^{norm} D^{-1/2}$

$U \Sigma U^T = M$; U orthogonal, Σ Diagonal

$\mu^{(2^0)} \leftarrow P \mu \leftarrow D^{-1/2} M D^{1/2} \mu$

for $k = 1$ **to** K **do**

$\mu^{(2^k)} \leftarrow P^{2^k} \mu \leftarrow D^{-1/2} U (\Sigma)^{2^k} U^T D^{1/2} \mu$

$b_{k-1} \leftarrow 2^{(K-k-1)\alpha} (\mu^{(2^k)} - \mu^{(2^{k-1})})$

end for

$b_K \leftarrow \mu^{(2^K)}$

$b \leftarrow [b_0, b_1, \dots, b_K]$

4.4. Efficient Computation of Dyadic Scales of the Diffusion Operator

The most computationally intensive step of Diffusion EMD requires computing dyadic scales of the diffusion operator P times μ to estimate the density of $\mu^{(t)}$ at multiple scales which we will call b . After this embedding, Diffusion EMD between two embeddings b_i, b_j is computed as $|b_i - b_j|$, i.e. the L^1 norm of the difference. Computing the embedding b naively by first powering P then right multiplying μ , may take up to 2^K matrix multiplications which is infeasible for even moderately sized graphs. We assume two properties of P that makes this computation efficient in practice. First, that P is sparse with order $\tilde{O}(n)$ non-zero entries. This applies when thresholding K_ϵ or when using a K -nearest neighbors graph to approximate the manifold (van der Maaten & Hinton, 2008; Moon et al., 2019). Second, that P^t is low rank for large powers of t .

While there are many ways to approximate dyadic scales of P , we choose from two methods depending on the number of distributions m compared to the number of points in the graph n . When $m \ll n$, we use a method based on Chebyshev approximation of polynomials of the eigenspectrum of P as shown in Alg. 1. This method is efficient for sparse P and a small number of distributions (Shuman et al., 2011). For more detail on the error incurred by using Chebyshev polynomials we refer the reader to Trefethen (2013, Chap. 3). In practice, this requires for the approximating polynomial of J terms, computation of mJ (sparse) matrix vector multiplications for a worst case time complexity of $O(Jmn^3)$, but in practice is $\tilde{O}(Jmn)$ where J is a small constant (see Fig. 5(e)). However, while asymptotically efficient, when $m \gg n$ in practice this can be inefficient as it requires many multiplications of the form $P\mu$.

In the case where $m \gg n$, approximating powers of P and

applying these to the $n \times m$ collection of distributions μ once is faster. This method is also useful when the full set of distributions is not known and can be applied to new distributions one at a time in a data streaming model. A naive approach for computing P^{2^k} would require K dense $n \times n$ matrix multiplications. However, as noted in Coifman & Maggioni (2006), for higher powers of P , we can use a much smaller basis. We use an algorithm based on interpolative decomposition (Liberty et al., 2007; Bermanis et al., 2013) to reduce the size of the basis, and subsequently the computation time, at higher scales. In this algorithm we first determine the approximate rank of P^{2^k} using an estimate of the density of the eigenspectrum of P as in Dong et al. (2019). We then alternate steps of downsampling the basis to the specified rank of P^{2^k} with (randomized) interpolative decomposition with steps of powering P on these bases. Informally, the interpolative decomposition selects a representative set of points that approximate the basis well. In the worst case this algorithm can take $\tilde{O}(mn^3)$ time to compute the diffusion density estimates, nevertheless with sparse P with a rapidly decaying spectrum, this algorithm is $\tilde{O}(mn)$ in practice. For more details see Alg. 4 and Sec. C of the Appendix.

4.5. Subsampling Density Estimates

The density estimates created for each distribution are both large and redundant with each distribution represented by a vector of $(K + 1) \times n$ densities. However, as noted in the previous section, P^{2^k} can be represented on a smaller basis, especially for larger scales. Intuitively, the long time diffusions of nodes that are close to each other are extremely similar. Interpolative decomposition (Liberty et al., 2007; Bermanis et al., 2013) allows us to pick a set of points to center our diffusions kernels such that they approximately cover the graph up to some threshold on the rank. In contrast, in other multiscale EMD methods the bin centers or clusters are determined randomly, making it difficult to select the number of centers necessary. Furthermore, the relative number of centers at every scale is fixed, for example, Quadtree or Haar wavelet based methods (Indyk & Thaper, 2003; Gavish et al., 2010) use 2^{d^k} centers at every scale, and a clustering based method (Le et al., 2019) selects C^k clusters at every scale for some constant C . Conversely, in Diffusion EMD, by analyzing the behavior of P^{2^k} , we intelligently select the number of centers needed at each scale based on the approximate rank of P^{2^k} up to some tolerance at each scale. This does away with the necessity of a fixed ratio of bins at every scale, allowing adaptation depending on the structure of the manifold and can drastically reduce the size representations (see Fig. 5(c)). For the Chebyshev polynomials method, this subsampling is done post computation of diffusion scales, and for the method based on approximating

P^{2^k} directly the subsampling happens during computation. To this point, we have described a method to embed distributions on a graph into a set of density estimates whose size depends on the data, and the spectrum decay of P . We will now explore how to use these estimates for exploring the Diffusion EMD metric between distributions.

4.6. Diffusion EMD Based Embeddings of Samples

Our main motivation for a fast EMD computed on related datasets is to examine the space of the samples or datasets themselves, i.e., the higher level manifold of distributions. In terms of the clinical data, on which we show this method, this would be the relationship between patients themselves, as determined by the EMD between their respective single-cell peripheral blood datasets. Essentially, we create a kernel matrix $K_{\mathcal{X}}$ and diffusion operator $P_{\mathcal{X}}$ between datasets where the samples are nodes on the associated graph. This diffusion operator $P_{\mathcal{X}}$ can be embedded using diffusion maps (Coifman & Lafon, 2006) or visualized with a method like PHATE (Moon et al., 2018) that collects the information into two dimensions as shown in Sec. 5. We note this higher level graph can be a sparse KNN graph, particularly given that a diffusion operator on the graph can allow for global connections to be reformed via t -step path probabilities.

Multiscale formulations of EMD as in Eq. 8 are especially effective when searching for nearest neighbor distributions under the Wasserstein metric (Indyk & Thaper, 2003; Backurs et al., 2020) as this distance forms a normed space, i.e., a space where the metric is induced by the L_1 norm of the distribution vectors and their differences. Data structures such as kd-trees, ball trees, locality sensitive hashing, are able to take advantage of such normed spaces for sub-linear neighbor queries. This is in contrast to network-flow or Sinkhorn type approximations that require a scan through all datapoints for each nearest neighbor query as this metric is not derived from a norm.

4.7. Gradients of the Earth Mover’s Distance

One of the hindrances in the use of optimal transport-based distances has been the fact that it cannot be easily incorporated into deep learning frameworks. Gradients with respect to the EMD are usually found using a trained Lipschitz discriminator network as in Wasserstein-GANs (Arjovsky et al., 2017), which requires unstable adversarial training, or by taking derivatives through a small number of iterations of the Sinkhorn algorithm (Frogner et al., 2015; Bonneel et al., 2016; Genevay et al., 2018; Liu et al., 2020), which scales with $O(n^2)$ in the number of points. Tree based methods that are linear in the number of points do not admit useful gradients due to their hard binning over space, giving a gradient of zero norm almost everywhere.

We note that, given the data diffusion operator, the com-

putation of Diffusion EMD is differentiable and, unlike Tree-based EMD, has smooth bins and therefore a non-zero gradient norm near the data (as visible in Fig. 2). Further, computation of the gradient only requires powers of the diffusion operator multiplied by the indicator vector describing the distribution on the graph. In fact, as mentioned in the supplementary material (Sec. C.1) for each $v \in V$, the gradient of the Diffusion EMD $\partial W_{\alpha, K}(X_i, X_j)/\partial v$ depends mainly on the gradients $\partial P_{\epsilon}^{2^k}/\partial v$ for $0 \leq K$ which can be expressed in terms of the gradient of the Gaussian kernel $\partial K_{\epsilon}/\partial v$. This last quantity is easy to compute. In Sec. C.1 of the Appendix, we give an exact process on computing the gradient of the Diffusion EMD.

5. Results

In this section, we first evaluate the Diffusion EMD on two manifolds where the ground truth EMD with a geodesic ground distance is known, a swiss roll dataset and spherical MNIST (Cohen et al., 2017). On these datasets where we have access to the ground truth geodesic distance we show that Diffusion EMD is both faster and closer to the ground truth than comparable methods. Then, we show an application to a large single cell dataset of COVID-19 patients where the underlying metric between cells is thought to be a manifold (Moon et al., 2018; Kuchroo et al., 2020). We show that the manifold of patients based on Diffusion EMD by capturing the graph structure, better captures the disease state of the patients.

Experimental Setup. We consider four baseline methods for approximating EMD: QuadTree(D) (Backurs et al., 2020) which partitions the dataspace in half in each dimension up to some specified depth D , ClusterTree(C, D) (Le et al., 2019) which recursively clusters the data with C clusters up to depth D using the distance between clusters to weight the tree edges, and the convolutional Sinkhorn distance (Solomon et al., 2015) with the same graph as used in Diffusion EMD. QuadTree and ClusterTree are fast to compute. However, because they operate in the ambient space, they do not represent the geodesic distances on the manifold in an accurate way. The convolutional Sinkhorn method represents the manifold well but is significantly slower even when using a single iteration. For more details on related work and the experimental setup see Sections B and D of the Appendix respectively.

1D data. We first illustrate the advantages of using Diffusion EMD over tree-based methods (ClusterTree and QuadTree) on a line graph with 500 points spaced in the interval $[0, 1]$. In Fig. 2 we depict the Wasserstein distance between an indicator function at each of these 500 points ($\mathbf{1}_x$) and an indicator function at $x = 0.5$. In the case of indicator functions the Wasserstein distance is exactly the




Figure 2. Wasserstein distance of indicator distributions $\mathbf{1}_x, x \in [0, 1]$ from $\mathbf{1}_{0.5}$ computed using linear EMD methods L^2 distance: (a) ClusterTree (b) QuadTree and (c) Diffusion EMD.

ground distance. We show three approximations of each method, varying the number of trees and the Chebyshev polynomial order respectively. It is clear that Diffusion EMD achieves a much better approximation of the ground truth primarily due to its use of smooth bins.




Figure 3. Swiss roll dataset embeddings of $m = 1000$ distributions with $n = 10,000$ total points rotated into 10D colored by ground truth 2D sheet axes. Diffusion EMD recreates the manifold better in similar time.

Swiss roll data. The next application we explore is to a dataset where we have distributions on a manifold for which the geodesic distance is easily computable. In this way we can compare to the ground truth EMD between distributions. We generate $m = 100$ Gaussians on the swiss roll with 100 points each for a total of $n = 10,000$ points on the graph. We compare each method on two metrics, the 10-nearest neighbor accuracy measured (P@10) where a P@10 of 1 means that the 10-nearest neighbors are the same as the ground truth. We compare the rankings of nearest neighbors over the entire dataset using the Spearman- ρ correlation coefficient, which measures the similarity of nearest neighbor rankings. This coefficient ranges between -1 for inversely ranked lists and 1 for the same ranks. This measures rankings over the entire dataset equally rather




Figure 4. Accuracy of methods measured via P@10 (left) and Spearman coefficient (right), against their (log scaled) computation time in seconds on the swiss roll dataset. Variations of methods are over Chebyshev approximation order for Diffusion EMD, # of trees for tree methods, and number of iterations for conv. Sinkhorn. Diffusion EMD is more accurate than tree methods and orders of magnitude faster than conv. Sinkhorn even with a single iteration.




Figure 5. Ablation study of major parameters for Chebyshev polynomial approximation on the swiss roll dataset. Mean and std. over 10 runs over (a) values of the maximum scale K , (b) the rank threshold in interpolative decomposition, (c) the total number of centers in the L^1 representation which drops with decomposition, (d-e) performance against the # of scales, and the order of the polynomial, both are very stable after a certain point, and (f) time vs. the Chebyshev order.

than only considering the nearest neighbors. Visually, we show embeddings of the swiss roll in Fig. 3, where the 2D manifold between distributions is best captured by Diffusion EMD given a similar amount of time.

In Fig. 4, we investigate the time vs. accuracy tradeoff of a number of fast EMD methods on the swiss roll. We compare against the ground truth EMD which is calculated with the exact EMD on the “unrolled” swiss roll in 2D. We find that Diffusion EMD is more accurate than tree methods for a given amount of time and is much faster than the convolutional Sinkhorn method and only slightly less accurate. To generate multiple models for each dataset we vary the number of trees for tree methods, the Chebyshev order for Diffusion EMD, and the number of iterations for convolutional Sinkhorn. We search over and fix other parameters using a grid search as detailed in Sec. D of the Appendix.

In Fig. 5, we vary parameters of the Chebyshev polynomial algorithm of Diffusion EMD. Regarding performance, we find Diffusion EMD is stable to the number of scales chosen after a certain minimum maximum scale K , Chebyshev

polynomial order, and the number of scales used. By performing interpolative decomposition with a specified rank threshold on P^{2^k} we can substantially reduce the embedding size at a small cost to performance Fig. 5(b,c).

Table 1. Classification accuracy, P@10, Spearman ρ and runtime (in minutes) on 70,000 distributions from Spherical MNIST.

	ACCURACY	P@10	SPEARMAN ρ	TIME
DIFF. EMD	95.94	0.611	0.673	34M
CLUSTER	91.91	0.393	0.484	30M
QUAD	79.56	0.294	0.335	16M

Spherical MNIST. Next, we use the Spherical MNIST dataset to demonstrate the efficacy of the interpolative decomposition based approximation to Diffusion EMD, as here $m \gg n$. Each image is treated as a distribution (of pixel intensities) over the sphere. To evaluate the fidelity of each embedding, we evaluate the 1-NN classification on the embedding vectors in addition to P@10 and Spearman coefficient in Tab. 1. Diffusion EMD creates embeddings that better approximate true EMD over the sphere than tree methods in a similar amount of time, which in this case also gives better classification accuracy.

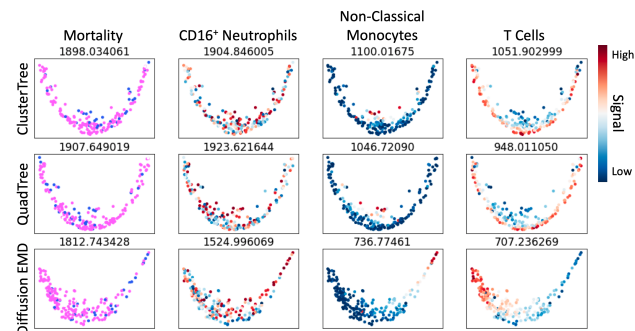


Figure 6. Embedding of 210 patients through different manifold constructions. Visualizing patient eventual mortality and cell types predictive of disease outcome on each manifold. Laplacian smoothness reported on each signal for each manifold.

Single cell COVID-19 patient data. The COVID-19 pandemic has driven biologists to generate vast amounts of cellular data on hospitalized patients suffering from severe disease. A major question in clinicians minds is determining a priori which patients may be at risk for worse outcomes, including requiring increased ventilatory support and increased risk of mortality. Certain cell types found in the blood, such as CD16⁺ Neutrophils, T cells and non-classical monocytes, have been associated with and predictive of mortality outcome. Ideally, a manifold of patients would find these cellular populations to occupy a region of high mortality for CD16⁺ Neutrophils and non-classical monocytes

and low mortality for T cells. In order to construct a manifold of patients suffering from COVID-19, we analyzed 210 blood samples from 168 patients infected with SARS-CoV-2 measured on a myeloid-specific flow cytometry panel, an expanded iteration of a previously published dataset (Lucas et al., 2020). We embedded 22 million cells from these patients into a common combined cell-cell graph with 27,000 nodes as defined in Sec. 4.1. We then computed Diffusion EMD and other methods on these datasets. Diffusion EMD is computed by using indicator vectors for each patient converted to density estimates as in Eq. 6.

On an informative embedding of patients, similar patients, with similar features (such as mortality) would localize on the manifold and thus the important features should be smooth over the manifold. Furthermore, cell types which are correlated with outcome either positively or negatively should also be smooth and either correlated or anticorrelated with outcome. To quantify this, we compute the smoothness with respect to the patient manifold by using a *Laplacian quadratic form* with respect to the 10-NN graph between patients. Convolutional Sinkhorn does not scale to this data, so we compare a patient manifold created with Diffusion EMD to ones created with QuadTree and ClusterTree. Diffusion EMD is able to use the manifold of cells where QuadTree and ClusterTree are built in the ambient space. In Fig. 6 we visualize relevant signals over the patients using PHATE (Moon et al., 2019) overlaid with the quadratic smoothness of the signal over the graph. While the mortality signal appeared enriched in the right branches of both the Diffusion EMD and QuadTree manifolds, it did not localize as well on the ClusterTree manifold. Both CD16⁺ Neutrophils and non-classical monocytes appeared smoother over the Diffusion EMD manifold than the comparison manifolds. Since both cell types are associated with mortality, it was interesting to see them both enriched in high mortality region of the Diffusion EMD manifold but not the others. Finally, T cells, which are negatively correlated with mortality appeared smoothly enriched in the Diffusion EMD manifold in a region with no mortality. In QuadTree and ClusterTree constructions, T cells appeared enriched throughout the manifold, no localizing smoothly to a region with low mortality. These experiments show that the patient manifold constructed with Diffusion EMD is more informative, smoothly localizing key signals, such as patient outcome and predictive cell types. For a more details see Sec. D of the Appendix.

6. Conclusion

In this work we have introduced Diffusion EMD, a multi-scale distance that uses heat kernel diffusions to approximate the earth mover’s distance over a data manifold. We showed how Diffusion EMD can efficiently embed many samples

on a graph into a manifold of samples in $\tilde{O}(mn)$ time more accurately than similarly efficient methods. This is useful in the biomedical domain as we show how to embed COVID-19 patient samples into a higher level patient manifold that more accurately represents the disease structure between patients. Finally, we also show how to compute gradients with respect to Diffusion EMD, which opens the possibility of gradients of the earth mover’s distance that scale linearly with the dataset size.

Acknowledgements

This research was partially funded by IVADO PhD Excellence Scholarship [A.N.]; IVADO Professor startup & operational funds, IVADO Fundamental Research Proj. grant PRF-2019-3583139727, Canada CIFAR AI Chair [G.W.]; Chan-Zuckerberg Initiative grants 182702 & CZF2019-002440 [S.K.]; and NIH grants R01GM135929 & R01GM130847 [G.W., S.K.]. The content provided here is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

References

- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70 of *PMLR*, pp. 214–223, 2017.
- Backurs, A., Dong, Y., Indyk, P., Razenshteyn, I., and Wagner, T. Scalable nearest neighbor search for optimal transport. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *PMLR*, pp. 497–506, 2020.
- Benamou, J.-D., Carlier, G., Cuturi, M., Nenna, L., and Peyré, G. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- Bermanis, A., Averbuch, A., and Coifman, R. R. Multi-scale data sampling and function extension. *Applied and Computational Harmonic Analysis*, 34(1):15–29, 2013.
- Bonneel, N., Peyré, G., and Cuturi, M. Wasserstein barycentric coordinates: Histogram regression using optimal transport. *ACM Transactions on Graphics*, 35(4), 2016.
- Burkhardt, D. B., Stanley, J. S., Tong, A., Perdigoto, A. L., Gigante, S. A., Herold, K. C., Wolf, G., Giraldez, A. J., van Dijk, D., and Krishnaswamy, S. Quantifying the effect of experimental perturbations at single-cell resolution. *Nature Biotechnology*, 39:619–629, 2021.
- Chen, W. S., Zivanovic, N., van Dijk, D., Wolf, G., Bodenmiller, B., and Krishnaswamy, S. Uncovering axes of

- variation among single-cell cancer specimens. *Nature Methods*, 17:302–310, 2020.
- Cohen, T., Geiger, M., Köhler, J., and Welling, M. Convolutional networks for spherical signals. In *ICML 2017 Workshop on Principled Approaches to Deep Learning*, 2017.
- Coifman, R. R. and Lafon, S. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
- Coifman, R. R. and Maggioni, M. Diffusion wavelets. *Applied and Computational Harmonic Analysis*, 21(1):53–94, 2006.
- Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems 26*, pp. 2292–2300, 2013.
- Demetci, P., Santorella, R., Sandstede, B., Noble, W. S., and Singh, R. Gromov-wasserstein optimal transport to align single-cell multi-omics data. bioRxiv, DOI: 10.1101/2020.04.28.066787, 2020.
- Dong, K., Benson, A. R., and Bindel, D. Network density of states. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1152–1161, 2019.
- Frogner, C., Zhang, C., Mobahi, H., Araya, M., and Poggio, T. A. Learning with a wasserstein loss. In *Advances in Neural Information Processing Systems 28 (NeurIPS 2015)*, 2015.
- Gavish, M., Nadler, B., and Coifman, R. R. Multiscale wavelets on trees, graphs and high dimensional data: Theory and applications to semi supervised learning. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pp. 367–374, 2010.
- Genevay, A., Peyre, G., and Cuturi, M. Learning generative models with sinkhorn divergences. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTAT)*, volume 84 of *PMLR*, pp. 1608–1617, 2018.
- Grigor'yan, A. and Liu, L. Heat kernel and Lipschitz-Besov spaces. *Forum Mathematicum*, 27(6):3567–3613, 2015.
- Grigor'yan, A., Hu, J., and Lau, K.-S. Heat kernels on metric measure spaces. In *Geometry and Analysis of Fractals*, volume 88 of *PROMS*, pp. 147–207, 2014.
- Indyk, P. and Thaper, N. Fast image retrieval via embeddings. In *the 3rd International Workshop on Statistical and Computational Theories of Vision*, 2003.
- Kuchroo, M., Huang, J., Wong, P., Grenier, J.-C., Shung, D., Tong, A., Lucas, C., Klein, J., Burkhardt, D., Gigante, S., Godavarthi, A., Israelow, B., Mao, T., Oh, J. E., Silva, J., Takahashi, T., Odio, C. D., Casanovas-Massana, A., Fournier, J., the Yale IMPACT Team, Farhadian, S., Dela Cruz, C. S., Ko, A. I., Wilson, F. P., Hussin, J., Wolf, G., Iwasaki, A., and Krishnaswamy, S. Multiscale PHATE exploration of SARS-CoV-2 data reveals multimodal signatures of disease. bioRxiv, DOI: 10.1101/2020.11.15.383661, 2020.
- Lafferty, J. and Lebanon, G. Diffusion kernels on statistical manifolds. *Journal of Machine Learning Research*, 6(5): 129–163, 2005.
- Le, T., Yamada, M., Fukumizu, K., and Cuturi, M. Tree-sliced variants of Wasserstein distances. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, 2019.
- Leeb, W. *Topics in Metric Approximation*. PhD thesis, Yale University, 2015.
- Leeb, W. The mixed lipschitz space and its dual for tree metrics. *Applied and Computational Harmonic Analysis*, 44(3):584–610, 2018.
- Leeb, W. and Coifman, R. Hölder-Lipschitz norms and their duals on spaces with semigroups, with applications to Earth mover's distance. *Journal of Fourier Analysis and Applications*, 22:910–953, 2016.
- Liberty, E., Woolfe, F., Martinsson, P.-G., Rokhlin, V., and Tygert, M. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, 104(51):20167–20172, December 2007.
- Liu, R., Zou, J., and Balsubramani, A. Learning transport cost from subset correspondence. In *the 8th International Conference on Learning Representations (ICLR 2020)*, 2020.
- Lucas, C., Wong, P., Klein, J., Castro, T. B. R., Silva, J., Sundaram, M., Ellingson, M. K., Mao, T., Oh, J. E., Israelow, B., Takahashi, T., Tokuyama, M., Lu, P., Venkataraman, A., Park, A., Mohanty, S., Wang, H., Wyllie, A. L., Vogels, C. B. F., Earnest, R., Lapidus, S., Ott, I. M., Moore, A. J., Muenker, M. C., Fournier, J. B., Campbell, M., Odio, C. D., Casanovas-Massana, A., Herbst, R., Shaw, A. C., Medzhitov, R., Schulz, W. L., Grubaugh, N. D., Cruz, C. D., Farhadian, S., Ko, A. I., Omer, S. B., and Iwasaki, A. Longitudinal analyses reveal immunological misfiring in severe COVID-19. *Nature*, 584:463–469, 2020.

- Moon, K. R., Stanley, J. S., Burkhardt, D., van Dijk, D., Wolf, G., and Krishnaswamy, S. Manifold learning-based methods for analyzing single-cell rna-sequencing data. *Current Opinion in Systems Biology*, 7:36–46, 2018.
- Moon, K. R., van Dijk, D., Wang, Z., Gigante, S., Burkhardt, D. B., Chen, W. S., Yim, K., van den Elzen, A., Hirn, M. J., Coifman, R. R., Ivanova, N. B., Wolf, G., and Krishnaswamy, S. Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology*, 37(12):1482–1492, 2019.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, É. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011.
- Peyré, G. and Cuturi, M. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5–6): 355–607, 2019.
- Sato, R., Yamada, M., and Kashima, H. Fast unbalanced optimal transport on a tree. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, pp. 19039–19051, 2020.
- Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube, P., Lee, L., Chen, J., Brumbaugh, J., Rigollet, P., Hochedlinger, K., Jaenisch, R., Regev, A., and Lander, E. S. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943.e22, 2019.
- Shifrut, E., Carnevale, J., Tobin, V., Roth, T. L., Woo, J. M., Bui, C. T., Li, P. J., Diolaiti, M. E., Ashworth, A., and Marson, A. Genome-wide CRISPR screens in primary human T cells reveal key regulators of immune function. *Cell*, 175(7):1958–1971.e15, 2018.
- Shirdhonkar, S. and Jacobs, D. W. Approximate Earth mover's distance in linear time. In *the 2008 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- Shuman, D. I., Vandergheynst, P., and Frossard, P. Chebyshev polynomial approximation for distributed signal processing. In *the 2011 International Conference on Distributed Computing in Sensor Systems and Workshops (DCOSS)*, pp. 1–8, 2011.
- Solomon, J., de Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., and Guibas, L. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics*, 34(4):1–11, 2015.
- Tong, A. and Krishnaswamy, S. Interpolating optimal transport barycenters of patient manifolds, 2020.
- Tong, A., Huang, J., Wolf, G., van Dijk, D., and Krishnaswamy, S. TrajectoryNet: A dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, volume 119 of PMLR, pp. 9526–9536, 2020.
- Trefethen, L. N. *Approximation Theory and Approximation Practice*. Society for Industrial and Applied Mathematics (SIAM), 2013.
- van der Maaten, L. and Hinton, G. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(86): 2579–2605, 2008.
- Yang, K. D. and Uhler, C. Scalable unbalanced optimal transport using generative adversarial networks. In *the 7th International Conference on Learning Representations (ICLR 2019)*, 2019.

Supplemental Material

We first analyze the theoretical framework of Diffusion EMD in Appendix A. Next we discuss related work in Appendix B, with a particular focus on multiscale methods for EMD. In Appendix C we provide further detail on the two algorithms for computing Diffusion EMD, the first based on Chebyshev approximation, and the second based on directly powering the diffusion operator while reducing the basis. We discuss gradients of Diffusion EMD in section C.1. Finally we provide further experimental details in Appendix D.

A. General framework and proofs

A.1. General framework

We now recall some useful results from Leeb & Coifman (2016). We will only present an overview, for a rigorous exposition, we suggest Leeb (2015), Leeb & Coifman (2016) and Grigor'yan & Liu (2015).

We let \mathcal{Z} be a sigma-finite measure space of dimension n , $d(\cdot, \cdot)$ its intrinsic distance and ξ its associated sigma-finite measure. In order to evaluate the EMD between two measures on \mathcal{Z} , one would need to know $d(\cdot, \cdot)$. Either directly, by using equation 1, or implicitly, to define the space of Lipschitz functions with respect to $d(\cdot, \cdot)$ in equation 2. One of the contributions of this paper is to define a kernel based metric $D_\alpha(\cdot, \cdot)$, where the kernel depends on $\alpha \in (0, 1)$, and show that the metrics $D_\alpha(\cdot, \cdot)$ and $d(\cdot, \cdot)$ are closely related. Next, the objective is to use $D_\alpha(\cdot, \cdot)$ as the ground distance to compute the EMD in its dual form. That is a norm between two distributions acting on the space of Lipschitz functions with respect to $D_\alpha(\cdot, \cdot)$. To do so, the authors define Λ_α the space of functions that are Lipschitz with respect to $D_\alpha(\cdot, \cdot)$, and its dual Λ_α^* ; the space of measures acting on $f \in \Lambda_\alpha$. Further, in order to use the Kantorovich–Rubinstein theorem, they show that the space \mathcal{Z} is separable with respect to the metric $D_\alpha(\cdot, \cdot)$. Thus, the norm of the dual space $\|\cdot\|_{\Lambda_\alpha^*}$ can be used to compute the EMD with $D_\alpha(\cdot, \cdot)$ as the ground distance. Lastly, they define two norms $\|\cdot\|_{\Lambda_\alpha^*}^{(1)}$ and $\|\cdot\|_{\Lambda_\alpha^*}^{(2)}$ that are equivalent to $\|\cdot\|_{\Lambda_\alpha^*}$ on Λ^* . In practice, these norms are much faster to compute.

The metric $D_\alpha(\cdot, \cdot)$ is defined using a family of kernels $\{a_t(\cdot, \cdot)\}_{t \in \mathbb{R}^+}$ on \mathcal{Z} . For each kernel, we define an operator \mathbf{A}_t as $(\mathbf{A}_t f)(x) = \int_{\mathcal{Z}} a_t(x, y) f(y) d\xi(y)$. These kernels must respect some properties:

- The semigroup property: for all $s, t > 0$, $\mathbf{A}_t \mathbf{A}_s = \mathbf{A}_{t+s}$;
- The conservation property: $\int_{\mathcal{Z}} a_t(x, y) d\xi(y) = 1$;
- The integrability property: there exists $C > 0$ such that $\int_{\mathcal{Z}} |a_t(x, y)| d\xi(y) < C$, for all $t > 0$ and $x \in \mathcal{Z}$.

Considering only the dyadic times, that is $t = 2^{-k}$, we define the kernel $p_k(\cdot, \cdot) := a_{2^{-k}}(\cdot, \cdot)$ and the operator $\mathbf{P}_k = \mathbf{A}_{2^{-k}}$, for all $k \in \mathbb{N}$. By leveraging the local geometric information gathered from the L^1 distance between two measures

$$D_k(x, y) := \|p_k(x, \cdot) - p_k(y, \cdot)\|_1,$$

the authors define the following multiscale metric

$$D_\alpha(x, y) := \sum_{k \geq 0} 2^{-k\alpha} D_k(x, y).$$

To interpret $D_k(\cdot, \cdot)$ in an intuitive way, consider the case where $a_t(\cdot, \cdot)$ defines a random walk on \mathcal{Z} . As a consequence, $a_t(x, B(y, r))$ is the probability to move from x to a point in $B(y, r)$ in t steps. Moreover, for any $x \in \mathcal{Z}$, $a_t(x, \cdot)$ defines a distribution, therefore $D_k(x, y)$ is the L^1 distance between two distributions induced by the points x and y . Since these distributions depend on the number of steps t , it is clever to consider a distance that includes many scales, just like $D_\alpha(\cdot, \cdot)$. Another property needs to be verified by the kernel $a_t(\cdot, \cdot)$, this property depends on the distance $D_\alpha(\cdot, \cdot)$. Namely, the geometric property: there exist $C > 0$ and $\alpha \in (0, 1)$ such that for all $k \in \mathbb{N}$ and $x \in \mathcal{Z}$

$$\int_{\mathcal{Z}} |p_k(x, y)| D_\alpha(x, y) d\xi(y) \leq C 2^{-k\alpha}.$$

We need to add three stronger regularity conditions on the kernel $a_t(\cdot, \cdot)$, for $D_\alpha(\cdot, \cdot)$ to be closely related to the intrinsic distance $d(\cdot, \cdot)$:

1. An upper bound on the kernel: there exist a non-negative, monotonic decreasing function $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}$ and $\beta > 0$ such that for any $\gamma < \beta$, the function verifies

$$\int_0^\infty \tau^{n+\gamma-1} \phi(\tau) d\tau$$

and

$$|a_t(x, y)| \leq t^{-\frac{n}{\beta}} \phi\left(\frac{d(x, y)}{t^{1/\beta}}\right).$$

2. Hölder continuity estimate: there exist $\Theta > 0$ sufficiently small, such that, for all $t \in (0, 1]$ and all $x, y \in \mathcal{Z}$, the distance verifies $d(x, y) \leq t^{1/\beta}$ and for all $u \in \mathcal{Z}$ the difference between kernels is bounded

$$|a_t(x, u) - a_t(y, u)| \leq t^{-\frac{n}{\beta}} \left(\frac{d(x, y)}{t^{1/\beta}}\right)^\Theta \phi\left(\frac{d(x, y)}{t^{1/\beta}}\right).$$

3. A local lower bound: there exist a monotonic decreasing function $\psi : \mathbb{R}^+ \rightarrow \mathbb{R}$ and $R > 0$, such that, for all $t \in (0, 1]$ and all x, y where $d(x, y) < R$, we have

$$|a_t(x, y)| \geq t^{-\frac{n}{\beta}} \psi\left(\frac{d(x, y)}{t^{1/\beta}}\right).$$

It is shown that the heat kernel on a closed Riemannian manifold respects all these conditions (Leeb & Coifman, 2016, see Sec. 3.3).

Definition 2. A distance $d(\cdot, \cdot)^b$ is a snowflake of a distance $d(\cdot, \cdot)$ if $b \in (0, 1)$. Moreover, the Hölder space is the space of functions that are Lipschitz continuous w.r.t. a snowflake distance, hence the terminology used in Leeb & Coifman (2016).

We now state an important theorem from Leeb & Coifman (2016, Thm. 2).

Theorem 4. Consider a sigma-finite measure space \mathcal{Z} of dimension n with metric $d(\cdot, \cdot)$ and a measure ξ such that $\xi(B(x, r)) \simeq r^n$. If the family of kernels $\{a_t(\cdot, \cdot)\}_{t \in \mathbb{R}^+}$ respect the condition 1, 2 and 3, then, for $0 < \alpha < \min(1, \Theta/\beta)$, the distance $D_\alpha(\cdot, \cdot)$ is equivalent to the thresholded snowflake distance $\min[1, d(\cdot, \cdot)^{\alpha\beta}]$.

Remark. In our case, because we used the heat kernel on a closed Riemannian manifold \mathcal{M} , we had $D_\alpha(\cdot, \cdot) \simeq d_{\mathcal{M}}(\cdot, \cdot)^{2\alpha}$ (Thm. 1). This can be justified from Leeb & Coifman (2016, Cor. 2). Using the same notation as in the corollary, we define $C := \max_{x, y} d_{\mathcal{M}}(x, y)^{2\alpha}$ (which is finite due to the assumptions on \mathcal{M}), thus we can bound the constant B

$$B = \frac{B}{d_{\mathcal{M}}(x, y)^{2\alpha}} d_{\mathcal{M}}(x, y)^{2\alpha} \geq \frac{B}{C} d_{\mathcal{M}}(x, y)^{2\alpha}.$$

The previous theorem closely links the two considered distances. It also motivated the goal to compute the EMD w.r.t. $D_\alpha(\cdot, \cdot)$. In Leeb & Coifman (2016)[Prop. 15], it is shown that \mathcal{Z} is a separable space with respect to the metric $D_\alpha(\cdot, \cdot)$. Hence, we can use the Kantorovich-Rubinstein theorem to express the EMD in its dual form.

First, we need to define the space of functions that are Lipschitz with respect to $D_\alpha(\cdot, \cdot)$. For a fix $\alpha \in (0, 1)$, we note this space by Λ_α . It corresponds to the set of functions f on \mathcal{Z} such that the norm

$$\|f\|_{\Lambda_\alpha} := \sup_x |f(x)| + \sup_{x \neq y} \frac{|f(x) - f(y)|}{D_\alpha(x, y)}$$

is finite. Next, we need to define the space dual to Λ_α , which is noted by Λ_α^* . For a function $f \in \Lambda_\alpha$ and a L^1 measure T , we define $\langle f, T \rangle := \int_{\mathcal{Z}} f dT$. The space Λ_α^* is the space of L^1 measure with the norm

$$\|T\|_{\Lambda_\alpha^*} := \sup_{\|f\|_{\Lambda_\alpha} \leq 1} \langle f, T \rangle.$$

In practice, this norm would still be computationally expensive. However, the authors show that the norms

$$\|T\|_{\Lambda_\alpha^*}^{(1)} := \|P_0^* T\|_1 + \sum_{k \geq 0} 2^{-k\alpha} \|(P_{k+1}^* - P_k^*) T\|_1$$

$$\|T\|_{\Lambda_\alpha^*}^{(2)} := \|P_0^* T\|_1 + \sum_{k \geq 0} 2^{-k\alpha} \|(P_k^* - P_0^*) T\|_1$$

are equivalent to the norm $\|\cdot\|_{\Lambda_\alpha^*}$ on Λ_α^* (Leeb & Coifman, 2016, Thm. 4). Where P_k^* is the adjoint of P_k . Finally, using the Kantorovich-Rubinstein theorem, we get

$$\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{Z} \times \mathcal{Z}} D_\alpha(x, y) \pi(dx, dy) \|\mu - \nu\|_{\Lambda_\alpha^*} \simeq \|\mu - \nu\|_{\Lambda_\alpha^*}^{(1)} \simeq \|\mu - \nu\|_{\Lambda_\alpha^*}^{(2)},$$

where $D_\alpha(\cdot, \cdot) \simeq \min[1, d(\cdot, \cdot)^{\beta\alpha}]$ and $d(\cdot, \cdot)$ is the ground distance of \mathcal{Z} .

In conclusion, using $\|\cdot\|_{\Lambda_\alpha^*}^{(1)}$ or $\|\cdot\|_{\Lambda_\alpha^*}^{(2)}$ yields a norm equivalent to $\|\cdot\|_{\Lambda_\alpha^*}$. The norm $\|\mu - \nu\|_{\Lambda_\alpha^*}$ is equal to the Wasserstein distance between the distributions μ and ν with respect to the ground distance $D_\alpha(\cdot, \cdot)$.

A.2. Proofs of section 4.2

Lemma 1. *Assuming that P converges (i.e. to its stationary distribution) in polynomial time w.r.t. $|V|$, then there exists a $K = O(\log |V|)$ such that $\mu_i^{(2^K)} \simeq P^{2^K} \mathbf{1}_{X_i} \approx \phi_0$ for every $i = 1, \dots, n$, where ϕ_0 is the trivial eigenvector of P associated with the eigenvalue $\lambda_0 = 1$.*

Proof. First we notice that P is reversible with respect to $\pi_i = D_{ii} / \sum_i \sum_j D_{ij}$. Since P is ergodic it converges to its unique stationary distribution π . Moreover, we assumed this convergence to be in polynomial time w.r.t. $|V|$, i.e. we define

$$\Delta_i(k) := \frac{1}{2} \sum_j \left| P_{ij}^{2^k} - \pi_j \right|$$

and the mixing time for a $\epsilon > 0$

$$\tau_i(\epsilon) := \min\{k : \Delta_i(k') < \epsilon, \forall k' \geq k\}.$$

Then, by our assumption, there exist $2^K = \tau_i(\epsilon) = O(|V|)$, thus $K = O(\log |V|)$. Intuitively, for all $k \geq K$, each row of the matrix P^{2^k} is approximately equal to π (w.r.t. ϵ), as a consequence $P^{2^k} \mathbf{1}_{X_i} \approx \phi_0$. □

Theorem 3. *Let $\alpha \in (0, 1/2)$ and $(\mathcal{M}, d_{\mathcal{M}})$ be a closed manifold with geodesic $d_{\mathcal{M}}$. The Wasserstein distance w.r.t. the diffusion ground distance $D_\alpha(\cdot, \cdot)$ is equivalent to the Wasserstein distance w.r.t. the snowflake distance $d_{\mathcal{M}}(\cdot, \cdot)^{2\alpha}$ on \mathcal{M} , that is $W_{D_\alpha} \simeq W_{d_{\mathcal{M}}^{2\alpha}}$.*

Proof. We will prove this for a more general framework, we let two metrics $d_1(\cdot, \cdot)$ and $d_2(\cdot, \cdot)$ on a sigma-finite measure space \mathcal{Z} , and let μ and ν be two measures. We assume $d_1 \simeq d_2$, that is there exist two constants $c, C > 0$ such that for all $x, y \in \mathcal{Z}$ we have $c d_1(x, y) \leq d_2(x, y) \leq C d_1(x, y)$. Using the same notation as in Eq. 1, for all $\pi \in \Pi(\mu, \nu)$ we have

$$c \int_{\mathcal{Z} \times \mathcal{Z}} d_1(x, y) \pi(dx, dy) \leq \int_{\mathcal{Z} \times \mathcal{Z}} d_2(x, y) \pi(dx, dy) \leq C \int_{\mathcal{Z} \times \mathcal{Z}} d_1(x, y) \pi(dx, dy)$$

then

$$c \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{Z} \times \mathcal{Z}} d_1(x, y) \pi(dx, dy) \leq \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{Z} \times \mathcal{Z}} d_2(x, y) \pi(dx, dy) \leq C \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{Z} \times \mathcal{Z}} d_1(x, y) \pi(dx, dy)$$

which is the same as $cW_{d_1}(\mu, \nu) \leq W_{d_2}(\mu, \nu) \leq CW_{d_1}(\mu, \nu)$, this proves that whenever d_1 and d_2 are equivalent then W_{d_1} and W_{d_2} are equivalent as well. □

Corollary 2.1. *Let P_ϵ be the continuous equivalent of the stochastic matrix in Eq. 5. For ϵ small enough, we have:*

$$\widehat{W}_{P_\epsilon^{t/\epsilon}} \simeq W_{D_\alpha}. \quad (13)$$

Proof. Note that by Theorem 2 it is equivalent to show that $\widehat{W}_{P_\epsilon^{t/\epsilon}} \simeq \widehat{W}_{H_t}$. First note that according to (Coifman & Lafon, 2006) we have $\left\| P_\epsilon^{t/\epsilon} - H_t \right\|_{L^2(\mathcal{M})} \rightarrow 0$ as $\epsilon \rightarrow 0$, note that since $\text{Vol}(\mathcal{M}) < \infty$ (closed Manifold) and according to

Cauchy-Schwarz inequality we get $\left\| \mathbf{P}_\epsilon^{t/\epsilon} - \mathbf{H}_t \right\|_{L^1(\mathcal{M})} \rightarrow 0$ as $\epsilon \rightarrow 0$. Generally speaking, convergence in L^2 implies convergence in L^1 when the manifold is closed. Now put $\mathbf{D}_{\epsilon,t} = \mathbf{P}_\epsilon^{t/\epsilon} - \mathbf{H}_t$, then we have $\|\mathbf{D}_{\epsilon,t}\|_{L^1(\mathcal{M})} \rightarrow 0$ as $\epsilon \rightarrow 0$. Let μ and ν be two measures and let $\delta > 0$, choose $\epsilon > 0$ small enough so that $\|\mathbf{D}_{t,\epsilon}\gamma\|_1 < \delta\|\gamma\|_1$ for all $t > 0$, where $\gamma = \mu - \nu$,

$$\begin{aligned} \widehat{W}_{\mathbf{D}_{\epsilon,t}}(\mu, \nu) &= \|\mathbf{D}_{\epsilon,1}\gamma\|_1 + \sum_{k \geq 0} 2^{-k\alpha} \left\| (\mathbf{D}_{\epsilon,2^{-(k+1)}} - \mathbf{D}_{\epsilon,2^{-k}}) \gamma \right\|_1 \\ &\leq \|\mathbf{D}_{\epsilon,1}\gamma\|_1 + \sum_{k \geq 0} 2^{-k\alpha} \|\mathbf{D}_{\epsilon,2^{-(k+1)}}\gamma\|_1 + 2^{-k\alpha} \|\mathbf{D}_{\epsilon,2^{-k}}\gamma\|_1 \\ &\leq \delta\|\gamma\|_1 + 2\delta\|\gamma\|_1 \sum_{k \geq 0} 2^{-k\alpha} \\ &= \delta\|\gamma\|_1 \left(1 + \frac{2}{1-2^{-\alpha}} \right). \end{aligned}$$

This proves that for all $t > 0$, for all $\delta > 0$ there exists an $\epsilon > 0$ sufficiently small such that

$$\widehat{W}_{\mathbf{D}_{\epsilon,t}}(\mu, \nu) \leq \delta\|\mu - \nu\|$$

Note also that using the reverse triangle inequality we can easily show that $\left| \widehat{W}_{\mathbf{P}_\epsilon^{t/\epsilon}}(\mu, \nu) - \widehat{W}_{\mathbf{H}_t}(\mu, \nu) \right| \leq \widehat{W}_{\mathbf{D}_{\epsilon,t}}(\mu, \nu)$, let $\delta > 0$ then for $\epsilon > 0$ small enough we get

$$\widehat{W}_{\mathbf{H}_t}(\mu, \nu) - \delta\|\mu - \nu\|_1 \leq \widehat{W}_{\mathbf{P}_\epsilon^{t/\epsilon}}(\mu, \nu) \leq \widehat{W}_{\mathbf{H}_t}(\mu, \nu) + \delta\|\mu - \nu\|_1,$$

according to (Wang et al., 1997) we can get lower bounds of the heat kernel (which implies lower bounds for the heat operator), we have $\widehat{W}_{\mathbf{H}_t}(\mu, \nu) \geq \|\mathbf{H}_1(\mu - \nu)\|_1 \geq C\|\mu - \nu\|_1$ for some $C > 0$. If $\delta < C/2$ then for sufficiently small $\epsilon > 0$,

$$\frac{1}{2}\widehat{W}_{\mathbf{H}_t}(\mu, \nu) \leq \widehat{W}_{\mathbf{P}_\epsilon^{t/\epsilon}}(\mu, \nu) \leq \frac{3}{2}\widehat{W}_{\mathbf{H}_t}(\mu, \nu)$$

which completes the proof. \square

Corollary 3.1. For each $1 \leq i, j \leq m$ let $X_i, X_j \in \mathcal{X}$ be two datasets with size n_i and n_j respectively, and let μ_i and μ_j be the continuous distributions corresponding to the ones of X_i and X_j , let K be the largest scale and put $N = \min(K, n_i, n_j)$. Then, for sufficiently big $N \rightarrow \infty$ (implying sufficiently small $\epsilon = 2^{-K} \rightarrow 0$):

$$W_{\alpha,K}(X_i, X_j) \simeq W_{d_{\mathcal{M}}^{2\alpha}}(\mu_i, \mu_j), \quad (14)$$

for all $\alpha \in (0, 1/2)$.

Proof. First define

$$W_{\alpha,K,\epsilon}(X_i, X_j) := \left\| \mu_i^{(1/\epsilon)} - \mu_j^{(1/\epsilon)} \right\|_1 + \sum_{k=0}^{K-1} 2^{-k\alpha} \left\| (\mu_i^{(2^{-(k+1)}/\epsilon)} - \mu_i^{(2^{-k}/\epsilon)}) - (\mu_j^{(2^{-(k+1)}/\epsilon)} - \mu_j^{(2^{-k}/\epsilon)}) \right\|_1, \quad (15)$$

note that for $\epsilon = 2^{-K}$ we have $W_{\alpha,K,\epsilon}(X_i, X_j) = W_{\alpha,K}(X_i, X_j)$. Since $n_i \rightarrow \infty$ and $n_j \rightarrow \infty$, then using Monte-Carlo integration we get

$$\lim_{n_i, n_j \rightarrow \infty} W_{\alpha,K,\epsilon}(X_i, X_j) = W_{\alpha,K,\epsilon}(\mu_i, \mu_j)$$

where $W_{\alpha,K,\epsilon}(\mu_i, \mu_j)$ has the same expression as in Eq. 15 (replacing the discrete measure μ_i 's with the continuous measures μ_i 's),

$$W_{\alpha,K,\epsilon}(\mu_i, \mu_j) = \left\| \mathbf{P}_\epsilon^{1/\epsilon} \mu_i - \mathbf{P}_\epsilon^{1/\epsilon} \mu_j \right\|_1 + \sum_{k=0}^{K-1} 2^{-k\alpha} \left\| (\mathbf{P}_\epsilon^{2^{-(k+1)}/\epsilon} \mu_i - \mathbf{P}_\epsilon^{2^{-k}/\epsilon} \mu_i) - (\mathbf{P}_\epsilon^{2^{-(k+1)}/\epsilon} \mu_j - \mathbf{P}_\epsilon^{2^{-k}/\epsilon} \mu_j) \right\|_1$$

note that by definition of $\widehat{W}_{P_\epsilon^{t/\epsilon}}$ we have $\lim_{K \rightarrow \infty} W_{\alpha, K, \epsilon}(\mu_i, \mu_j) = \widehat{W}_{P_\epsilon^{t/\epsilon}}(\mu_i, \mu_j)$ and thus

$$\lim_{N \rightarrow \infty} W_{\alpha, K, \epsilon}(X_i, X_j) = \widehat{W}_{P_\epsilon^{t/\epsilon}}(\mu_i, \mu_j)$$

combining this result with Corollary 2.1 yields $W_{\alpha, K, \epsilon}(X_i, X_j) \simeq W_{D_\alpha}(\mu_i, \mu_j)$ for N large enough and ϵ small enough. Now if we take $\epsilon = 2^{-K}$ and apply Theorem 3 we get $W_{\alpha, K}(X_i, X_j) \simeq W_{d_{\mathcal{M}}^{2\alpha}}(\mu_i, \mu_j)$ as required. \square

B. Related Work

Wavelet-based linear time approximations of the earth mover's distance have been investigated by Indyk & Thaper (2003) who used randomly shifted multi-scale grids to compute EMD between images. Shirdhonkar & Jacobs (2008) expanded this to wavelets over \mathbb{R}^n showing the benefit of using smooth bins over the data, and the relationship to the dual problem. Leeb & Coifman (2016) investigated using wavelets over more general metric spaces. We build off of this theory and efficiently approximate this distance on discrete manifolds represented by sparse kernels.

Another line of work in approximations to the Wasserstein distance instead works with tree metrics over the data (Leeb, 2018; Le et al., 2019; Backurs et al., 2020; Sato et al., 2020). These tree methods also linear time however are built in the ambient dimension and do not represent graph distances as well as graph based methods as shown in Sec. 5. Averaging over certain tree families can be linked to approximating the EMD with a Euclidean ground distance in \mathbb{R}^n (Indyk & Thaper, 2003). Many trees may be necessary to smooth out the effects of binning over a continuous space as is evident in Figure 2. Diffusion EMD uses multiple scales of smooth bins to reduce this effect, thus giving a smoother distance which is approximately equivalent to a snowflake of the L^2 distance on the manifold.

A third line of work considers entropy regularized Wasserstein distances (Cuturi, 2013; Benamou et al., 2015; Solomon et al., 2015). These methods show that the transportation plan of the regularized 2-Wasserstein distance is a rescaling of the heat kernel accomplished by the iterative matrix rescaling algorithm known as the Sinkhorn algorithm. In particular, Solomon et al. (2015) links the time parameter t in the heat kernel H_t to the entropy regularization parameter and performs Sinkhorn scaling with H_t . While applying the heat kernel is efficient, to embed m distributions with the entropy regularized Wasserstein distance is $O(m^2)$ as all pairwise distances must be considered. Next we explore the linear time algorithms based on multiscale smoothing for compute the earth mover's distance.

B.1. Multiscale Methods for Earth Mover's Distance

Let a (possibly randomized) transformation T map distributions on Ω to a set of multiscale bins \mathbf{b} , $T : \mu(\Omega) \rightarrow \mathbf{b}$, these methods define a T such that

$$W_d(\mu, \nu) \approx \mathbb{E} \|T(\mu) - T(\nu)\|_1. \quad (16)$$

Where the approximation, the randomness, and the transform depend on the exact implementation. All of these methods have two things in common, they smooth over multiple scales

Indyk & Thaper (2003) presented one of the early methods of this type. They showed that by computing averages over a set of randomly placed grids at dyadic scales. These grids work well for images where the pixels form a discrete set of coordinates.

Shirdhonkar & Jacobs (2008) showed how to generalize this work over images to more general bin types, more specifically they showed how wavelets placed centered on a grid could replace the averages in Indyk & Thaper (2003). This work linked the earth mover's distance in the continuous domain to that of the discrete through wavelets in \mathbb{R}^d , showing that for some wavelets the EMD approximation was better than the previous grid method. In Diffusion EMD we generalize these wavelets to the graph domain using diffusion wavelets on the graph allowing Wasserstein distances with a geodesic ground distance.

Kolouri et al. (2016) in Sliced Wasserstein distances showed how the Wasserstein distance could be quickly approximated by taking the Wasserstein distance along many one dimensional slices of the data. This can be thought of binning along one dimension where the cumulative distribution function is represented along n bins (one for each point) with each bin encompassing one more point than the last.

Le et al. (2019) generalized the Sliced Wasserstein distance back to trees, giving a new method based on multi-level clustering where the data is partitioned into smaller and smaller clusters where the bins are the clusters. They demonstrated this method on high dimensional spaces where previous methods that had d^{2^k} bins at scale k were inefficiently using bins. By clustering based on the data, their ClusterTree performs well in high dimensions. However, by clustering they lose the convergence to the Wasserstein distance with Euclidean ground distance of previous methods for efficiency in high dimensions.

Diffusion EMD uses smooth diffusion wavelets over the graph which can give geodesic distances unlike previous multiscale methods. Furthermore, Diffusion EMD selects a number of bins that depends on the rank of the data at each dyadic scale, which can lead to smaller representations depending on the data. This is summarized in Table 2, which details these differences.

Table 2. Comparison of Multiscale Methods for Earth Mover’s Distance

METHOD	SCALES	BINS PER SCALE	DATA DEPENDENT CENTERS	SMOOTH BINS	GEODESIC DISTANCES
INDYK & THAPER (2003)	DYADIC	d^{2^k}	NO	NO	NO
SHIRDHONKAR & JACOBS (2008)	DYADIC	d^{2^k}	NO	YES	NO
KOLOURI ET AL. (2016)	n	1	YES	NO	NO
LE ET AL. (2019)	SPECIFIED	C^k	YES	NO	NO
DIFFUSION EMD	DYADIC	RANK DEPENDENT	YES	YES	YES

C. Algorithm Details

In this section we present two algorithms for computing the Diffusion EMD, the Chebyshev approximation method and the interpolative decomposition method. The Chebyshev method is more effective when the number of distributions is relatively small. The interpolative decomposition method is more effective when the number of distributions is large relative to the size of the manifold. We also detail how to subsample the normed space based on the spectrum of \mathbf{P} which can be approximated quickly.

First we define the approximate rank up to precision δ of a matrix $A \in \mathbb{R}^{n \times n}$ as:

$$R_\delta(A) := \# \{i : \sigma_i(A)/\sigma_0(A) \geq \delta\} \quad (17)$$

Where $\sigma_i(A)$ is the i th largest singular value of the matrix A . The approximate ranks of dyadic powers of \mathbf{P}^{2^k} are useful for determining the amount of subsampling to do at each scale either after an application of the Chebyshev method or during the computation of (approximate) dyadic powers of \mathbf{P} . We note that based on the density of singular values of \mathbf{P} , which is quickly computable using the algorithm presented in Dong et al. (2019), the approximate rank of all dyadic powers of \mathbf{P} can be calculated without computing powers of \mathbf{P} .

Chebyshev Approximation of Diffusion EMD We first note that \mathbf{P}^{2^k} can be computed spectrally using a filter on its eigenvalues. Let $\mathbf{M} = \mathbf{D}^{1/2} \mathbf{P} \mathbf{D}^{-1/2}$ be the symmetric conjugate of \mathbf{P} . Then \mathbf{M} is symmetric and has eigenvalues lying in the range $-1 \leq \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_n \leq 1$. \mathbf{M} can be decomposed into $\mathbf{U} \mathbf{\Sigma} \mathbf{U}^T$ for orthonormal \mathbf{U} and $\mathbf{\Sigma}$ a diagonal matrix of $[\lambda_0, \lambda_1, \dots, \lambda_n]$. We then express \mathbf{P}^{2^k} as a filter on the eigenvalues of the \mathbf{M} :

$$\mathbf{P}^{2^k} = \mathbf{D}^{-1/2} \mathbf{U} (\mathbf{\Sigma})^{2^k} \mathbf{U}^T \mathbf{D}^{1/2} \quad (18)$$

We compute the first J Chebyshev polynomials of $\mathbf{P}^j \boldsymbol{\mu}$ then use the polynomial filter on the eigenvalues $h(\sigma) = \sigma^{2^k}$ to compute diffusions of the distributions, and reweight these diffusion bins as specified in Algorithm 1. This algorithm requires J sparse matrix multiplications of \mathbf{P} and $\boldsymbol{\mu}$. For a total time complexity of $\tilde{O}(Jmn)$ when \mathbf{P} is sparse.

Interpolative Decomposition Approximation of Diffusion EMD Our second method proceeds by directly approximating multiplication by the matrix \mathbf{P}^{2^k} . The naive solution of computing dyadic powers of \mathbf{P} on the original basis quickly leads to a dense $n \times n$ matrix. Coifman & Maggioni (2006) observed that \mathbf{P}^{2^k} is of low rank, and therefore multiplication by \mathbf{P}^{2^k} can be approximated on a smaller basis. In that work they introduced a method of iteratively reducing the size of the basis using rank-revealing pivoted sparse QR decomposition, then computing $\mathbf{P}^{2^{k+1}} = \mathbf{P}^{2^k} \mathbf{P}^{2^k}$. By reducing the size of

Algorithm 2 Interpolative Decomposition

Input: $m \times n$ matrix A and number of subsamples k s.t. $k < \min\{m, n\}$.

Output: $m \times k$ matrix B and $k \times n$ matrix C s.t. $\|AP - BC\|_2 \leq \sqrt{4k(n-k) + 1}\sigma_{k+1}(A)$, where B consists of k columns of A .

Perform pivoted QR decomposition on A s.t. $AP = QR$, where P is a permutation matrix of A .

Where Q and R are decomposed as the following with Q_1 is a $m \times k$ matrix, Q_2 is $m \times (n-k)$, etc.

$$Q = [Q_1 \mid Q_2]; \quad R = \begin{bmatrix} R_{11} & R_{12} \\ \mathbf{0}_{k \times n-k} & R_{22} \end{bmatrix}$$

$$B \leftarrow Q_1 R_{11} \\ C \leftarrow [I_k \mid R_{11}^{-1} R_{12}]$$

the basis and orthogonalizing the basis is kept small and P^{2^k} is kept sparse on its corresponding basis. For sparse P this gives an algorithm that is $O(n \log^2 n)$ for computing dyadic powers of a sparse matrix.

Our algorithm follows the same lines as theirs however we use interpolative decomposition to reduce the size of the basis, and do not do so at early stages where the size of the basis may still be a significant fraction of n . Interpolative decomposition allows for an easy interpretation: we are selecting a well spaced set of points that acts as a basis for higher levels of P^{2^k} . For more details on interpolative decomposition we refer the reader to [Liberty et al. \(2007\)](#); [Bermanis et al. \(2013\)](#). The algorithm in [Coifman & Maggioni \(2006\)](#) also computes powers of P^{2^k} in the original basis by projecting the multiplication in the reduced basis back to the original. We do not bother computing $P^{2^k} \mu$ in the original basis, but instead use the reduced basis directly to compare distributions against. Denote the basis at level k as ϕ_k , which is really a subset of ϕ_{k-1} , then ϕ_k can be thought of as the centers of the bins at level k of our multiscale embedding. We only compare distributions at these representative centers rather than at all datapoints. This is summarized in [Algorithm 4](#), creating an embedding for each distribution whose length is dependent on the rank of P .

Sampling the Diffusions at Larger Scales Interpolative decomposition is an integral part of [Algorithm 4](#). However, it can also be useful in reducing the size of the distribution representations of [Algorithm 1](#) which uses the Chebyshev approximation. Without sampling the Chebyshev algorithm centers a bin at every datapoint at every scale. However, for larger scales this is unnecessary, and we can take advantage of the low rank of P^{2^k} with interpolative decomposition. In practice, we use the top 6 largest scales, in fact the distance shows very little sensitivity to small scales (see [Fig. 5\(d\)](#)). In fact, this representation can be compressed without a significant loss in performance as shown in [Fig. 5\(b,c\)](#). We apply the rank threshold to the spectrum of P^{2^k} , which can be computed from the density of states algorithm described above, to determine the number and location of centers to keep. With $\delta = 10^{-6}$ this reduces the number of centers from 60,000 to $< 4,000$ with a small loss in performance. Rather than selecting a fixed number of centers per scale as in previous methods, this allows us to vary the number of centers per scale as necessary to preserve the rank at each scale.

Time Complexity Here, we assume that P has at most $\tilde{O}(n) = O(n \log^c n)$ nonzero entries. In this case, interpolative decomposition in [Algorithm 2](#) has complexity $O(k \cdot m \cdot n \log n)$, and the randomized version in [Algorithm 3](#) has complexity $O(km + k^2 n)$ ([Liberty et al., 2007](#)). [Algorithm 1](#) has complexity $O(Jmn \log^c n)$ which is dominated by calculation of the J powers of $P^j \mu$. The computation time of [Algorithm 4](#) is dominated by the first randomized interpolative decomposition step. When we set $\gamma = O(\log^c n)$, which controls when the application of RandID occurs, this gives a total time complexity for [Algorithm 4](#) of $\tilde{O}(mn) = O(n \log^{2c} n + mn \log^b n)$, where b and c are constants controlled by γ . As γ increases, b increases and c decreases. The first term is dominated by the first interpolative decomposition setup, and the second is the calculation of $\mu^{(2^k)}$. It follows that when $m \gg n$, it is helpful to set γ larger for a better tradeoff of b and c . In practice we set $\gamma = n/10$ which seems to offer a reasonable tradeoff between the two steps. For small m in practice [Algorithm 1](#) is significantly faster than [Algorithm 4](#).

C.1. Gradients of the Diffusion EMD

In this section we will develop the computation of the gradient of the Diffusion EMD $W_{\alpha, K}(X_i, X_j)$, we will show that its gradient depend only on the gradient of the Gaussian kernel matrix K_ϵ which is easy to compute. We start by recalling

Algorithm 3 Randomized Interpolative Decomposition: RandID(A, j, k, l)

Input: $m \times n$ matrix A , random sample sizes j, l , and number of subsamples k s.t. $\min\{m, n\} > j > k > l$.

Output: $m \times k$ matrix B and $k \times n$ matrix C s.t. $\|A - BC\|_2 \leq \sqrt{4k(n-k)} + 1\sigma_{k+1}(A)$ w.h.p., where B consists of k columns of A .

$G_1 \leftarrow \mathcal{N}(0, 1)^{j \times m}$

$W \leftarrow G_1 A$

Perform the (deterministic) interpolative decomposition on W in Algorithm 2 to get \tilde{B}, C s.t. $\|W - \tilde{B}C\|_2$ small.

$G_2 \leftarrow \mathcal{N}(0, 1)^{l \times j}$

$W_2 \leftarrow G_2 W$

$\tilde{B}_2 \leftarrow G_2 \tilde{B}$

Find the k indices of \tilde{B}_2 that are approximately equal to columns in W_2 , $[i_1, i_2, \dots, i_k]$.

$\mathbf{b} \leftarrow [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_K]$

some basic facts about the gradient of matrices,

1. For a matrix P and a vector μ we have $\nabla \|P\mu\|_1 = \text{sign}(P\mu)\mu^T \nabla P$.
2. For an invertible matrix P , the gradient of the inverse of P is $\nabla(P^{-1}) = -P^{-1} \nabla(P) P^{-1}$.
3. Let $n \in \mathbb{N}^*$ the gradient of P^n is

$$\nabla P^n = \sum_{\ell=0}^{n-1} P^\ell \nabla(P) P^{n-\ell-1}.$$

Using the definition of the Diffusion EMD, we have

$$\begin{aligned} \nabla W_{\alpha, K}(X_i, X_j) &= \nabla \left\| \mathbf{P}_\epsilon^{2^K} (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \right\|_1 + \sum_{k=0}^{K-1} 2^{-(K-k-1)\alpha} \nabla \left\| \left(\mathbf{P}_\epsilon^{2^{k+1}} - \mathbf{P}_\epsilon^{2^k} \right) (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \right\|_1 \\ &= \text{sign} \left(\mathbf{P}_\epsilon^{2^K} (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \right) (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i)^T \nabla \left(\mathbf{P}_\epsilon^{2^K} \right) + \sum_{k=0}^{K-1} 2^{-(K-k-1)\alpha} \text{sign} \left(\mathbf{P}_\epsilon^{2^{k+1}} - \mathbf{P}_\epsilon^{2^k} \right) (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i)^T \nabla \left(\mathbf{P}_\epsilon^{2^{k+1}} - \mathbf{P}_\epsilon^{2^k} \right). \end{aligned}$$

The last formula tells us that the gradient of $W_{\alpha, K}(X_i, X_j)$ depends only on the gradient of the powers of \mathbf{P}_ϵ which in turn can be expressed in terms of the gradients $\nabla \mathbf{P}_\epsilon$. We have:

$$\begin{aligned} \nabla \mathbf{P}_\epsilon &= \nabla (D_\epsilon^{-1} M_\epsilon D_\epsilon^{-1}) \\ &= \nabla (D_\epsilon^{-1}) M_\epsilon + D_\epsilon^{-1} \nabla (M_\epsilon) \\ &= -D_\epsilon^{-1} \nabla (D_\epsilon) D_\epsilon^{-1} M_\epsilon + D_\epsilon^{-1} \nabla (M_\epsilon) \\ &= -D_\epsilon^{-1} \nabla (\text{diag}(M_\epsilon \mathbf{1})) D_\epsilon^{-1} M_\epsilon + D_\epsilon^{-1} \nabla (M_\epsilon) \\ &= -D_\epsilon^{-1} \text{diag}(\nabla(M_\epsilon) \mathbf{1}) D_\epsilon^{-1} M_\epsilon + D_\epsilon^{-1} \nabla (M_\epsilon), \end{aligned}$$

therefore in order to compute ∇P_ϵ we need to compute ∇M_ϵ ,

$$\begin{aligned} \nabla M_\epsilon &= \nabla (Q_\epsilon^{-1} K_\epsilon Q_\epsilon^{-1}) \\ &= \nabla (Q_\epsilon^{-1}) K_\epsilon Q_\epsilon^{-1} + Q_\epsilon^{-1} \nabla (K_\epsilon) Q_\epsilon^{-1} + Q_\epsilon^{-1} K_\epsilon \nabla (Q_\epsilon^{-1}) \\ &= -Q_\epsilon^{-1} \nabla (Q_\epsilon) Q_\epsilon^{-1} K_\epsilon Q_\epsilon^{-1} + Q_\epsilon^{-1} \nabla (K_\epsilon) Q_\epsilon^{-1} - Q_\epsilon^{-1} K_\epsilon Q_\epsilon^{-1} \nabla (Q_\epsilon) Q_\epsilon^{-1} \\ &= -Q_\epsilon^{-1} \nabla (\text{diag}(K_\epsilon \mathbf{1})) Q_\epsilon^{-1} K_\epsilon Q_\epsilon^{-1} + Q_\epsilon^{-1} \nabla (K_\epsilon) Q_\epsilon^{-1} - Q_\epsilon^{-1} K_\epsilon Q_\epsilon^{-1} \nabla (\text{diag}(K_\epsilon \mathbf{1})) Q_\epsilon^{-1} \\ &= -Q_\epsilon^{-1} \text{diag}(\nabla(K_\epsilon) \mathbf{1}) Q_\epsilon^{-1} K_\epsilon Q_\epsilon^{-1} + Q_\epsilon^{-1} \nabla (K_\epsilon) Q_\epsilon^{-1} - Q_\epsilon^{-1} K_\epsilon Q_\epsilon^{-1} \text{diag}(\nabla(K_\epsilon) \mathbf{1}) Q_\epsilon^{-1} \end{aligned}$$

Algorithm 4 Interpolative Decomposition diffusion densities $T_{id}(\mathbf{K}, \boldsymbol{\mu}, K, \alpha, \delta, \gamma)$

Input: $n \times n$ graph kernel \mathbf{K} , $n \times m$ distributions $\boldsymbol{\mu}$, maximum scale K , snowflake constant α , rank threshold δ , and basis recomputation threshold γ .

Output: Distribution embeddings \mathbf{b}

$\mathbf{Q} \leftarrow \text{Diag}(\sum_i \mathbf{K}_{ij})$

$\mathbf{K}^{norm} \leftarrow \mathbf{Q}^{-1} \mathbf{K} \mathbf{Q}^{-1}$

$\mathbf{D} \leftarrow \text{Diag}(\sum_i \mathbf{K}_{ij}^{norm})$

$\mathbf{M} \leftarrow \mathbf{D}^{-1/2} \mathbf{K}^{norm} \mathbf{D}^{-1/2}$

$\boldsymbol{\mu}^{(2^0)} \leftarrow \mathbf{P} \boldsymbol{\mu} \leftarrow \mathbf{D}^{-1/2} \mathbf{M} \mathbf{D}^{1/2} \boldsymbol{\mu}$

$\boldsymbol{\mu}_0 \leftarrow \boldsymbol{\mu}$

$n_0 \leftarrow n$

for $k = 1$ **to** K **do**

if $R_\delta(\mathbf{M}^{2^k}) < \gamma$ **then**

$\mathbf{B}_k, \mathbf{C}_k \leftarrow \text{RandID}(\mathbf{M}_{k-1}, R_\delta(\mathbf{M}^{2^k}) + 8, R_\delta(\mathbf{M}^{2^k}), 5)$

$n_k \leftarrow R_\delta(\mathbf{M}^{2^k})$

else

$\mathbf{B}_k, \mathbf{C}_k, n_k \leftarrow \mathbf{B}_{k-1}, \mathbf{I}_{n_{k-1}}, n_{k-1}$

end if

$\mathbf{M}_k \leftarrow \mathbf{C}_k \mathbf{M}_{k-1} \mathbf{M}_{k-1}^T \mathbf{C}_k^T$

$\boldsymbol{\mu}_k \leftarrow \mathbf{C}_k \boldsymbol{\mu}_{k-1}$

$\boldsymbol{\mu}^{(2^k)} \leftarrow \mathbf{D}^{-1/2} \mathbf{M}_k \mathbf{D}^{1/2} \boldsymbol{\mu}_k$

$\mathbf{b}_{k-1} \leftarrow 2^{(K-k-1)\alpha} (\boldsymbol{\mu}^{(2^k)} - \mathbf{C}_k \boldsymbol{\mu}^{(2^{k-1})})$

end for

$\mathbf{b}_K \leftarrow \boldsymbol{\mu}^{(2^K)}$

$\mathbf{b} \leftarrow [\mathbf{b}_0, \mathbf{b}_1, \dots, \mathbf{b}_K]$

hence the gradient of M_ϵ can be written only in terms of the gradient of K_ϵ . This process shows that the gradient of the EMD diffusion can be expressed in terms of the gradient of the Gaussian kernel only, which may make it useful in future applications where fast gradients of the earth mover's distance are necessary. For example, in distribution matching where previous methods use gradients with respect to the Sinkhorn algorithm (Frogner et al., 2015), which scales with $O(n^2)$. In this application, computation of gradients of Diffusion EMD would be $O(n)$. What makes this possible is a smooth kernel and smooth density bins over the graph. In contrast to Diffusion EMD, multiscale methods that use non-smooth bins have non-smooth gradients with respect to K_ϵ , and so are not useful for gradient descent type algorithms. We leave further investigation into the gradients of Diffusion EMD to future work.

D. Experiment details

We ran all experiments on a 36 core Intel(R) Xeon(R) CPU E5-2697 v4 @ 2.30GHz with 512GB of RAM. We note that our method is extremely parallelizable, consisting of only matrix operations, thus a GPU implementation could potentially speed up computation. For fair comparison we stick to CPU implementations and leave GPU acceleration to future work. We provide custom implementations of all methods in the `DiffusionEMD` package, which is available on Github. We base our Sinkhorn implementation on that in the python optimal transport package (POT) (Flamary et al., 2021), and our tree methods off of the sklearn nearest neighbors trees (Pedregosa et al., 2011).

D.1. Metrics used

P@10. P@10 is a metric that measures the overlap in the 10-nearest neighbors between the nearest neighbors of the EMD method and the ground truth nearest neighbors, which is calculated using exact OT on the geodesic ground distance. P@10 is useful for distinguishing the quality of distance calculations locally, whereas the next metric measures the quality more globally. This may be particularly important in applications where only the quality of nearest neighbors matters as in Sect 4.6, or in applications of Wasserstein nearest neighbor classification (Indyk & Thaper, 2003; Backurs et al., 2020).

Spearman- ρ . The Spearman’s rank correlation coefficient (also known as Spearman- ρ) measures the similarity in the order of two rankings. This is useful in our application because while the 1-Wasserstein and 2-Wasserstein may have large distortion as metrics, they will have a high Spearman- ρ in all cases. Spearman- ρ is defined for two rankings of the data where d_i is the difference in the rankings of each for each datapoint i as

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)}$$

This is useful for quantifying the global ranking of distributions in the graph, which shows the more global behavior of EMD methods. Since locally manifolds with low curvature look very similar to \mathbb{R}^d , where d is the intrinsic dimension of the manifold, for distributions very close together on the manifold using a Euclidean ground distance may be acceptable, however for distributions far away from each other according to a geodesic ground distance, approximation with a Euclidean ground distance may perform poorly.

D.2. Line Example

In Fig 2 we used a cluster tree with 4 levels of 4 clusters for the cluster tree, and a quad tree of depth 4. These are mostly for illustrative purposes so we there was no parameter search for ClusterTree or QuadTree. As the number of trees gets larger ClusterTree and QuadTree start to converge to some smooth solution. ClusterTree has a number of bumps caused by the number of clusters parameter that do not disappear even with many trees. Quadtree converges to a smooth solution as the number of trees gets large in this low dimensional example.

D.3. Swiss Roll

On the swiss roll example we compared Diffusion EMD to ClusterTree, QuadTree, and the convolutional Sinkhorn method. Here we describe the parameter setup in each of these methods for Figure 3, 4, and 5.

For Fig. 3 we chose parameters that approximately matched the amount of time between methods on this dataset to the Diffusion EMD with default settings using the Chebyshev algorithm. This was 150 cluster trees of depth 8 with 3 cluster, and 20 quad trees of depth 4. We noticed that the speed of quadtree diminishes exponentially with dimension.

For Fig. 4 we did a grid search over parameters for QuadTree and Cluster Tree to find the best performing settings of depth and number of clusters. We searched over depth $d \in [2 \dots 10]$ and number of clusters $C \in [2 \dots 5]$ for ClusterTree. We found a depth of 4 for QuadTree and 3 Clusters of depth 5 for ClusterTree gave the best tradeoff of performance vs. time on this task. These parameters were tried while fixing the number of trees to 10. We fixed this depth and number of clusters for subsequent experiments varying the number of trees in the range $[1 \dots 100]$ for QuadTree and in the range $[1 \dots 15]$ for ClusterTree. The maximum of this range exceeded the time allotment of Diffusion EMD in both cases. For the convolutional Sinkhorn method we fixed parameter $t = 50$ for H_t , and fixed the maximum number of iterations in the range $[1 \dots 100]$. This took orders of magnitude longer than the multiscale methods in all cases, with comparable performance to Diffusion EMD on the P@10 metric and better performance using the Spearman- ρ metric. We note that a similar multi-step refinement scheme as implemented in Backurs et al. (2020) could be used here.

D.4. MNIST

To construct the Spherical MNIST dataset (Cohen et al., 2017), we projected the 70,000 images in classical MNIST onto a spherical graph containing 25,088 nodes. Every pixel in each MNIST image was projected onto four nodes on the northern hemisphere. These projections were treated as signals over a spherical graph.

We ran Diffusion EMD, ClusterTree, and QuadTree on this dataset to generate embeddings in L1 space. Diffusion EMD was run with the diffusion wavelet approximations described in the paper, using an epsilon value of 0.001. To best compare accuracy, we chose parameters for QuadTree and ClusterTree that approximately matched the runtime of Diffusion EMD. For QuadTree, we used a depth of 20 and 10 clusters. For ClusterTree, we used a depth of 5. However, we note that both ClusterTree and QuadTree experienced a variance of 10-30 minutes between subsequent runtimes with the same parameters, probably due to differences of random initialization. We did not run Convolutional Sinkhorn on this dataset, as the projected runtime exceeded 24 hours.

We calculated a ground truth EMD between 100 of the projected MNIST digits using the exact EMD implementation of Python Optimal Transport. More accurate comparisons might be obtained using more points, but computation of this 100 by

100 distance matrix took over 24 hours, necessitating severe subsampling, and once again highlighting the need for fast and accurate approximations of EMD.

For each algorithm, we obtained an accuracy score by running a kNN classifier on the embeddings using an even train/test split considering only the nearest neighbors. We then computed the Spearman ρ and P@10 score between each method and the exact EMD ground truth.

D.5. Single cell COVID-19 Dataset

One hundred sixty eight patients with moderate to severe COVID-19 (Marshall et al., 2020) were admitted to YNHH and recruited to the Yale Implementing Medical and Public Health Action Against Coronavirus CT (IMPACT) study. From each patient, blood samples were collected to characterize patient cellular responses across timepoints to capture the spectrum of disease. In total, the composition of peripheral blood mononuclear cell (PBMC) was measured by a myeloid-centric flow cytometry on 210 samples. Finally, clinical data was extracted from the electronic health record corresponding to each biosample timepoint to allow for clinical correlation of findings as shown previously. In the main analysis, poor or adverse outcomes were defined by patient death, while good outcomes were defined by patient survived. In order to analyze the 22 million cells generated in this study, we performed k-means clustering, setting k to 27,000, and took the resultant cluster centroids as a summarization of the cellular state space as done previously in (Kuchroo et al., 2020).

To test the quality of the organization of patients using EMD methods, we test the Laplacian smoothness on the 10-nearest neighbors graph created for each method. The Laplacian smoothness of a function on the nodes is denoted $f^T L f$ where $L = D - A$ is the (combinatorial) Laplacian of the adjacency matrix A . This is equivalent to computing the following sum:

$$\sum_{(i,j) \in E} (f_i - f_j)^2$$

where E is the edge set of the kNN graph. This measures the sum of squared differences of f along the edges of the kNN graph, which measures the smoothness of f over the graph.

Various cell types have been shown to be associated with mortality from COVID-19 infection, including CD16⁺ Neutrophils, T cells and non-classical monocytes. Previously, T cells-to-CD16⁺ Neutrophils ratios have been reported to be predict of outcome, with low ratios predicting mortality and high ratios predicting survival (Li et al., 2020). Even outside of ratios, the absolute counts to T cells (Chen & Wherry, 2020) and CD16⁺ Neutrophils (Meizlish et al.) have been shown to be associated with mortality. Finally, non-classical monocytes have also been shown to be enriched in patients with more severe outcomes (Pence, 2020), further supporting our findings.

Supplement References

- Chen, Z. and Wherry, E. J. T cell responses in patients with COVID-19. *Nature Reviews Immunology*, 20(9):529–536, 2020.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N. T. H., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D. J., Tavenard, R., Tong, A., and Vayer, T. POT: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- Kolouri, S., Zou, Y., and Rohde, G. K. Sliced Wasserstein kernels for probability distributions. In *the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5258–5267, 2016.
- Li, X., Liu, C., Mao, Z., Xiao, M., Wang, L., Qi, S., and Zhou, F. Predictive values of neutrophil-to-lymphocyte ratio on disease severity and mortality in COVID-19 patients: a systematic review and meta-analysis. *Critical Care*, 24(1), 2020.
- Marshall, J. C., Murthy, S., Diaz, J., Adhikari, N. K., Angus, D. C., Arabi, Y. M., Baillie, K., Bauer, M., Berry, S., Blackwood, B., Bonten, M., Bozza, F., Brunkhorst, F., Cheng, A., Clarke, M., Dat, V. Q., de Jong, M., Denholm, J., Derde, L., Dunning, J., Feng, X., Fletcher, T., Foster, N., Fowler, R., Gobat, N., Gomersall, C., Gordon, A., Glueck, T., Harhay, M., Hodgson, C., Horby, P., Kim, Y., Kojan, R., Kumar, B., Laffey, J., Malvey, D., Martin-Loeches, I., McArthur, C., McAuley, D., McBride, S., McGuinness, S., Merson, L., Morpeth, S., Needham, D., Netea, M., Oh, M.-D., Phyu, S., Piva, S., Qiu, R., Salisu-Kabara, H., Shi, L., Shimizu, N., Sinclair, J., Tong, S., Turgeon, A., Uyeki, T., van de Veerdonk, F., Webb, S., Williamson, P., Wolf, T., and Zhang, J. A minimal common outcome measure set for COVID-19 clinical research. *The Lancet Infectious Diseases*, 20(8):e192–e197, 2020.

Meizlish, M. L., Pine, A. B., Bishai, J. D., Goshua, G., Nadelmann, E. R., Simonov, M., Chang, C.-H., Zhang, H., Shallow, M., Bahel, P., Owusu, K., Yamamoto, Y., Arora, T., Atri, D. S., Patel, A., Gbyli, R., Kwan, J., Won, C. H., Cruz, C. D., Price, C., Koff, J., King, B. A., Rinder, H. M., Wilson, F. P., Hwa, J., Halene, S., Damsky, W., van Dijk, D., Lee, A. I., and Chun, H. J. A neutrophil activation signature predicts critical illness and mortality in COVID-19. medRxiv, DOI: 10.1101/2020.09.01.20183897.

Pence, B. D. Severe COVID-19 and aging: are monocytes the key? *GeroScience*, 42(4):1051–1061, 2020.

Wang, F.-Y. et al. Sharp explicit lower bounds of heat kernels. *Annals of Probability*, 25(4):1995–2006, 1997.