

Gromov-Wasserstein Averaging in a Riemannian Framework

Samir Chowdhury
Stanford University

Department of Psychiatry and Behavioral Sciences

samirc@stanford.edu

Tom Needham

Florida State University
Department of Mathematics

tneedham@fsu.edu

Abstract

We introduce a theoretical framework for performing statistical tasks—including, but not limited to, averaging and principal component analysis—on the space of (possibly asymmetric) matrices with arbitrary entries and sizes. This is carried out under the lens of the Gromov-Wasserstein (GW) distance, and our methods translate the Riemannian framework of GW distances developed by Sturm into practical, implementable tools for network data analysis. Our methods are illustrated on datasets of letter graphs, asymmetric stochastic blockmodel networks, and planar shapes viewed as metric spaces. On the theoretical front, we supplement the work of Sturm by producing additional results on the tangent structure of this “space of spaces”, as well as on the gradient flow of the Fréchet functional on this space.

1. Introduction

In a variety of data analysis contexts, one often obtains matrices which are square and asymmetric. Often these matrices arise when studying *networks* [20] where the relationships between nodes cannot be measured directly, but have to be inferred from the activity of the nodes themselves. This is the case for biological networks such as the brain, gene regulation pathways, and protein interaction networks.

Inspired by this connectivity paradigm, we refer to arbitrary square matrices as *networks*. The row/column labels are referred to as *nodes*, and the matrix entries are referred to as *edge weights*. Such matrix datasets commonly arise in many other use cases. For example, a practitioner is typically confronted with an $n \times p$ data matrix X where each row is an observation and each column is a variable, from which the covariance matrix is formed. If the dataset is Euclidean, then there is a well-understood duality between the covariance of the variables and the pairwise distances between the observations. More generally, the dataset could be sampled from a Riemannian manifold (or from a distribution whose high density regions live near such a manifold), and the distances between the points could be given by the geodesic

distances on the manifold. Even more generally, it may be the case that the data is sampled from a Finsler manifold, and one has access to the quasimetric defined by the asymmetric length structure of the manifold. This may occur when one is sampling data from a dynamical system driven by some potential function: the asymmetry arises because traveling up the potential function is costlier than traveling down [3].

In the interest of performing statistics on such data, it is natural to ask how one obtains a *mean* of such matrices. Simply taking a coordinatewise mean does not work in many cases, e.g. when the matrices are of different sizes or are unlabeled. In such situations, one needs to first perform an alignment/registration task that optimally matches the nodes of one network to the nodes of the other. If the matrices are the same size, then the most obvious approach would be to search for an optimal permutation to match nodes between the networks. However, this idea is too restrictive as real-world datasets are frequently of unequal size. Moreover, for large matrices, searching over permutations is prohibitively computationally expensive. For these reasons, one introduces the idea of “probabilistic matchings”. Here, each node is assigned a weight, so that the total weight of the network is one (i.e., a probability measure is assigned to the nodes of the network). Instead of searching over permutations to match nodes between a pair of networks, we can then instead search over the convex set of *couplings* of their probability measures (that is, joint probability distributions whose marginals agree with the original distributions on the input networks). This is the essential idea of *Gromov-Wasserstein distance*, which is defined below.

The goal of this paper is to introduce a theoretical framework for statistical computations on the space of networks. This is achieved by fusing theoretical results on Gromov-Wasserstein distance [26], algorithms for statistics on Riemannian manifolds [22], and recent algorithmic advances for the computation of Gromov-Wasserstein distance [23]. Using this framework, we are able to perform not just averaging, but a plethora of statistical tasks such as principal component analysis and support vector machine classification.

1.1. Previous Work

A *metric measure (mm) space* is a compact metric space endowed with a Borel probability measure. *Gromov-Wasserstein (GW) distance* was first introduced as a metric on the space of all (isomorphism classes of) mm spaces. Theoretical aspects of the GW distance were explored in [17, 18, 26]. The work in [17, 18] was already focused on applications to object matching, while [26] explored the Riemannian-like structures induced by GW distance.

In recent years, GW distance has garnered interest in data science communities as a way to compare unlabeled datasets, or datasets containing samples from different ambient spaces. For example, GW distance has been used to explore a variety of network datasets [14], as a metric alignment layer in deep learning algorithms for object classification [11], to align word embedding spaces for translation applications [1], for several tasks in analysis of large graphs and networks [30, 29], and has been incorporated into generative models across incomparable spaces [4]. Several specialized variants of GW distance have also been recently introduced [19, 27, 28].

The problem of computing GW distance was studied from the algorithmic viewpoint in [23], where a projected gradient descent algorithm was introduced. The main focus was on using GW distance to compute a *Fréchet mean with prescribed size* of distance (or kernel) matrices. The main idea of the present paper is to recast the work in [23] using the theoretical Riemannian framework of [26] together with statistical algorithms on Riemannian manifolds [22]. By using this viewpoint, we are able to generalize the work of [23] to a gradient flow that theoretically prescribes the required size for a Fréchet mean, while providing a flexible general framework for machine learning tasks on network-valued datasets. This includes the case of asymmetric networks.

An approach similar to that of the present paper to studying statistics on the space of graphs via Riemannian geometry was initiated in [15, 16]. Recently, these ideas were applied to formulate a theory of statistical shape analysis of embedded graphs in [13]. These works perform analysis on graph space by aligning graph nodes over permutations or “hard matchings”, whereas our approach aligns networks via measure couplings or “soft matchings”. The differences between these theories are interesting, and we expect that the correct formalism to use is highly dependent on the particular application.

1.2. Contributions

Our specific contributions are as follows. We first provide the gradient of the GW functional on asymmetric networks. This complements a similar result for symmetric matrices in [23]. On the metric geometry side, we provide a concrete exposition of the tangent space structure on this space of asymmetric networks and of the construction of geodesics in the space of networks. This includes the, to our knowledge,




Figure 1. Networks from Example 2. **Top row:** One-node network X and two-node network Y . **Bottom left:** Blown-up form \hat{X} of X . Dotted line shows networks to align. **Inset:** “average” of X and Y .

first computationally feasible algorithm to produce Sturm geodesics. We explicitly formulate the iterative Fréchet mean algorithm of Pennec as gradient descent of the Fréchet functional on the space of networks. The tangent structure provides a framework for vectorizing collections of networks in order to apply standard ML algorithms. We exemplify this by performing averaging and principal component analysis on a database of planar shapes. Our methods can also be used for network compression, and we illustrate this on a toy example of an asymmetric stochastic blockmodel network.

2. Preliminaries on the GW distance and Fréchet means

2.1. Networks and the GW distance

A *measure network* is a triple (X, ω_X, μ_X) where X is a Polish space (i.e. separable, completely metrizable), μ_X is a fully supported Borel probability measure, and $\omega_X : X \times X \rightarrow \mathbb{R}$ is a square integrable function. The collection of all networks is denoted \mathcal{N} . When no confusion will arise, we abuse notation and denote the triple (X, ω_X, μ_X) by X .

The notion of a measure network is quite general and includes several types of spaces which arise in applications. A graph can be represented as a measure network: X is set of n nodes, while ω_X provides relational information and could be represented by a weighted adjacency matrix, a graph Laplacian, or a matrix of graph distances (for connected graphs or strongly connected digraphs)—see Example 2. The probability measure μ_X can be taken to be uniform, or could more generally give higher weight to nodes deemed more important by a particular application. The notion of a measure network is also a strict generalization of that of a metric measure space, as defined in the previous section. For a finite metric measure space, ω_X can be represented as by its distance matrix. Our definition also includes infinite spaces—this is necessary for certain theoretical completeness results

to hold, but the reader more interested in applications can safely restrict their attention to finite spaces.

A *coupling* between two probability measures μ_X and μ_Y supported on Polish spaces X and Y is a probability measure on $X \times Y$ with marginals μ_X and μ_Y . Stated differently, $\mu(A \times Y) = \mu_X(A)$ for all $A \in \text{Borel}(X)$ and $\mu(X \times B) = \mu_Y(B)$ for all $B \in \text{Borel}(Y)$. The collection of all such couplings is denoted $\mathcal{C}(\mu_X, \mu_Y)$.

Intuitively, for finite measure networks X and Y and a coupling μ , the value of the coupling $\mu(x_i, y_j)$ can be understood as an assignment of a certain proportion of the mass $\mu_X(x_i)$ to the point y_j . In the graph setting, this is a ‘‘soft matching’’ of the nodes of X and Y . One then wishes to find a soft matching which reflects similarity of X and Y as well as possible. This is formalized by the notion of *distortion*, defined below.

Given two networks (X, ω_X, μ_X) and (Y, ω_Y, μ_Y) in \mathcal{N} , one defines the *distortion functional* to be the map

$$\begin{aligned} \text{dis} : \mathcal{C}(\mu_X, \mu_Y) &\rightarrow \mathbb{R} \\ \mu &\mapsto \|\omega_X - \omega_Y\|_{L^2(\mu \otimes \mu)}. \end{aligned}$$

More explicitly, $\text{dis}(\mu)$ is the quantity

$$\int |\omega_X(x, x') - \omega_Y(y, y')|^2 d\mu(x', y') d\mu(x, y),$$

where the integral is taken over the space $X \times Y \times X \times Y$. The *Gromov-Wasserstein distance* between networks (X, ω_X, μ_X) and (Y, ω_Y, μ_Y) is then defined by

$$d_{\mathcal{N}}(X, Y) := \frac{1}{2} \inf_{\mu \in \mathcal{C}(\mu_X, \mu_Y)} \text{dis}(\mu).$$

One may similarly define the p -GW distance by taking the L^p norm, but the constructions in this paper rely on the special structure of the L^2 case.

The following lemma shows that the infimum in the definition above is always achieved. Minimizers of dis are referred to as *optimal couplings*. The proof follows directly from [26, Lemma 1.2].

Lemma 1 (Optimality of couplings, [26]). *Let $(X, \omega_X, \mu_X), (Y, \omega_Y, \mu_Y) \in \mathcal{N}$. Then there always exists a minimizer of dis in $\mathcal{C}(\mu_X, \mu_Y)$.*

In the finite setting, the notation of [23] admits some useful insights into this minimization problem. We present this notation now. Let $L : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a loss function. In our case, this will always be defined as $L(a, b) := |a - b|^2$. Next we switch to matrix notation: given a finite space $X = \{x_1, x_2, \dots, x_n\}$, we write X_{ik} to denote $\omega_X(x_i, x_k)$ for $1 \leq i, k \leq n$. Suppose Y is another finite space of size m . The collection of couplings $\mathcal{C}(\mu_X, \mu_Y)$ then consists of $n \times m$ matrices $C = (C_{ij})_{ij}$ such that

$$\sum_i C_{ij} = \mu_Y(y_j) \quad \text{and} \quad \sum_j C_{ij} = \mu_X(x_i).$$

Then one defines the 4-way tensor $\mathcal{L}(X, Y) := (L(X_{ik}, Y_{jl}))_{ijkl}$. Given a 4-way tensor \mathcal{L} and a matrix $(C_{ij})_{ij}$, one defines the tensor-matrix multiplication

$$\mathcal{L} \otimes C := \left(\sum_{kl} \mathcal{L}_{ijkl} C_{kl} \right)_{ij}.$$

Next, given two real-valued matrices A and B of the same dimensions, one writes $\langle A, B \rangle$ to denote the Frobenius inner product $\sum_{ij} A_{ij} B_{ij}$. As observed in [23], the GW distance between two finite networks X and Y can be written as:

$$d_{\mathcal{N}}(X, Y) = \frac{1}{2} \min_{C \in \mathcal{C}(\mu_X, \mu_Y)} \langle \mathcal{L}(X, Y) \otimes C, C \rangle^{1/2}.$$

For the reader’s convenience, we verify that the dimensions are consistent. If X is an n -point space and Y is an m -point space. Then C is an $n \times m$ coupling matrix, $\mathcal{L}(X, Y)$ is an $n^2 \times m^2$ tensor, and the product $\mathcal{L}(X, Y) \otimes C$ is an $n \times m$ matrix. An alternative matrix formulation [25] of the term inside the min is the following:

$$\left(\langle \mu_X, X.^{\wedge 2} \mu_X \rangle + \langle \mu_Y, Y.^{\wedge 2} \mu_Y \rangle - 2 \text{tr}(C^T X^T C Y) \right)^{\frac{1}{2}}. \quad (1)$$

Here $.^{\wedge 2}$ denotes the elementwise square. While this formulation is well-known, we provide details in the appendix for the reader’s convenience.

2.2. Weak isomorphism: From transport plans to transport maps via blow-ups

The space $(\mathcal{N}, d_{\mathcal{N}})$ is a pseudometric space. Networks (X, ω_X, μ_X) and (Y, ω_Y, μ_Y) satisfy $d_{\mathcal{N}}(X, Y) = 0$ if and only if there exists a Borel probability space (Z, μ_Z) with maps $\pi_X : Z \rightarrow X$ and $\pi_Y : Z \rightarrow Y$ such that the pushforward measures satisfy $(\pi_X)_{\#} \mu_Z = \mu_X$, $(\pi_Y)_{\#} \mu_Z = \mu_Y$, and the pullbacks $(\pi_X)^* \omega_X$, $(\pi_Y)^* \omega_Y$ satisfy $\|(\pi_X)^* \omega_X - (\pi_Y)^* \omega_Y\|_{\infty} = 0$, where $(\pi_X)^* \omega_X(z, z') := \omega_X(\pi_X(z), \pi_X(z'))$. In this case, X and Y are said to be *weakly isomorphic* and write $X \cong^w Y$. The space Z is referred to as a *common expansion* of X and Y . We write $[X] = [X, \omega_X, \mu_X]$ to denote the weak isomorphism class of $X = (X, \omega_X, \mu_X)$ in \mathcal{N} . The collection of equivalence of classes of networks will be denoted $[\mathcal{N}]$. See also [8] for more details on weak isomorphism.

In the case of finite networks, weak isomorphism is especially useful as it allows one to convert transport plans (i.e., couplings between the networks’ measures) to transport maps (i.e., measure-preserving maps between networks). This observation plays a key role in both our theory and algorithms. We present this construction next.

Definition 1 (Blowups). Let X, Y be finite networks, and let $\mu \in \mathcal{C}(\mu_X, \mu_Y)$. Let $\mathbf{u} := (u_x)_{x \in X}$ be a vector where

$u_x := |\{y \in Y : \mu(x, y) > 0\}|$. Also define $\mathbf{v} := (v_y)_{y \in Y}$ by setting $v_y := |\{x \in X : \mu(x, y) > 0\}|$. Next define $X[\mathbf{u}]$ to be the node set $\bigcup_{x \in X} \{(x, i) : 1 \leq i \leq u_x\}$. Fix $x \in X$ and let y_1, y_2, \dots, y_{u_x} denote the $y \in Y$ such that $\mu(x, y) > 0$. Define $\mu_{X[\mathbf{u}]}((x, i)) := \mu(x, y_i)$. Finally, for $x, x' \in X$ and $1 \leq i \leq u_x, 1 \leq j \leq u_{x'}$, define $\omega_{X[\mathbf{u}]}((x, i), (x', j)) = \omega_X(x, x')$. Similarly define $(Y[\mathbf{v}], \omega_{Y[\mathbf{v}]}, \mu_{Y[\mathbf{v}]})$. The crux of this construction is that while $X[\mathbf{u}], Y[\mathbf{v}]$ are weakly isomorphic to X and Y , respectively, the initial transport plan μ naturally expands to a transport map from $X[\mathbf{u}]$ to $Y[\mathbf{v}]$. We refer to the process of constructing $X[\mathbf{u}]$ from X as a *blow-up*.

Definition 2 (Alignment). Let X, Y be finite networks on n and m nodes, respectively, and let μ be an optimal coupling. We refer to the $n \times m$ binary matrix $\mathbf{1}_{\mu > 0}$ as the *binarization* of μ : this matrix has the same dimensions as μ , has a 1 where $\mu > 0$, and 0 elsewhere. By taking appropriate blow-ups, we obtain (possibly enlarged) networks \hat{X}, \hat{Y} and an optimal coupling $\hat{\mu}$ such that the binarization $\mathbf{1}_{\hat{\mu} > 0}$ of $\hat{\mu}$ is a permutation matrix. Then we may *align* \hat{Y} to \hat{X} by defining $\hat{Y} \leftarrow \mathbf{1}_{\hat{\mu} > 0} \hat{Y} \mathbf{1}_{\hat{\mu} > 0}^T$. The corresponding realignment of the optimal coupling is given by $\hat{\mu} \leftarrow \mathbf{1}_{\hat{\mu} > 0} \hat{\mu}$. Note that we then have $\hat{\mu} = \text{diag}(\mu_{\hat{X}})$. We refer to this process of blowing up and realigning as *aligning* Y to X . After aligning, $d_{\mathcal{N}}(X, Y)$ is given by

$$\sum_{i,j=1}^n |\omega_{\hat{X}}(x_i, x_j) - \omega_{\hat{Y}}(y_i, y_j)|^2 \mu_{\hat{X}}(x_i) \mu_{\hat{X}}(x_j).$$

Example 2 (Blowing up and aligning simple networks). Consider the weighted networks X and Y shown in the top row of Figure 1. We represent these as measure networks by taking ω_X and ω_Y to be weighted adjacency matrices. Concretely, let $X = \{x\}$, $\omega_X = (1)$, and $\mu_X(x) = 1$. Also let $Y = \{y_1, y_2\}$, $\omega_Y = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, and $\mu_Y(y_1) = \mu_Y(y_2) = 1/2$. Since X is a one-node network, the unique coupling $\mu \in \mathcal{C}(\mu_X, \mu_Y)$ is given by $\mu(x, y_1) = \mu(x, y_2) = 0.5$. To convert μ to a transport map, X is blown-up to $\hat{X} = \{x_1, x_2\}$, with $\omega_{\hat{X}} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ and $\mu_{\hat{X}}(x_1) = \mu_{\hat{X}}(x_2) = 1/2$, whence $\hat{\mu} = \begin{pmatrix} 0.5 & 0 \\ 0 & 0.5 \end{pmatrix}$. Intuitively, the average of \hat{X} and Y should be the network $Z = \{z_1, z_2\}$, $\omega_Z = \begin{pmatrix} 0.5 & 1 \\ 1 & 0.5 \end{pmatrix}$, and $\mu_Z(z_1) = \mu_Z(z_2) = 0.5$. This intuition will be formalized below.

2.3. Computing GW distance

It was implicitly observed in Section 2.1 that for finite measure networks X of size n and Y of size m , the squared distortion of a coupling matrix $C \in \mathcal{C}(\mu_X, \mu_Y) \subset \mathbb{R}^{n \times m}$ is given by

$$\text{dis}(C)^2 = \langle \mathcal{L}(X, Y) \otimes C, C \rangle.$$

For fixed X and Y , let A_{XY} denote the linear map from $\mathbb{R}^{n \times m}$ to itself given by $A_{XY}C := \mathcal{L}(X, Y) \otimes C$. The GW

optimization problem seeks a minimizer of the map

$$C \mapsto \langle A_{XY}C, C \rangle$$

over the convex polytope $\mathcal{C}(\mu_X, \mu_Y)$ and is therefore an instance of a *quadratic programming* problem.

Following [23], we approximate GW distance by finding local minimizers for the GW optimization problem via projected gradient descent. Since we are allowing asymmetric weight functions, the linear map A_{XY} may be asymmetric. This distinguishing feature from the setting of [23] must be accounted for when computing the gradient.

Proposition 3. *The gradient of $C \mapsto \langle A_{XY}C, C \rangle$ is given by*

$$(A_{XY} + A_{XY}^*)C,$$

with A_{XY}^* denoting the adjoint of A_{XY} .

2.4. Fréchet means

Given a collection of networks $S = \{X_1, X_2, \dots, X_n\}$, a *Fréchet mean* of S is a minimizer of the functional

$$F_S(Z) := \frac{1}{n} \sum_{i=1}^n d_{\mathcal{N}}(X_i, Z)^2.$$

In [23], the approach for calculating a Fréchet mean was as follows: (1) fix a cardinality N for the target space Z , (2) minimize over choices of ω_Z , and (3) optimize over couplings $C_i \in \mathcal{C}(\mu_{X_i}, \mu_Z)$. The last two steps are repeated until convergence.

In contrast, the scheme we present in this paper follows ideas of Pennec on averaging in (finite-dimensional) Riemannian manifolds [22] coupled with the work of Sturm on developing the Riemannian structure of generalizations of metric measure spaces. Informally, the idea is as follows: start with a “seed” network X , use log maps to lift geodesics $X \rightarrow X_i$ to vectors in the tangent space at X , average the vectors, use an exponential map to map down to \mathcal{N} , and iterate this procedure until convergence. Theoretically, we justify this procedure by showing that it agrees with the downward gradient flow of the Fréchet functional.

3. Metric geometry of $[\mathcal{N}]$

Given a network (X, ω_X, μ_X) , the only requirement on ω_X is that it needs to be square integrable; i.e., we need $\omega_X \in L^2(X^2, \mu_X^{\otimes 2})$. We will show that this flexibility allows us to define structures on $[\mathcal{N}]$ analogous to those of a Riemannian manifold, such as geodesics, tangent spaces and exponential maps. These structures are defined using language from the theory of analysis on metric spaces. In our setting, they can be defined in a surprisingly concrete way, allowing us to sidestep the need to invoke any deep concepts or results—see [5, 2] for general introductions to the theory.

3.1. Geodesics

In the metric geometry sense, a *geodesic* from $[X]$ to $[Y]$ in $[\mathcal{N}]$ is a continuous map γ from a closed interval $[S, T]$ into $[\mathcal{N}]$ satisfying the property

$$d_{\mathcal{N}}(\gamma(s), \gamma(t)) = \frac{|t-s|}{T-S} \cdot d_{\mathcal{N}}([X], [Y]) \quad (2)$$

for all $s, t \in [S, T]$. The geodesic is *unit speed* if $T - S = d_{\mathcal{N}}([X], [Y])$. We can of course assume without loss of generality that our domain interval is always of the form $[0, T]$. We say that a geodesic $\gamma : [0, T] \rightarrow [\mathcal{N}]$ *emanates from* $[X]$ if $\gamma(0) = [X]$.

It follows from work in [26] that geodesics can always be constructed in $[\mathcal{N}]$ (although they need not be unique). Let $X, Y \in \mathcal{N}$, and let μ be an optimal coupling (cf. Lemma 1). For each $t \in [0, 1]$, define

$$\gamma(t) := [X \times Y, \Omega_t, \mu], \quad (3)$$

where

$$\Omega_t((x, y), (x', y')) := (1-t)\omega_X(x, x') + t\omega_Y(y, y').$$

It is easy to see that $\gamma(0) = [X]$ and $\gamma(1) = [Y]$. A relatively straightforward computation then shows that γ satisfies (2) (cf. [26, Theorem 3.1]). Note that the underlying set of this geodesic is always $X \times Y$. In particular, $\gamma(0)$ is the triple $[X \times Y, \omega_X, \mu]$ where ω_X is defined (by abuse of notation) on $X \times Y \times X \times Y$ as $\omega_X((x, y), (x', y')) := \omega_X(x, x')$. Here the ω_X on the right hand side is the original function defined on $X \times X$.

For networks X and Y of sizes m and n , respectively, a naive implementation of the geodesic described above is represented by a measure network of size $m \cdot n$. We later consider iterative algorithms where this size blowup would quickly become intractable. Instead, we offer the following minimal size representation of a geodesic. We compute an optimal coupling of X and Y , then blow up and align the networks as in Definitions 1 and 2; an example for networks coming from simple graphs is shown in Figure 2. The geodesic described above is then represented (up to weak isomorphism) by interpolating the blown up and aligned weight matrices for X and Y ; such a geodesic for graph networks is shown in Figure 3. We have observed empirically that such a minimal size geodesic is typically represented as a path of matrices of size proportional to $m+n$, rather than the size $m \cdot n$ naive representation. Experimental evidence of this size reduction is provided in the supplementary materials, Section C.

3.2. Tangent space

For a point $[X] \in [\mathcal{N}]$, we define the *tangent space* to be:

$$T_{[X]} := \bigcup_{Z \in [X]} L^2(Z^2, \mu_Z^{\otimes 2}) / \sim,$$




Figure 2. **Left column:** Two graphs to be matched. Each graph is a measure network with all edge weights equal to one and uniform node weights. The lighter arrows depict an optimal coupling between the graphs; the mass from each node from the first graph is distributed evenly to two nodes in the second graph as indicated. **Middle column:** Graphs are represented as measure networks by their adjacency matrices of size 3×3 and 6×6 , respectively. **Right column:** Matrix representations of the measure networks after they have been blown up and aligned according to the optimal coupling. Resulting matrices are both size 6×6 . The matrix for the first graph is doubled in size; at a graph level, copies of each node are created. The matrix for the second graph has remained the same size but has been permuted; this is just a relabelling of the nodes based on how they are matched with nodes of the first graph.

where \sim is defined as follows. For (Y, ω_Y, μ_Y) and (Z, ω_Z, μ_Z) in $[X]$ and functions $f \in L^2(Y^2, \mu_Y^{\otimes 2})$ and $g \in L^2(Z^2, \mu_Z^{\otimes 2})$, we declare $f \sim g$ if and only if there exists a coupling μ of μ_Y and μ_Z such that

$$\omega_Y(y, y') = \omega_Z(z, z') \quad \text{and} \quad f(y, y') = g(z, z')$$

for $\mu^{\otimes 2}$ -a.e. $((y, z), (y', z')) \in (Y \times Z)^2$. Elements of $T_{[X]}$ are called *tangent vectors to* $[X]$ and are denoted $[f]$, where f is an L^2 function defined on some representative of $[X]$.

We define an inner product $\langle \cdot, \cdot \rangle_{[X]}$ on each tangent space $T_{[X]}$ as follows: for $f \in L^2(Y^2, \mu_Y^{\otimes 2})$ and $g \in L^2((Y')^2, \mu_{Y'}^{\otimes 2})$ with $Y, Y' \in [X]$,

$$\langle [f], [g] \rangle_{[X]} := \langle (\pi_Y)^* f, (\pi_{Y'})^* g \rangle_{L^2(Z^2, \mu_Z^{\otimes 2})},$$

where Z is any measure network realizing the ‘tripod’ in the definition of weak isomorphism between Y and Y' . One can check that this value does not depend on any of the choices made and therefore gives a well-defined inner product. The norm induced by this inner product is denoted $\| \cdot \|_{[X]}$. One can check that it reduces to the formula

$$\|[f]\|_{[X]} := \|f\|_{L^2(Y^2, \mu_Y^{\otimes 2})},$$

where $f \in L^2(Y^2, \mu_Y^{\otimes 2})$ and $Y \in [X]$.

Remark 4. The tangent space $T_{[X]}$ is not a bona fide vector space, but is a vector space quotiented out by a ‘symmetry group’ consisting of optimal self-couplings of a certain




Figure 3. Geodesic between the graphs of Figure 2. The adjacency matrices are blown up and realigned. The blown up graph second from the left is drawn with split nodes superimposed. The geodesic interpolates matched edge weights. At the graph level, the geodesic is illustrated with node sizes corresponding to node weights and edge thickness corresponding to edge weights.

minimal representative of X [26, Section 6]. When $[X]$ has a representative with only the trivial symmetry, i.e. the diagonal coupling, the tangent space is a Hilbert space. In particular, this phenomenon endows $[\mathcal{N}]$ with the structure of a Riemannian orbifold.

3.3. Exponential map

We define the *exponential map* at $[X]$,

$$\exp_{[X]} : T_{[X]} \rightarrow [\mathcal{N}],$$

as follows. For $f \in L^2(Z, \mu_Z^{\otimes 2})$, with $Z \in [X]$, let

$$\exp_{[X]}([f]) := [Z, \omega_Z + f, \mu_Z].$$

After unwrapping the various notions of equivalence involved, one is able to show that this map is well-defined. This map is analogous to the exponential map in a Riemannian manifold. We demonstrate this concretely in the finite setting with the next proposition.

Proposition 5. *Let X be an finite measure network. There exists $\epsilon_{[X]} > 0$ such that for any tangent vector represented by $f \in L^2(Z^2, \mu_Z^{\otimes 2})$ with $Z \in [X]$ satisfying $|f(z, z')| < \epsilon_{[X]}$ for all $(z, z') \in Z \times Z$, $\exp_{[X]}([f])$ is the endpoint of a geodesic emanating from $[X]$.*

3.4. Log map

We wish to show that $\exp_{[X]}$ has a local inverse, called the *log map* at $[X]$. Let Y be a finite measure network and let μ be an optimal coupling of X and Y . Define the *log map* with respect to μ as follows. Use μ to expand and align the measure networks to

$$\hat{X} = (\hat{X}, \omega_{\hat{X}}, \mu_{\hat{X}}) \quad \text{and} \quad \hat{Y} = (\hat{X}, \omega_{\hat{Y}}, \mu_{\hat{X}}) \quad (4)$$

so that the identity map on the set \hat{X} induces an optimal coupling of \hat{X} with \hat{Y} . We then define

$$\log_{[X]}^{\mu}([Y]) := [\omega_{\hat{Y}} - \omega_{\hat{X}}]. \quad (5)$$

It immediately follows that

$$\exp_{[X]} \left(\log_{[X]}^{\mu}([Y]) \right) = [Y].$$

This provides a surjectivity result for the exponential map. On the other hand, the following lemma provides an injectivity result. Its proof is similar to that of Proposition 5.

Lemma 6. *Let X be a finite measure network. The exponential map $\exp_{[X]}$ is injective on the set of $[f]$ with $f \in L^2(Z^2, \mu_Z^{\otimes 2})$ such that Z is finite and f satisfies $|f(z, z')| < \epsilon_{[X]}/2$ for all $z, z' \in Z$.*

We now define the *log map* at $[X]$, $\log_{[X]}$, to be the local inverse of $\exp_{[X]}$ on finite measure networks near $[X]$. For a finite measure network Y , we define $\log_{[X]}([Y]) = \log_{[X]}^{\mu}([Y])$ as in (5), where μ is any optimal coupling of X with Y . The lemma then provides a certificate to check that the image of the log map did not depend on a choice of optimal coupling.

3.5. Gradients

Let $F : [\mathcal{N}] \rightarrow \mathbb{R}$ be a functional, let $[X] \in [\mathcal{N}]$ and let $[f] \in T_{[X]}$. Define the *directional derivative* of F at $[X]$ in the direction $[f]$ to be the limit

$$D_{[f]}F([X]) := \lim_{t \rightarrow 0^+} \frac{1}{t} \left(F \left(\exp_{[X]}([t \cdot f]) \right) - F([X]) \right),$$

provided it exists. We say that F is *differentiable* at $[X]$ if all directional derivatives exist.

For a differentiable functional F , a *gradient* of F at $[X]$ is a tangent vector $\nabla F([X])$ satisfying

$$D_{[f]}F([X]) = \langle [f], \nabla F([X]) \rangle_{[X]}$$

for all $[f] \in T_{[X]}$.

The next lemma follows from [26, Lemma 6.24].

Lemma 7. *Let F be a differentiable functional. If the gradient of F at $[X]$ exists, then it is unique and satisfies*

$$\|\nabla F([X])\|_{[X]} = \sup \{ D_{[f]}F([X]) \mid \| [f] \|_{[X]} = 1 \}.$$




Figure 4. Examples of geodesics between simple graphs. See Section 4.1 for explanations.

We are now able to derive an explicit expression for the gradient of the Fréchet functional for finite networks.

Proposition 8. *Let $S = \{Y_1, Y_2, \dots, Y_n\}$ be a collection of finite networks, and let X be another finite network. Suppose each Y_i has been aligned to X , so that each of $\{X, Y_1, \dots, Y_n\}$ has m nodes. Then the gradient of the Fréchet functional F_S at $[X]$ is represented by the $m \times m$ matrix $\nabla F_S(X)$ defined by*

$$(\nabla F_S(X))_{ij} = 2 \left(\omega_X(x_i, x_j) - \frac{1}{n} \sum_{k=1}^n \omega_{Y_k}(y_i, y_j) \right).$$

Remark 9. The preceding proposition gives us a meaningful description of a Fréchet mean. Specifically, let $S = \{Y_1, \dots, Y_n\}$ be a collection of finite networks, and let X be such that $\nabla F_S(X) = 0$. Suppose also that X is aligned to each Y_i . Then X has the property that $\omega_X(x_i, x_j) = \frac{1}{n} \sum_{k=1}^n \omega_{Y_k}(y_i, y_j)$ for each $x_i, x_j \in X$. In other words, ω_X comprises arithmetic means of entries in the ω_{Y_k} .

4. Experiments

We now provide details of our computational experiments. Algorithms and numerous empirical remarks are provided in the supplementary materials. Python implementations are available on GitHub [9]. Our code makes heavy use of the Python Optimal Transport Library [12].

4.1. Geodesic Examples

Figure 4 shows several examples of geodesics between simple graphs. Each graph X is a measure network with ω_X the graph adjacency matrix. Except for the example in the first row, each μ_X is a uniform node measure.



Figure 5. Samples from five shape classes in the classification experiment.

The first row in the figure shows a geodesic between graphs with the same edge weight structure as the graphs in Figure 3, but with different node weights indicated by node sizes. Observe that the difference in node weights changes the geodesic path drastically. The second row in the figure shows a geodesic between a graph with two disconnected nodes, each with a self-loop (not shown) and a connected graph with large clusters. The geodesics in the third and fourth rows of the figure are each between graphs from the “Letter Graphs” graph classification benchmark dataset [24]; the first between letters in the same class and the second between letters in different classes. Each geodesic is displayed with lower opacity on some nodes and edges—these nodes with weight less than a user-defined threshold (50% (respectively, 40%) of maximum node weight in the first (respectively, second) example) and edges with at least one endpoint meeting this criteria. This technique allows us to understand common graph features at multiple resolutions.

4.2. Shape Classification

As a proof-of-concept for incorporating this framework into machine learning pipelines, we present a simple shape classification experiment. The data consists of 20 object classes with 20 samples from each class from the well known MPEG-7 computer vision database (see Figure 5). Each shape consists of 100 planar points. The input data for the experiment consists of pairwise distance matrices for each shape, yielding 400 matrices of size 100×100 . The ordering of the points was randomized when constructing the distance matrices. Weights on the nodes are uniform.

We consider three methodologies for classifying the shapes. In each experiment the same 80% of the shapes were used as a training set. In a completely naive approach, a support vector machine was trained on the permuted distance matrices. With this method, the classification rate on the testing set was 19%. In the second approach, one of the permuted distance training matrices X was fixed and all other training matrices were aligned to X . The matrices were then “centered” on X , which can be understood as pulling them back to tangent vectors in $T_{[X]}$ via the log map (or, rather, the coupling-dependent log map $\log_{[X]}^\mu([Y])$ here). An SVM was then trained on these tangent vectors. Classification was tested by aligning and centering the test matrices with X , yielding a classification rate of 84%. In the final method, the Fréchet mean \bar{X} of all samples in the



Figure 6. Apple dataset. Each shape is a measure network X containing 100 points, with μ_X uniform measure and ω_X the pairwise (Euclidean) distance matrix between the points.



Figure 7. First three principal directions of variance for tangent PCA on the apple dataset.

training set was computed. Then all training matrices were aligned to \bar{X} and pulled back to tangent vectors, where an SVM was trained. Classification was once again performed by aligning and centering test matrices with \bar{X} , where the classification rate was improved to 94%. This approach illustrates a template for vectorization of network data.

4.3. Tangent PCA on planar shapes

The Riemannian framework allows us to do other machine learning computations by pulling networks back to a tangent space; for example, we now present results of a tangent PCA experiment. Figure 6 illustrates the *apple dataset* that we used. Each shape is represented as a measure network as in the previous subsection. To perform tangent PCA, we first computed a Fréchet mean for these 20 shapes. Next we used log maps based at the chosen Fréchet mean to pull back the 20 shapes to vectors in the tangent space. Here we performed PCA as usual.

The first three principal directions explained 85% of the variance in the data, and they are visualized in Figure 7 via MDS embeddings. The first direction captures variance in the size of the apple; the second captures surface irregularities and the size and shape of the leaves, and the third captures the presence of a “bite” on the apple.

4.4. Compressing an Asymmetric SBM Network

To illustrate our constructions on asymmetric networks, we generated a 100×100 asymmetric stochastic block model (SBM) network Y following the model provided in [8]. Here Y consisted of five blocks B_1, \dots, B_5 of 20 nodes each. For $y \in B_i$ and $y' \in B_j$, we sampled $\omega_Y(y, y') \sim N(\mu_{ij}, 5)$, where $\mu_{ij} \in \{0, 25, 50, 75, 100\}$. Negative values were allowed. See Figure 8 for an illustration of Y . Note that Y is intuitively represented by a 5×5 “ground-truth” matrix.




Figure 8. **Left:** 100×100 SBM Y with entries drawn from $N(\mu, \sigma^2)$, where $\sigma^2 := 5$. The five colors correspond to $\mu = 0$ (blue), 25, 50, 75, 100(yellow). **Middle:** 5×5 compressed average of Y and the 5×5 all-zeros matrix X . Both μ_X, μ_Y were taken to be uniform. Colors range in $\{0, 12.5, 25, 37.5, 50\} \pm 0.1$. **Right:** Permuted form of middle figure.

We averaged Y with a 5×5 all-zeros matrix X using the network compression approach given in the supplementary materials to see if our method would recover the ground truth matrix. The output of our method is shown in the middle panel of Figure 8—up to a permutation (shown in the right panel), this accurately recovered the matrix of μ_{ij} values.

5. Discussion

In this paper, we followed the seminal work of Sturm [26] on Riemannian structures induced by GW distances and produced a Riemannian framework for performing data analysis on collections of arbitrary matrices. There are many applications in data science which can be reframed using this formalism, such as network clustering and sketching and future work will focus on making these formulations precise.

There are several open challenges left to be explored from both theoretical and practical perspectives. On the theoretical side, one would like to obtain estimates on injectivity radii for measure networks with special properties (this amounts to replacing the L^∞ bounds in Proposition 5 with L^2 bounds). It would also be interesting to determine conditions where the Fréchet gradient flow is guaranteed to converge—in our applications, this either happened naturally or was enforced by a compression step. On the practical side, one would like to improve the scalability of our algorithms by incorporating entropic regularization [10] and the more sophisticated GW algorithm of [29]. Several steps in our framework relied on the empirical observation of sparsity in optimal couplings, so incorporating entropic regularization will bring its own collection of theoretical challenges.

Acknowledgments

We would like to thank Facundo Mémoli for many useful discussions about Gromov-Wasserstein distance and for introducing us to the work of Sturm on the topic. We would also like to thank Justin Solomon for a helpful conversation about the matrix formulation of the GW problem for asymmetric networks. Finally, we thank the anonymous referees for their thoughtful comments.

References

- [1] David Alvarez-Melis and Tommi Jaakkola. Gromov-Wasserstein alignment of word embedding spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1881–1890, 2018. [2](#)
- [2] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008. [4](#)
- [3] David Bao, S-S Chern, and Zhongmin Shen. *An introduction to Riemann-Finsler geometry*, volume 200. Springer Science & Business Media, 2012. [1](#)
- [4] Charlotte Bunne, David Alvarez-Melis, Andreas Krause, and Stefanie Jegelka. Learning generative models across incomparable spaces. In *International Conference on Machine Learning*, pages 851–861, 2019. [2](#)
- [5] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A course in metric geometry*, volume 33. American Mathematical Soc., 2001. [4](#)
- [6] Xin Chen and Jiming Peng. New analysis on sparse solutions to random standard quadratic optimization problems and extensions. *Mathematics of Operations Research*, 40(3):725–738, 2015. [11](#)
- [7] Xin Chen, Jiming Peng, and Shuzhong Zhang. Sparse solutions to random standard quadratic optimization problems. *Mathematical Programming*, 141(1-2):273–293, 2013. [11](#)
- [8] Samir Chowdhury and Facundo Mémoli. The gromov-wasserstein distance between networks and stable network invariants. *Information and Inference: A Journal of the IMA*, 8(4):757–787, 2019. [3](#), [8](#)
- [9] Samir Chowdhury and Tom Needham. Gromov-Wasserstein statistics Github repository. <https://github.com/trneedham/gromov-wasserstein-statistics>. [7](#)
- [10] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013. [8](#)
- [11] Danielle Ezuz, Justin Solomon, Vladimir G Kim, and Mirela Ben-Chen. GWCNN: A metric alignment layer for deep shape analysis. In *Computer Graphics Forum*, volume 36, pages 49–57. Wiley Online Library, 2017. [2](#)
- [12] Rémi Flamary and Nicolas Courty. POT: Python Optimal Transport library, 2017. [7](#), [14](#)
- [13] Xiaoyang Guo, Anuj Srivastava, and Sudeep Sarkar. A quotient space formulation for statistical analysis of graphical data. *arXiv preprint arXiv:1909.12907*, 2019. [2](#)
- [14] Reigo Hendrikson. Using Gromov-Wasserstein distance to explore sets of networks. *University of Tartu, Master Thesis*, 2016. [2](#)
- [15] Brijnesh J Jain and Klaus Obermayer. Structure spaces. *Journal of Machine Learning Research*, 10(Nov):2667–2714, 2009. [2](#)
- [16] Brijnesh J Jain and Klaus Obermayer. Learning in Riemannian orbifolds. *arXiv preprint arXiv:1204.4294*, 2012. [2](#)
- [17] Facundo Memoli. On the use of Gromov-Hausdorff Distances for Shape Comparison. In M. Botsch, R. Pajarola, B. Chen, and M. Zwicker, editors, *Eurographics Symposium on Point-Based Graphics*. The Eurographics Association, 2007. [2](#)
- [18] Facundo Mémoli. Gromov-Wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics*, pages 1–71, 2011. 10.1007/s10208-011-9093-5. [2](#)
- [19] Facundo Mémoli and Tom Needham. Gromov-Monge quasi-metrics and distance distributions. *arXiv preprint arXiv:1810.09646*, 2018. [2](#)
- [20] Mark Newman. *Networks: an introduction*. Oxford university press, 2010. [1](#)
- [21] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006. [12](#)
- [22] Xavier Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127, 2006. [1](#), [2](#), [4](#)
- [23] Gabriel Peyré, Marco Cuturi, and Justin Solomon. Gromov-Wasserstein averaging of kernel and distance matrices. In *International Conference on Machine Learning*, pages 2664–2672, 2016. [1](#), [2](#), [3](#), [4](#), [10](#), [13](#), [14](#)
- [24] Kaspar Riesen and Horst Bunke. IAM graph database repository for graph based pattern recognition and machine learning. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*, pages 287–297. Springer, 2008. [7](#)
- [25] Justin Solomon, Gabriel Peyré, Vladimir G Kim, and Suvrit Sra. Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)*, 35(4):72, 2016. [3](#)
- [26] Karl-Theodor Sturm. The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces. *arXiv preprint arXiv:1208.0434*, 2012. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)
- [27] Vayer Titouan, Nicolas Courty, Romain Tavenard, and Rémi Flamary. Optimal transport for structured data with application on graphs. In *International Conference on Machine Learning*, pages 6275–6284, 2019. [2](#)
- [28] Vayer Titouan, Rémi Flamary, Nicolas Courty, Romain Tavenard, and Laetitia Chapel. Sliced Gromov-Wasserstein. In *Advances in Neural Information Processing Systems*, pages 14726–14736, 2019. [2](#)
- [29] Hongteng Xu, Dixin Luo, and Lawrence Carin. Scalable Gromov-Wasserstein learning for graph partitioning and matching. In *Advances in Neural Information Processing Systems*, pages 3046–3056, 2019. [2](#), [8](#)
- [30] Hongteng Xu, Dixin Luo, Hongyuan Zha, and Lawrence Carin. Gromov-Wasserstein learning for graph matching and node embedding. In *International Conference on Machine Learning*, pages 6932–6941, 2019. [2](#)