

Gromov-Hausdorff Stable Signatures for Shapes using Persistence

Frédéric Chazal¹ and David Cohen-Steiner² and Leonidas J. Guibas³ and Facundo Mémoli⁴ and Steve Y. Oudot¹

¹INRIA Saclay ²INRIA Sophia ³Stanford University, Dept. Computer Science ⁴Stanford University, Dept. Mathematics

Abstract

We introduce a family of signatures for finite metric spaces, possibly endowed with real valued functions, based on the persistence diagrams of suitable filtrations built on top of these spaces. We prove the stability of our signatures under Gromov-Hausdorff perturbations of the spaces. We also extend these results to metric spaces equipped with measures. Our signatures are well-suited for the study of unstructured point cloud data, which we illustrate through an application in shape classification.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modelling—

1. Introduction

In recent years, great advances have been made in the fields of data acquisition and shape modelling, and huge collections of digital models have been obtained. With the goal of organizing these collections, it is very important to be able to define and compute meaningful notions of *similarity* between objects that exhibit invariance to different deformations or poses of the objects represented by the data. Problems of this nature arise in areas such as molecular biology, metagenomics, face recognition, matching of articulated objects, graph matching, and pattern recognition in general.

In many practical applications, data are endowed with a notion of distance between their points, which turns them into metric spaces. The information contained in such data sets typically takes the form of *metric invariants*, which it is of interest to capture and characterize. For instance, when the data have been obtained experimentally, the analysis of these invariants can provide insights about the nature of the underlying phenomenological science.

Relevant work. In the context of shape analysis, one typically wishes to be able to discriminate shapes under different notions of invariance. Many approaches have been proposed for the problem of (pose invariant) shape classification and recognition, including the *size theory* of Frosini and collaborators [Fro90, CFL06a, CFL06b], the work of Hilaga et al. [HSKK01], the *shape contexts* [BMP02], the

integral invariants of [MCH*06], the eccentricity functions of [HK03], the *shape distributions* of [OFCD02], the *canonical forms* of [EK03], and the *shape DNA* and *global point signatures* based spectral methods in [RWP05] and [Rus07], respectively. The common underlying idea revolves around the computation and comparison of certain metric invariants, or *signatures*, so as to ascertain whether two given data sets represent in fact the same object, up to a certain notion of invariance.

The question of proving that a given family of signatures is indeed able to signal proximity or similarity of shapes in a reasonable way has been hardly addressed. In particular, the degree to which two shapes with similar signatures are forced to be similar is in general not well understood. Conversely, one can ask the more basic question of whether the similarity between two shapes forces their signatures to be similar. These questions cannot be completely well formulated until one agrees on: (1) a notion of *equality* between shapes, and (2) a notion of *(dis)similarity* between shapes.

By regarding sampled shapes as finite metric spaces, one can use the Gromov-Hausdorff distance [Gro99] as a measure of dissimilarity between shapes [MS04, MS05]. By endowing the point clouds with different metrics, one obtains a great deal of flexibility in the different degrees of invariance that one can encode in the measure of dissimilarity. For instance, using the Euclidean metric in \mathbb{R}^d makes the Gromov-Hausdorff distance invariant under ambient rigid isome-

tries [Mém08]. In contrast, using intrinsic metrics within the shapes makes the Gromov-Hausdorff distance blind to intrinsic isometries, such as when a same animal is represented in different poses (Figure 1). Bronstein *et al.* [BBK09] suggested to combine intrinsic and extrinsic metrics to increase the discriminating power of shape signatures. We illustrate this idea in Figure 6, where the first two shapes (circle and helicoidal curve on a torus) have same length and are therefore isometric embeddings of a same object into \mathbb{R}^3 . As a result, signatures obtained from geodesic distances are identical whereas signatures based on Euclidean distances differ.

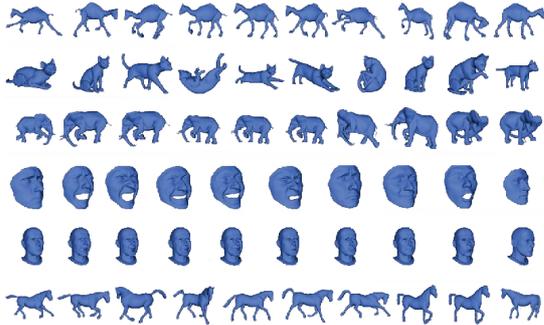


Figure 1: The database of shapes used in our experiments.

Once the finite metric space representation of shapes has been adopted, the question becomes whether a given family of signatures is *stable* under perturbations of the shape in the Gromov-Hausdorff distance. A somewhat different version of the problem is when the spaces are endowed with measures, a case in which perturbations are understood in the sense of the (l^∞) Gromov-Wasserstein distance [Mém07]. These questions about stability of invariants are addressed in [Mém07] for a limited set of signatures [HK03, MCH*06, BMP02, OFCD02].

Our contributions. Using tools from *persistent homology* [ELZ00, ZC05], we construct a novel class of Gromov-Hausdorff stable metric invariants. Our approach is in the spirit of [CZCG05] and [dFL06, dFL05] in that we try to devise certain topologically motivated signatures that naturally combine the classificatory power of topology with the flexibility of metrics.

We enrich the pool of shape signatures available in the literature in two ways (see Section 3):

- First, we extend zero-dimensional connectivity-based signatures such as the *denograms* of [CM09] or the invariants computed by size theory [dFL06, dFL05] to higher-dimensional homology classes, and we prove their stability under small Gromov-Hausdorff perturbations. This greatly increases the discriminative power of the approach, as illustrated in Figure 6, where all shapes have been ϵ -sampled for a same value of ϵ and therefore all zero-dimensional diagrams are the same (rows 2 and 3).

- Second, we extend the approach to the case of metric spaces endowed with real-valued functions. This latter step enables the design of new signatures for finite metric spaces by considering canonically-defined functions on these spaces, such as eccentricities (Section 4). As a result, more shapes can be discriminated using our approach. For instance, in Figure 6, the sphere and ellipsoid cannot be discriminated using standard Rips filtrations (see Definition 2.7) because their diagrams are the same up to rescaling and small noise (rows 2 and 3). In contrast, using a fraction of the eccentricity as function modifies the shapes of the 1- and 2-dimensional diagrams of the ellipsoid significantly, while the ones for the sphere hardly change due to the fact that the eccentricity is constant (row 4).

Our signatures are also adapted to deal with *measure metric spaces* [Gro99, Chapter 3.1/2], where additional information is provided in the form of a collection of weights, one per point of the shape. In this case, stability is guaranteed under small perturbations in the Gromov-Wasserstein distance (Section 5).

Since the computation of the Gromov-Hausdorff distance leads to hard combinatorial problems akin to the *quadratic assignment problem* (which is NP hard), the construction of easily computable Gromov-Hausdorff stable invariants is very important. We show in Section 6 that the computation of our invariants yields simple *bottleneck assignment problems*, which can be solved in polynomial time. In contrast with other approaches that generally involve some delicate numerical algorithms and/or require the availability of a mesh of the underlying object [MS04, BBK06], the algorithm for computing and comparing our signatures is entirely combinatorial and can be applied on the sole input of a distance matrix.

Finally, our invariants prove to be useful for discriminating between different classes of shapes in a classification task performed on the database of Figure 1. The results are presented in Section 7.

2. Mathematical background

2.1. Metric spaces and Gromov-Hausdorff distance(s)

Recall that a metric space is a pair (X, d_X) where X is a set and $d_X : X \times X \rightarrow \mathbb{R}$ is a non-negative map such that for any $x, y, z \in X$, $d_X(x, y) = 0$ if and only if $x = y$, $d_X(x, y) = d_X(y, x)$ and $d_X(x, z) \leq d_X(x, y) + d_X(y, z)$.

In the sequel \mathcal{X} denotes the collection of all compact metric spaces. Two compact spaces (X, d_X) and (Y, d_Y) are said to be *isometric* if there exists a bijection $\phi : X \rightarrow Y$ that preserves distances, namely: $\forall x, x' \in X$, $d_Y(\phi(x), \phi(x')) = d_X(x, x')$. Such a map is called an *isometry*. The set of isometry classes of compact metric spaces can be endowed with a distance called the Gromov-Hausdorff distance [BBI01, Thm. 7.3.30], whose definition uses the following notion of correspondence between sets [BBI01, Def. 7.3.17]:

Definition 2.1 A *correspondence* between two sets X and Y is a subset $C \subset X \times Y$ such that: $\forall x \in X, \exists y \in Y$ s.t. $(x, y) \in C$, and $\forall y \in Y, \exists x \in X$ s.t. $(x, y) \in C$. The set of all correspondences between X and Y is denoted $\mathcal{C}(X, Y)$.

Definition 2.2 The *Gromov-Hausdorff distance* between compact metric spaces $(X, d_X), (Y, d_Y)$ is:

$$d_{\text{GH}}((X, d_X), (Y, d_Y)) = \frac{1}{2} \inf_{C \in \mathcal{C}(X, Y)} \|\Gamma_{X, Y}\|_{l^\infty(C \times C)}, \quad (1)$$

where $\Gamma_{X, Y} : X \times Y \times X \times Y \rightarrow \mathbb{R}^+$ is defined by $(x, y, x', y') \mapsto |d_X(x, x') - d_Y(y, y')|$ and the notation $\|\Gamma_{X, Y}\|_{l^\infty(C \times C)}$ stands for $\sup_{(x, y), (x', y') \in C} \Gamma_{X, Y}(x, y, x', y')$.

Another way of defining the Gromov-Hausdorff distance is *via* common isometric embeddings of the two spaces. Specifically, given two compact subsets X, Y of a same metric space (Z, d_Z) , the *Hausdorff distance* between X and Y is the quantity:

$$d_{\text{H}}^Z(X, Y) = \max\{\max_{x \in X} \min_{y \in Y} d_Z(x, y), \max_{y \in Y} \min_{x \in X} d_Z(x, y)\}.$$

The Gromov-Hausdorff distance is then defined as follows:

Definition 2.3 The *Gromov-Hausdorff distance* between compact metric spaces $(X, d_X), (Y, d_Y)$ is:

$$d_{\text{GH}}((X, d_X), (Y, d_Y)) = \inf_{Z, \gamma_X, \gamma_Y} d_{\text{H}}^Z(\gamma_X(X), \gamma_Y(Y)), \quad (2)$$

where γ_X, γ_Y range over all the isometric embeddings of X, Y into some same metric space (Z, d_Z) .

Definitions 2.2 and 2.3 are known to be equivalent [BBI01, Thm. 7.3.25]. In particular, when the spaces are compact, the proof of equivalence shows that the infima in Eqs. (1) and (2) are in fact minima.

We now turn our focus to compact metric spaces endowed with continuous real-valued functions. Let \mathcal{X}_1 be the collection of all such spaces:

$$\mathcal{X}_1 = \{(X, d_X, f_X) \mid (X, d_X) \in \mathcal{X}, f_X : X \rightarrow \mathbb{R} \text{ continuous}\}.$$

We deem two spaces $X, Y \in \mathcal{X}_1$ *isomorphic* whenever there exists an isometry $\phi : X \rightarrow Y$ s.t. $f_X = f_Y \circ \phi$. We then extend the Gromov-Hausdorff distance as follows:

Definition 2.4 Given $X, Y \in \mathcal{X}_1$, $d_{\text{GH}}^1(X, Y) =$

$$\inf_{C \in \mathcal{C}(X, Y)} \max\left\{\frac{1}{2}\|\Gamma_{X, Y}\|_{l^\infty(C \times C)}, \|f_X - f_Y\|_{l^\infty(C)}\right\} \quad (3)$$

where $\|f_X - f_Y\|_{l^\infty(C)}$ stands for $\sup_{(x, y) \in C} |f_X(x) - f_Y(y)|$.

Note that if f_X and f_Y are equal constant functions, then one recovers d_{GH} .

As in the case of the standard Gromov-Hausdorff distance, the infimum in Eq. (3) is a minimum when the spaces are compact. Moreover, the proof of [BBI01, Thm. 7.3.25] shows that, for any correspondence $C \in \mathcal{C}(X, Y)$, the disjoint union $X \sqcup Y$ can be endowed with a metric d_C that coincides with d_X over $X \times X$ and with d_Y over $Y \times Y$, and such that

we have $d_C(x, y) \leq \frac{1}{2}\|\Gamma_{X, Y}\|_{l^\infty(C \times C)}$ and $|f_X(x) - f_Y(y)| \leq \|f_X - f_Y\|_{l^\infty(C)}$ whenever $(x, y) \in C$.

Theorem 2.5 d_{GH}^1 defines a metric on the set of isomorphism classes of \mathcal{X}_1 .

The proof of this result is a direct adaptation of the proof [BBI01, Thm. 7.3.30].

2.2. Filtrations, persistence diagrams and stability

Throughout the paper we consider simplicial homology with coefficients in a field, and we make a large use of topological persistence theory [ELZ00, ZC05]. Thorough introductions to these subjects can be found in [Mun84, EH, CCS07]. Recall that a simplicial complex K is a finite collection of simplices such that every face of a simplex of K is also in K and the intersection of any two simplices is either empty or a common face of each of them. A filtration \mathcal{K} of a simplicial complex K is a nested sequence of subcomplexes $\emptyset = K_{\alpha_0} \subseteq K_{\alpha_1} \subseteq \dots \subseteq K_{\alpha_m} = K$, where $\alpha_0 < \alpha_1 < \dots < \alpha_m$ is an ordered sequence of real numbers. The inclusion maps between the subcomplexes induce a directed system of vector spaces, called a *persistence module*, involving their k -dimensional homology groups:

$$H_k(K_{\alpha_0}) \xrightarrow{\phi_0^k} H_k(K_{\alpha_1}) \xrightarrow{\phi_1^k} \dots \xrightarrow{\phi_{m-1}^k} H_k(K_{\alpha_m}). \quad (4)$$

The structure of this persistence module can be encoded as a multi-set of points $D_k \mathcal{K}$, called the *k-th persistence diagram* of \mathcal{K} (see Figure 2). Intuitively, each point $p \in D_k \mathcal{K}$ encodes the lifespan of some k -dimensional homological feature appearing at time p_x and dying at time p_y in the filtration \mathcal{K} . More formally, $D_k \mathcal{K}$ is a multi-subset of the extended plane \mathbb{R}^2 , where $\mathbb{R} = \mathbb{R} \cup \{-\infty, +\infty\}$, contained in the union of the extended diagonal $\Delta = \{(x, x) : x \in \mathbb{R}\}$ and of the grid $\{\alpha_0, \dots, \alpha_m\} \times \{\alpha_0, \dots, \alpha_m, \alpha_\infty = +\infty\}$. The multiplicity of the points of Δ is set to $+\infty$, while the multiplicities of the (α_i, α_j) , $0 \leq i < j \leq +\infty$, are defined in terms of the ranks of the homomorphisms $\phi_i^j = \phi_{j-1}^j \circ \dots \circ \phi_i^{i+1}$. See [CCSG*09] for a precise definition.

The concept of persistence extends to the case of a filtration defined by the sublevel-sets of a continuous function $f : X \rightarrow \mathbb{R}$, under some *tameness* condition [CSEH05] stating that the family of inclusions $f^{-1}((-\infty, \alpha]) \subseteq f^{-1}((-\infty, \beta])$ induces at k -dimensional homology level a persistence module of the same (finite) type as in Eq. (4). The definition of k -th persistence diagram of such a function carries over. The use of persistence to devise signatures is motivated by the following stability property:

Theorem 2.6 (Stability [CSEH05]) Let X be a triangulable space and $f, g : X \rightarrow \mathbb{R}$ two tame continuous functions. Then, the bottleneck distance between the persistence diagrams of f and g in the extended plane \mathbb{R}^2 is at most $\|f - g\|_\infty$.

Recall that the bottleneck distance $d_{\text{B}}^\infty(A, B)$ between two

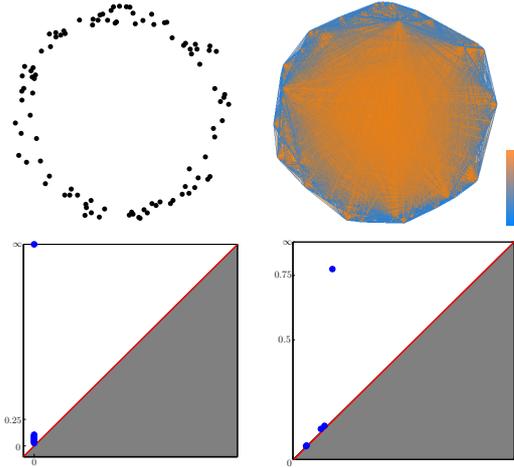


Figure 2: Top row, left: a set X of 100 points sampled uniformly at random in i.i.d. fashion from an offset of the unit circle in the plane. The point cloud is endowed with the ambient Euclidean metric d_E . Top row, right: the 2-skeleton of the Rips filtration $\mathcal{R}(X, d_E)$, containing 4,950 edges and 161,700 triangles. The simplices are colored according to their time of appearance in the filtration, from light blue to orange. Bottom row: the 0-dimensional (left) and 1-dimensional (right) persistence diagrams of $\mathcal{R}(X, d_E)$.

multi-sets in (\mathbb{R}^2, l^∞) is the quantity $\min_\gamma \max_{p \in A} \|p - \gamma(p)\|_\infty$, where γ ranges over all bijections from A to B .

Theorem 2.6 has been recently extended to wider classes of functions and spaces using a new notion of proximity between persistence modules called ε -interleaving [CCSG*09]. This extended stability result will be instrumental in the proof of Theorem 3.2.

From a practical point of view, the persistence diagrams of a given filtration can be efficiently computed using the so-called persistence algorithm [ELZ00, ZC05]. In this paper we use a particular family of complexes, called Rips complexes, to build filtrations on top of finite metric spaces (possibly endowed with functions), and we use their persistence diagrams as signatures for the spaces, as shown in Figure 2.

Definition 2.7 Given a finite metric space (X, d_X) and a parameter $\alpha > 0$, the *Rips complex* $R_\alpha(X, d_X)$ is the abstract simplicial complex of vertex set X , whose simplices correspond to the non-empty subsets of X of diameter less than 2α in the metric d_X [†]. The *Rips filtration* of (X, d_X) , noted $\mathcal{R}(X, d_X)$, is the nested family of Rips complexes obtained by varying parameter α from 0 to $+\infty$. When X is

[†] Note that we depart from the usual definition, where the Rips simplices correspond to subsets of X of diameter less than α . The reason for this choice is purely formal: it boils down to eliminating factors of 2 in the indices.

endowed with a real valued function f , we define a filtration $\mathcal{R}(X, d_X, f) = \{R_\alpha(X, d_X, f)\}_{\alpha > 0}$ where X_α denotes the set $f^{-1}((-\infty, \alpha]) \subseteq X$.

It has been shown in [CO08] that for well-chosen values of α , the Rips complex $R_\alpha(X, d_X)$ of a finite metric space (X, d_X) can be used to recover the homology of an unknown shape that is sufficiently densely sampled by X . In this paper, we adopt a more general point of view by studying the stability properties of the entire Rips filtration $\mathcal{R}(X, d_X)$.

2.3. Technical results

The following results will be used in the proofs but are not necessary to understand the general approach. We begin with a standard result (see e.g. example 3.5.3 and exercise 3.5.4 of [BBI01]) that will be key in the next section:

Lemma 2.8 Any finite metric space of cardinality n can be isometrically embedded into (\mathbb{R}^n, l^∞) .

In view of this lemma, Rips filtrations of finite subsets X of \mathbb{R}^n endowed with the l^∞ -distance are of particular interest for the rest of the paper. In particular they are closely related to the so-called Čech filtration of X within (\mathbb{R}^n, l^∞) , which is also defined as a nested family of abstract simplicial complexes of vertex set X . Specifically, for all $\alpha > 0$, the simplices of $C_\alpha(X, \mathbb{R}^n, l^\infty)$ correspond to the non-empty subsets of X that are centers of open l^∞ -balls of same radius α with a non-empty common intersection in \mathbb{R}^n .

Lemma 2.9 ([GM05]) For any finite set $X \subset \mathbb{R}^n$ and any $\alpha > 0$, one has $C_\alpha(X, \mathbb{R}^n, l^\infty) = R_\alpha(X, l^\infty)$.

Note that the Čech complex $C_\alpha(X, \mathbb{R}^n, l^\infty)$ is nothing but the nerve of the collection of open l^∞ -balls of same radius α about the points of X . The topological structures of unions of open balls in (\mathbb{R}^n, l^∞) are related to the ones of their nerves through the following extension of the Nerve Lemma which, more generally, relates *good* open covers to their nerves. Given a topological space S and a family $\mathcal{U} = \{U_x\}_{x \in X}$ of open subsets covering S , the family defines a *good* cover if for every finite subset Y of X the common intersection $\bigcap_{y \in Y} U_y$ is either empty or contractible. The *nerve* $\mathcal{N}\mathcal{U}$ is the abstract simplicial complex of vertex set X such that x_0, \dots, x_k form a simplex if and only if $U_{x_0} \cap \dots \cap U_{x_k} \neq \emptyset$.

Lemma 2.10 (Persistent Nerve [CO08]) Let $S \subseteq S'$ be two paracompact spaces, and let $\mathcal{U} = \{U_x\}_{x \in X}$ and $\mathcal{U}' = \{U'_{x'}\}_{x' \in X'}$ be good open covers of S and S' respectively, based on finite parameter sets $X \subseteq X'$ such that $U_x \subseteq U'_{x'}$ for all $x \in X$. Then, the homotopy equivalences $\mathcal{N}\mathcal{U} \rightarrow S$ and $\mathcal{N}\mathcal{U}' \rightarrow S'$ provided by the Nerve Theorem [Hat01, §4G] commute with the canonical inclusions $S \hookrightarrow S'$ and $\mathcal{N}\mathcal{U} \hookrightarrow \mathcal{N}\mathcal{U}'$ at homology level.

3. Persistence based lower bounds for Gromov-Hausdorff distance

Computing the Gromov-Hausdorff distance between two finite metric spaces appears to be practically intractable in most cases. To distinguish such spaces from a metric point of view, it is thus important to exhibit reasonable lower bounds whose computation is tractable. The results of this section show that the bottleneck distances between persistence diagrams of Rips filtrations of finite metric spaces are reasonable lower bounds.

Theorem 3.1 For any finite metric spaces (X, d_X) and (Y, d_Y) , for any $k \in \mathbb{N}$,

$$d_B^\infty(D_k \mathcal{R}(X, d_X), \mathcal{R}(Y, d_Y)) \leq d_{GH}((X, d_X), (Y, d_Y)).$$

Note that this lower bound on the Gromov-Hausdorff distance is tight. For instance, take for X a set of two points at distance 2 and for Y a set of two points at distance $2 + 2\varepsilon$. Then, (X, d_X) and (Y, d_Y) can be isometrically embedded into the real line, with X mapped to $\{0, 2\}$ and Y mapped to $\{-\varepsilon, 2 + \varepsilon\}$, which shows that their Gromov-Hausdorff distance is at most ε . Now, the 0-dimensional persistence diagram of the Rips filtration of (X, d_X) is made of two points, namely $(0, +\infty)$ and $(0, 1)$, while the 0-dimensional persistence diagram of the Rips filtration of (Y, d_Y) is made of $(0, +\infty)$ and $(0, 1 + \varepsilon)$, hence their bottleneck distance is ε .

Proof of Theorem 3.1. Let $\varepsilon = d_{GH}((X, d_X), (Y, d_Y))$. As mentioned after Definition 2.3, the infimum in Eq. (2) is in fact a minimum, which means that there is a metric space (Z, d_Z) and two isometric embeddings $\gamma_X : X \rightarrow Z$ and $\gamma_Y : Y \rightarrow Z$ such that the Hausdorff distance $d_H^Z(\gamma_X(X), \gamma_Y(Y))$ is ε . Consider now the subspace $\gamma_X(X) \cup \gamma_Y(Y) \subseteq Z$, endowed with the induced metric. Since this subspace is finite, it can be isometrically embedded into (\mathbb{R}^n, l^∞) , where $n = \#X + \#Y$, by Lemma 2.8. Let γ be the isometric embedding. We then have $d_H^\infty(\gamma \circ \gamma_X(X), \gamma \circ \gamma_Y(Y)) = d_H^\infty(\gamma_X(X), \gamma_Y(Y)) = \varepsilon$. Hence, inside (\mathbb{R}^n, l^∞) , the distance function δ_X to $\gamma \circ \gamma_X(X)$ and the distance function δ_Y to $\gamma \circ \gamma_Y(Y)$ are ε -close in the max norm. Moreover, since δ_X and δ_Y are lower envelopes of finitely many piecewise-linear functions, they are themselves piecewise-linear and therefore tame and continuous. As a consequence, their persistence diagrams are ε -close in the bottleneck distance, by Theorem 2.6.

Recall now that for all $\alpha \in \mathbb{R}$, the open α -sublevel-set of δ_X is the α -offset of $\gamma \circ \gamma_X(X)$, i.e. the union of the open l^∞ -balls of same radius α about the points of $\gamma \circ \gamma_X(X)$, noted $\gamma \circ \gamma_X(X)^\alpha$. Since the l^∞ -balls are hypercubes, they are convex and therefore their intersections are either empty or contractible. As a result, the Persistent Nerve Lemma 2.10 ensures that δ_X and the Čech filtration $\mathcal{C}(\gamma \circ \gamma_X(X), \mathbb{R}^n, l^\infty)$ have identical persistence diagrams. Furthermore, Lemma 2.9 guarantees that $\mathcal{C}(\gamma \circ \gamma_X(X), \mathbb{R}^n, l^\infty)$ coincides with the Rips filtration $\mathcal{R}(\gamma \circ \gamma_X(X), l^\infty)$, which in return is equal to $\mathcal{R}(X, d_X)$ since $\gamma \circ \gamma_X$ is an isometric embedding. As a consequence, the persistence diagram of $\mathcal{R}(X, d_X)$ coincides with

the persistence diagram of δ_X . The same is true for Y , therefore the persistence diagrams of $\mathcal{R}(X, d_X)$ and $\mathcal{R}(Y, d_Y)$ lie at bottleneck distance at most ε of each other. \square

Theorem 3.1 extends to finite metric spaces endowed with real valued functions as follows:

Theorem 3.2 Let (X, d_X) and (Y, d_Y) be two finite metric spaces endowed respectively with $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$. Then, for any $k \in \mathbb{N}$, the bottleneck distance between the persistence diagrams of the filtrations $\mathcal{R}(X, d_X, f)$ and $\mathcal{R}(Y, d_Y, g)$ is at most $d_{GH}^1((X, d_X, f), (Y, d_Y, g))$.

In the limit case where $f = g \equiv 0$, $\mathcal{R}(X, d_X, f)$ and $\mathcal{R}(Y, d_Y, g)$ become respectively $\mathcal{R}(X, d_X)$ and $\mathcal{R}(Y, d_Y)$, and we have $d_{GH}^1 = d_{GH}$. Thus, we obtain the statement of Theorem 3.1 as a special case of the above result.

Proof of Theorem 3.2. Let $\varepsilon = d_{GH}^1((X, d_X, f), (Y, d_Y, g))$. We follow the same strategy as in the proof of Theorem 3.1. For all $\alpha \in \mathbb{R}$, we let $X_\alpha = f^{-1}((-\infty, \alpha]) \subseteq X$ and $Y_\alpha = g^{-1}((-\infty, \alpha]) \subseteq Y$. As mentioned after definition 2.4, the infimum in Eq. (3) is in fact a minimum, realized by some correspondence $C \in \mathcal{C}(X, Y)$, and the disjoint union $Z = X \sqcup Y$ can be endowed with a suitable metric d_Z such that the canonical inclusions $\gamma_X : X \hookrightarrow Z$ and $\gamma_Y : Y \hookrightarrow Z$ are isometries onto their images, and such that for every pair $(x, y) \in C$ we have $d_Z(\gamma_X(x), \gamma_Y(y)) \leq \frac{1}{2} \|\Gamma_{X,Y}\|_{l^\infty(C \times C)} \leq \varepsilon$ and $|f(x) - g(y)| \leq \|f - g\|_{l^\infty(C)} \leq \varepsilon$. By Lemma 2.8, we can then embed the finite subspace $(\gamma_X(X) \cup \gamma_Y(Y), d_Z)$ into (\mathbb{R}^n, l^∞) , where $n = \#X + \#Y$, through some isometric embedding γ . We thus get two isometric embeddings $\gamma \circ \gamma_X$ and $\gamma \circ \gamma_Y$ of X and Y respectively into (\mathbb{R}^n, l^∞) , such that for every pair $(x, y) \in C$ we have $\|\gamma \circ \gamma_X(x) - \gamma \circ \gamma_Y(y)\|_\infty \leq \varepsilon$ and $|f(x) - g(y)| \leq \varepsilon$. It follows that the offsets filtrations $\{\gamma \circ \gamma_X(X_\alpha)^\alpha\}_{\alpha > 0}$ and $\{\gamma \circ \gamma_Y(Y_\alpha)^\alpha\}_{\alpha > 0}$ are ε -interleaved, in the following sense: $\forall \alpha > 0$,

$$\gamma \circ \gamma_X(X_\alpha)^\alpha \subseteq \gamma \circ \gamma_Y(Y_{\alpha+\varepsilon})^{\alpha+\varepsilon} \subseteq \gamma \circ \gamma_X(X_{\alpha+2\varepsilon})^{\alpha+2\varepsilon}.$$

Indeed, for any point $p \in \gamma \circ \gamma_X(X_\alpha)^\alpha$, there is an $x \in X$ such that $\|p - \gamma \circ \gamma_X(x)\|_\infty \leq \alpha$. Letting $y \in Y$ be such that $(x, y) \in C$, our embedding ensures that $\|\gamma \circ \gamma_X(x) - \gamma \circ \gamma_Y(y)\|_\infty \leq \varepsilon$ and $g(y) \leq f(x) + \varepsilon \leq \alpha + \varepsilon$. Hence, y belongs to $Y_{\alpha+\varepsilon}$, and by the triangle inequality we have $\|p - \gamma \circ \gamma_Y(y)\|_\infty \leq \alpha + \varepsilon$. It follows that $p \in \gamma \circ \gamma_Y(Y_{\alpha+\varepsilon})^{\alpha+\varepsilon}$. The second inclusion in the above equation follows by symmetry.

Since the two offsets filtrations are ε -interleaved, the extended stability Theorem [CCSG*09, Thm 4.4] guarantees that their persistence diagrams are ε -close to each other in the bottleneck distance. By Lemma 2.10, so are the persistence diagrams of the Čech filtrations $\{C_\alpha(\gamma \circ \gamma_X(X_\alpha), \mathbb{R}^n, l^\infty)\}_{\alpha > 0}$ and $\{C_\alpha(\gamma \circ \gamma_Y(Y_\alpha), \mathbb{R}^n, l^\infty)\}_{\alpha > 0}$. By Lemma 2.9, so are the persistence diagrams of the Rips filtrations $\{R_\alpha(\gamma \circ \gamma_X(X_\alpha), l^\infty)\}_{\alpha > 0}$ and $\{R_\alpha(\gamma \circ \gamma_Y(Y_\alpha), l^\infty)\}_{\alpha > 0}$. Finally, so are the diagrams of $\mathcal{R}(X, d_X, f_X)$ and $\mathcal{R}(Y, d_Y, f_Y)$, since these filtrations coincide with

$\{R\alpha(\gamma \circ \gamma_X(X\alpha), l^\infty)\}_{\alpha>0}$ and $\{R\alpha(\gamma \circ \gamma_Y(Y\alpha), l^\infty)\}_{\alpha>0}$ respectively. \square

4. Persistence-based signatures for finite metric spaces

The results of section 3 suggest that the persistence diagrams of Rips filtrations built on top of finite metric spaces endowed with *well-chosen* functions (see Definition 4.1) provide stable signatures to distinguish between different shapes.

Definition 4.1 Let \mathcal{H}_L be the class of maps $h : \mathcal{X} \rightarrow \mathcal{X}_1$ that preserve the metric and satisfy for some $L > 0$

$$\|f_X - f_Y\|_{l^\infty(C)} \leq L \cdot \frac{1}{2} \|\Gamma_{X,Y}\|_{l^\infty(C \times C)} \quad (5)$$

for all $(X, d_X), (Y, d_Y) \in \mathcal{X}$ and all $C \in \mathcal{C}(X, Y)$, where $h(X, d_X) = (X, d_X, f_X)$ and $h(Y, d_Y) = (Y, d_Y, f_Y)$.

From Definitions 2.4 and 4.1 and Theorem 3.2 we obtain:

Corollary 4.2 Let $X, Y \in \mathcal{X}_1$ be finite. Then, for any $h \in \mathcal{H}_L$,

$$\frac{d_B^\infty(D_k \mathcal{R}(f_X), D_k \mathcal{R}(f_Y))}{\max(1, L)} \leq d_{GH}(X, Y)$$

where $h(X, d_X) = (X, d_X, f_X)$, $h(Y, d_Y) = (Y, d_Y, f_Y)$ and $\mathcal{R}(f_X) = \mathcal{R}(X, d_X, f_X)$, $\mathcal{R}(f_Y) = \mathcal{R}(Y, d_Y, f_Y)$.

When the context is clear, we will abuse notations and omit some of the arguments in the expression $\mathcal{R}(X, d_X, f_X)$. Below we provide examples of maps that are in \mathcal{H}_L :

The diameter map. Our easiest example is the map $h_{\text{diam}} : (X, d_X) \mapsto (X, d_X, f_X)$ where $f_X : X \rightarrow \mathbb{R}$ is the constant function $f_X \equiv \text{diam}(X)$. It is easily seen that $h_{\text{diam}} \in \mathcal{H}_2$.

The eccentricity map. The map $h_{e_1} : (X, d_X) \mapsto (X, d_X, e_1^X)$ where $e_1^X : X \rightarrow \mathbb{R}$ is given by $e_1^X(x) = \max_{x' \in X} d_X(x, x')$ is in \mathcal{H}_2 . Indeed, assume $X, Y \in \mathcal{X}$ and $C \in \mathcal{C}(X, Y)$, and let $\eta = \|\Gamma_{X,Y}\|_{l^\infty(C \times C)}$. Pick $(x, y), (x', y') \in C$ and notice that

$$d_X(x, x') \leq d_Y(y, y') + \eta \leq \max_{y' \in Y} d_Y(y, y') + \eta.$$

Hence, $e_1^X(x) \leq e_1^Y(y) + \eta$ for all $(x, y) \in C$. By symmetry, $|e_1^X(x) - e_1^Y(y)| \leq \eta$ for all $(x, y) \in C$.

Higher-order eccentricities. A more general construction is as follows: given $m \in \mathbb{N}$, let $h_{e_m} : \mathcal{X} \rightarrow \mathcal{X}_1$ be the map $(X, d_X) \mapsto (X, d_X, e_m^X)$, where $e_m^X : X \rightarrow \mathbb{R}$ is defined by

$$\forall x_0 \in X, e_m^X(x_0) = \max_{x_1, \dots, x_m} \min_{0 \leq i < j \leq m} d_X(x_i, x_j),$$

where x_1, \dots, x_m range over X . It turns out that h_{e_m} is in \mathcal{H}_2 . In addition, observe that for $m > 1$, e_m^X provides more information than just e_1^X : consider for instance S^n endowed with the usual intrinsic metric. Then, $e_1^{S^n} \equiv \pi$ for all $n \in \mathbb{N}$, whereas $e_4^{S^1} \equiv \pi/4$ is different from $e_4^{S^2} \equiv 2 \arcsin(\sqrt{2/3}) > \pi/4$. Note however that the complexity of computing e_m^X on a finite metric space with N points is roughly $O(N^{m+1})$.

An alternative to e_m^X . Instead of the eccentricities, one may consider using the following inductive family of functions: $f_0^X \equiv 0$, and $f_{m+1}^X(x) = \max_{x' \in X} (d_X(x, x') + f_m^X(x'))$. It can be shown that for each $m \in \mathbb{N}$ the map $h_m : (X, d_X) \mapsto (X, d_X, f_m^X)$ is in \mathcal{H}_m . The complexity associated with computing h_m is $O(m \cdot N^2)$, where N is the cardinality of the metric space.

Remark 1 The above examples give rise to a large variety of maps in \mathcal{H}_L . Indeed, given $h : (X, d_X) \mapsto (X, d_X, f_X)$ and $h' : (X, d_X) \mapsto (X, d_X, f'_X)$ in \mathcal{H}_L , observe that $\max(h, h')$ and $\sup(h) : (X, d_X) \mapsto (X, d_X, \max(f_X))$ are both in \mathcal{H}_L . Moreover, $h + h' \in \mathcal{H}_{2L}$, and for any $\lambda \in \mathbb{R}$, $\lambda h \in \mathcal{H}_{|\lambda|L}$. It is important to note that the persistence diagrams of $\mathcal{R}(f_X)$ and $\mathcal{R}(\lambda f_X)$ are not related by a simple transformation: different choices of λ give different signatures.

5. Introducing measures

For the sake of generality, it is important to make the class of admissible functions as large as possible. Unfortunately, it seems that certain canonical functions such as the average distance to a point are not directly controllable by the GH distance. This is intuitively clear since in averaging the values $\sum_{x' \neq x} d_X(x, x') \alpha(x')$ there is a choice to be made for the *weights* $\{\alpha(x')\}_{x' \in X}$, the choice $\alpha(x') = \frac{1}{|X|}$ being one of many possible. In order to incorporate this extra structure (weights) into our formulas, we define a new metric on the class of spaces that arise in [Mém07, Mém08]. This metric belongs to the class of *Gromov-Wasserstein* metrics on metric spaces enriched with weights.

A *measure metric space* (mm space, for short) is a triple (X, d_X, μ_X) where $(X, d_X) \in \mathcal{X}$ and μ_X is a Borel probability measure such that $\text{supp}[\mu_X] = X$. Here, $\text{supp}[\mu_X]$ denotes the *support* of X , i.e. the minimal closed set Z s.t. $\mu_X(X \setminus Z) = 0$. Let \mathcal{X}^w denote the collection of all mm spaces, and \mathcal{X}_1^w the collection of all mm spaces endowed with continuous real-valued functions. Two mm spaces (X, d_X, μ_X, f_X) and (Y, d_Y, μ_Y, f_Y) are said to be *isomorphic* if there exists an isometry $\phi : X \rightarrow Y$ s.t. $f_X = f_Y \circ \phi$ and $\mu_X(\phi^{-1}(Z)) = \mu_Y(Z)$ for all measurable sets $Z \subset Y$. It turns out that we can metrize the set of isomorphism classes of \mathcal{X}_1^w by replacing correspondences between X and Y in \mathcal{X}^w by *couplings*, that is, probability measures on $X \times Y$ whose marginals on the first and second factors are μ_X and μ_Y respectively. Since X, Y are compact and μ_X, μ_Y have full support, the support of a coupling is a correspondence [Mém07, Lemma 1].

Given $X, Y \in \mathcal{X}_1^w$ we then define $d_{GW, \infty}^1(X, Y)$ to be the infimal $\eta > 0$ such that there exists a coupling μ for which both $\frac{1}{2} \|\Gamma_{X,Y}\|_{l^\infty(C(\mu) \times C(\mu))} \leq \eta$ and $\|f_X - f_Y\|_{l^\infty(C(\mu))} \leq \eta$, where $C(\mu)$ denotes the correspondence $\text{supp}[\mu]$.

For any $X, Y \in \mathcal{X}^w$, we define the *Gromov-Wasserstein distance* $d_{GW, \infty}(X, Y)$ to be equal to $d_{GW, \infty}^1(X, Y)$, where we regard X, Y as elements in \mathcal{X}_1^w by endowing them

with the zero functions. One can prove that $d_{GH}(X, Y) \leq d_{GW, \infty}(X, Y)$ for all $X, Y \in \mathcal{X}^w$ [Mém07, Proposition 6].

The approach of Sections 3 and 4 generalizes almost verbatim to this measured setting. In particular, we have the following analogous of Corollary 4.2. Let \mathcal{H}_L^w be the class of maps $h: \mathcal{X}^w \rightarrow \mathcal{X}_1^w$ that preserve the metric and the measure and satisfy, for some $L > 0$,

$$\|f_X - f_Y\|_{l^\infty(C(\mu))} \leq L \cdot \frac{1}{2} \|\Gamma_{X,Y}\|_{l^\infty(C(\mu) \times C(\mu))} \quad (6)$$

for all $(X, d_X), (Y, d_Y) \in \mathcal{X}^w$ and all $\mu \in \mathcal{M}(\mu_X, \mu_Y)$. Here, $h(X, d_X, \mu_X) = (X, d_X, \mu_X, f_X)$, $h(Y, d_Y, \mu_Y) = (Y, d_Y, \mu_Y, f_Y)$, and $\mathcal{M}(\mu_X, \mu_Y)$ denotes the set of all couplings between (X, μ_X) and (Y, μ_Y) .

Corollary 5.1 Let $X, Y \in \mathcal{X}^w$ be finite. Then for any $h \in \mathcal{H}_L^w$,

$$\frac{d_B^\infty(D_k \mathcal{R}(h(X)), D_k \mathcal{R}(h(Y)))}{\max(1, L)} \leq d_{GW, \infty}(X, Y).$$

The class \mathcal{H}_L^w is actually larger than \mathcal{H}_L . For instance, given $(X, d_X, \mu_X) \in \mathcal{X}^w$ and $p \geq 1$ we define the p -eccentricity function as

$$s_{X,p}(x) := \left(\int_X d_X(x, x')^p \mu_X(dx') \right)^{1/p},$$

and for $p = \infty$, $s_{X,\infty}(x) := \max_{x' \in X} d_X(x, x')$.

Lemma 5.2 The p -eccentricity map $h_{s_p}: \mathcal{X}^w \rightarrow \mathcal{X}_1^w$ defined by $(X, d_X, \mu_X) \mapsto (X, d_X, \mu_X, s_{X,p})$ is in \mathcal{H}_L^w .

Proof. let $X, Y \in \mathcal{X}^w$ and $\mu \in \mathcal{M}(\mu_X, \mu_Y)$. Then, for all $x \in X$ and $y \in Y$, the triangle inequality in the $l^p(\mu)$ -norm implies:

$$\left| \|d_X(x, \cdot)\|_{l^p(\mu)} - \|d_Y(y, \cdot)\|_{l^p(\mu)} \right| \leq \|\Gamma_{X,Y}(x, y, \cdot, \cdot)\|_{l^p(\mu)}.$$

Since $\mu \in \mathcal{M}(\mu_X, \mu_Y)$, it follows easily that $\|d_X(x, \cdot)\|_{l^p(\mu)} = s_{X,p}(x)$ and $\|d_Y(y, \cdot)\|_{l^p(\mu)} = s_{Y,p}(y)$. Hence,

$$\begin{aligned} |s_{X,p}(x) - s_{Y,p}(y)| &\leq \|\Gamma_{X,Y}(x, y, \cdot, \cdot)\|_{l^p(\mu)} \\ &\leq \|\Gamma_{X,Y}(x, y, \cdot, \cdot)\|_{l^\infty(C(\mu))}. \end{aligned}$$

In particular, for all $(x, y) \in C(\mu)$,

$$|s_{X,p}(x) - s_{Y,p}(y)| \leq \|\Gamma_{X,Y}\|_{l^\infty(C(\mu) \times C(\mu))},$$

hence

$$\|s_{X,p} - s_{Y,p}\|_{l^\infty(C(\mu))} \leq \|\Gamma_{X,Y}\|_{l^\infty(C(\mu) \times C(\mu))},$$

which means that h_{s_p} is in \mathcal{H}_L^w (cf. Eq. (6)). \square

6. Algorithmic aspects

In implementing the lower bounds given by Corollaries 4.2 and 5.1, one must take into account the facts that (1) the computation of the full persistence diagrams may be too expensive and (2) the bottleneck matching of the diagrams may be time consuming. This section addresses these two issues.

6.1. Thresholding filtrations

To deal with problem (1) above, we restrict the computation of our filtrations as follows: given a filtration $\emptyset = K_{\alpha_0} \subseteq \dots \subseteq K_{\alpha_m} = K$ and a threshold $t \in \mathbb{R}$, we let m be such that $\alpha_{m-1} < t \leq \alpha_m$ and we consider only the t -thresholded filtration $\emptyset = K_{\alpha_0} \subseteq \dots \subseteq K_{\alpha_{m-1}} \subseteq K$. The reason why we replace K_m by K is to preserve bottleneck distances, as will be shown below. The drawback is that in principle we should still have to compute the full complex K . Nevertheless, we will see that this is not necessary in practice. The thresholding operation affects persistence diagrams in the following way (see also Figure 3):

Lemma 6.1 Given a filtration \mathcal{K} and its t -thresholded variant $\bar{\mathcal{K}}^t$, there is a multi-bijection $\gamma: D_k \mathcal{K} \rightarrow D_k \bar{\mathcal{K}}^t$ such that:

- the restriction of γ to the lower-left quadrant $(-\infty, t] \times (-\infty, t]$ is the identity;
- the restriction of γ to the upper-left quadrant $(-\infty, t] \times (t, +\infty)$ is the vertical projection onto the line $y = t$;
- the restriction of γ to the line $y = +\infty$ is the horizontal projection onto the extended half-plane $(-\infty, t] \times \mathbb{R}$;
- finally, the restriction of γ to the half-plane $(t, +\infty) \times \mathbb{R}$ is the projection onto the diagonal point (t, t) .

The proof of this rather intuitive result is omitted. Lemma 6.1 implies that for any filtrations \mathcal{F}, \mathcal{G} and any threshold t :

$$d_B^\infty(D_k \bar{\mathcal{F}}^t, D_k \bar{\mathcal{G}}^t) \leq d_B^\infty(D_k \mathcal{F}, D_k \mathcal{G}). \quad (7)$$

This means that the results of Sections 3 through 5 still hold when the filtrations are replaced by their t -thresholded versions. In addition, Lemma 6.1 provides a way of computing $D_k \bar{\mathcal{R}}^t(X, d_X, f_X)$ without building $R_\infty(X, d_X, f_X)$ explicitly. Indeed, we know that $R_\infty(X, d_X, f_X)$ is the full simplex on the vertex set X , therefore the filtration $\mathcal{R}(X, d_X, f_X)$ has only one essential homology class $[c]$, easily identified as the connected component created by the point of X with minimum f_X -value. It follows that $D_k \bar{\mathcal{R}}^t(X, d_X, f_X)$ can be obtained first by computing the diagram of the truncated filtration $\{R_\alpha(X_\alpha)\}_{0 < \alpha \leq t}$, and then by projecting the resulting multi-set (save the point corresponding to $[c]$ when $k = 0$) vertically onto the lower-left quadrant $(-\infty, t] \times (-\infty, t]$.

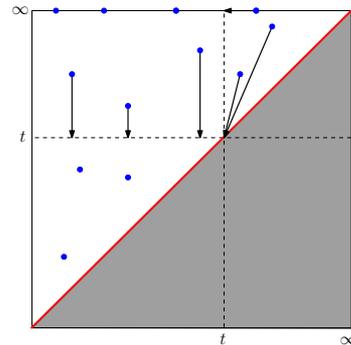


Figure 3: Effect of thresholding on a persistence diagram.

6.2. Simplifying persistence diagrams

To address problem (2) above, we simplify each considered persistence diagram $D_k\mathcal{F}$ as follows: given a parameter $\varepsilon \geq 0$, we snap all points of $D_k\mathcal{F}$ lying within l^∞ -distance ε of the diagonal Δ onto Δ . The resulting simplified diagram is called $D_k^\varepsilon\mathcal{F}$. The triangle inequality implies that for any persistence diagrams $D_k\mathcal{F}, D_k\mathcal{G}$ and any parameters $\varepsilon, \varepsilon' \geq 0$,

$$\left| d_B^\infty(D_k^\varepsilon\mathcal{F}, D_k^{\varepsilon'}\mathcal{G}) - d_B^\infty(D_k\mathcal{F}, D_k\mathcal{G}) \right| \leq \max\{\varepsilon, \varepsilon'\}. \quad (8)$$

In our experiments, we specify a maximal number N of off-diagonal points to be kept per diagram. Specifically, we keep the points with largest l^∞ -distance to Δ . For every diagram $D_k\mathcal{F}$ considered, we perform this simplification and record the incurred projection error $\varepsilon \geq 0$, which remains small if N is chosen sufficiently large.

6.3. Computing the bottleneck distance

Given two simplified persistence diagrams $D_k^\varepsilon\mathcal{F}$ and $D_k^{\varepsilon'}\mathcal{G}$, each containing at most N off-diagonal points, we proceed as follows. Firstly, we discard the diagonal and consider only $D_k^\varepsilon\mathcal{F} \setminus \Delta$ and $D_k^{\varepsilon'}\mathcal{G} \setminus \Delta$, whose total multiplicities are finite, bounded from above by N . Secondly, we add arbitrary diagonal points to $D_k^\varepsilon\mathcal{F} \setminus \Delta$ and $D_k^{\varepsilon'}\mathcal{G} \setminus \Delta$, so that the resulting multi-sets $\{f_i\}_i$ and $\{g_j\}_j$ contain exactly N points each. Finally, we solve the following *Bottleneck Assignment Problem*: compute $c^* = \min_\pi \max_i \vartheta(f_i, g_{\pi(i)})$, where π ranges over all permutations of the set $\{1, \dots, N\}$, and where

$$\vartheta(p, q) = \min \left\{ \|p - q\|_\infty, \frac{1}{2} \max\{|p_y - p_x|, |q_y - q_x|\} \right\}.$$

In plain words, $\vartheta(p, q)$ measures the cost of matching p with q against the cost of matching both points with the diagonal Δ . It can then be easily checked that $c^* = d_B^\infty(D_k^\varepsilon\mathcal{F}, D_k^{\varepsilon'}\mathcal{G})$. In addition, we have that $\vartheta(p, q) = \vartheta(p, q')$ whenever $q, q' \in \Delta$, which makes the choice of the locations of the diagonal points inserted during the second step of our procedure unimportant. In practice, we compute c^* using the *thresholding algorithm* of [BC99, Section 5], whose running time is $O(N^{2.5} \log N)$.

7. Application to shape classification

We exemplify the use of our framework on a shape classification problem under invariances. We used the publicly available database of (triangulated) shapes [SP04]. This database comprises 62 shapes from six different classes: camel, cat, elephant, face, head and horse. Each class contains several different *poses* of the same shape. These poses are richer than just rigid isometries, as can be seen from Figure 1. The number of vertices in the models range from $7K$ to $30K$. We equip each model X_i with an *intrinsic distance* d_i between pairs of points in X_i using Dijk-

stra's algorithm on the 1-skeleton graph $G(X_i)$ of the mesh[‡]. We then select a uniform subset X_i^r of $4K$ vertices using the Euclidean *farthest point sampling* procedure. Briefly, we first randomly choose a point from the vertex set. Then, we choose the second point as the one at maximal distance from the first one. Subsequent points are chosen always to maximize the distance to the points already chosen. Let X_i^r denote this reduced model with $4K$ vertices. We further subsample X_i^r using the farthest point procedure, this time in the distance d_i computed using $G(X_i)$, and retain only 300 points. We denote the resulting coarse set by X_i^{rr} , which we endow with the distance d_i , properly normalized so that the diameter of X_i^{rr} is 1. We also equip X_i^{rr} with a probability measure μ_i based on intrinsic Voronoi partitions, where the measure at point $x \in X_i^{rr}$ equals the proportion of points of X_i^r that are closer to x than to any other point of X_i^{rr} . From each model X_i we thus obtain a discrete mm-space (X_i^{rr}, d_i, μ_i) . Let S denote the collection of all 62 such mm-spaces.

Let $P := \{1, 2, 3, \infty\}$, $\mathcal{H}^+ = \{h_{s_p}, p \in P\} \subset \mathcal{H}_2^w$ and $\mathcal{H}^- = \{\sup(h_{s_p}) - h_{s_p}, p \in P\} \subset \mathcal{H}_4^w$ (recall Remark 1). Let \mathcal{H} denote the union of these two sets of functions, and let

$$\Lambda := \{0, 0.01, 0.1, 0.2, \dots, 0.9, 1, 2, 3, 4, 5, 10, 30, 50, 70, 90\}.$$

For each pair of mm-spaces X and Y in S , for a choice of $\lambda \in \Lambda$ and a function $h \in \mathcal{H}$ we estimate the lower bound given by Corollary 5.1 as follows:

1. We first construct the thresholded filtrations $\overline{\mathcal{R}(\lambda \cdot h(X))}^t$ and $\overline{\mathcal{R}(\lambda \cdot h(Y))}^t$ up to their 3-dimensional skeleta, as discussed in Section 6.1. The thresholding parameter t is chosen so as to guarantee that the 1-skeleta of the underlying complexes of $\overline{\mathcal{R}(\lambda \cdot h(X))}^t$ and $\overline{\mathcal{R}(\lambda \cdot h(Y))}^t$ contain no more than a fraction s of the (quadratic) number of edges of the complete graph. The value of s chosen in our experiments is 7%, which allows a reasonable execution time of approximately 1.5 seconds per pair of mm-spaces, on average.
2. We use the software `jplex [JPI]` and the procedure of Section 6.1 to compute the k th persistence diagrams of each thresholded filtration, for $k = 0, 1, 2$.
3. We apply the simplification step described in Section 6.2 with $N = 30$. For each $k = 0, 1, 2$, let $E_k(X) = D_k^{\varepsilon_k(X)} \overline{\mathcal{R}(\lambda \cdot h(X))}^t$ and $E_k(Y) = D_k^{\varepsilon_k(Y)} \overline{\mathcal{R}(\lambda \cdot h(Y))}^t$ be the simplified diagrams, where $\varepsilon_k(X)$ and $\varepsilon_k(Y)$ are the projection errors.
4. We use the procedure of Section 6.3 to compute the quantity $c(h, \lambda; X, Y) =$

$$\frac{\max_k \max \left\{ d_B^\infty(E_k(X), E_k(Y)) - \max\{\varepsilon_X(k), \varepsilon_Y(k)\}, 0 \right\}}{\max(1, \lambda \cdot L(h))},$$

[‡] In this particular application we take advantage of the fact that a mesh is available for each of the shapes. In full generality, one can replace $G(X_i)$ by some neighborhood graph in the ambient Euclidean metric, with similar guarantees.

where $L(h)$ is 2 if $h \in \mathcal{H}^+$ and 4 if $h \in \mathcal{H}^-$. By Corollary 5.1, Lemma 6.1 and Eq. (8), we know that $d_{GW,\infty}(X, Y) \geq c(h, \lambda; X, Y)$.

For any pair $X, Y \in S$, we can now bound $d_{GW,\infty}(X, Y)$ from below by $c(X, Y) := \max_{h \in \mathcal{H}, \lambda \in \Lambda} c(h, \lambda; X, Y)$. Let M denote the matrix with elements $c(X, Y)$, $X, Y \in S$. Figure 4 (left) shows the result of these computations on our database, while Figure 5 shows the result of applying the metric multi-dimensional scaling procedure implemented by the Matlab command `mdscale` to M .

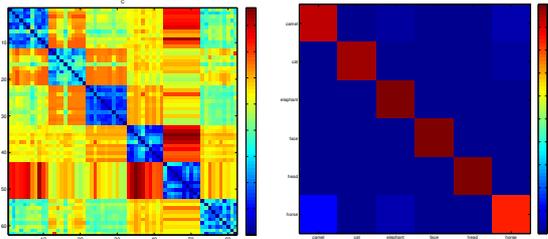


Figure 4: Left: estimate of the Gromov-Wasserstein distance computed on the database S . Right: estimated confusion matrix for the 1-nearest neighbor classification problem described in the text. The overall error rate is 4%.

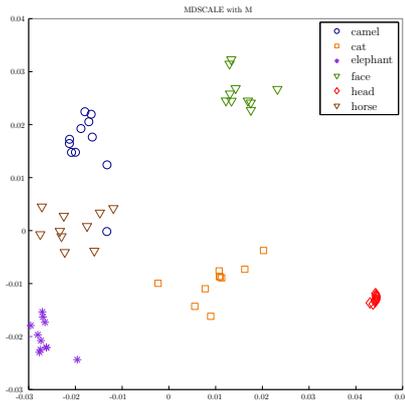


Figure 5: MDS plot of the matrix M with labels corresponding to each class. The overall error is estimated to be 4%. Notice how there is excellent clustering of the different classes, save for a single shape of the class `camel` that lies very close to the cluster formed by the class `horse`. Removal of this shape reduces the classification error to 2%.

In order to evaluate the discriminative power contained in the matrix M , we consider a classification task as follows: We randomly select one shape from each class, form a training set T and use it for performing 1-nearest neighbor classification (where *nearest* is with respect to the metric defined by M) of the remaining shapes. By simple comparison between the class predicted by the classifier and the actual class to which the shape belongs we obtain an estimate of

the probability $P_e(M)$ of mis-classification. We repeat this procedure for $2K$ random choices of the training set. Using the same randomized procedure we obtain an estimate of the confusion matrix for this problem, whose entry (i, j) is the probability that the classifier will assign class j to a shape when the actual class is i — see Figure 4 (right). We obtained $P_e(M) = 4\%$. As Figure 5 shows, there is a shape of class `camel` that lies very close to the cluster formed by the class `horse`. Removal of this shape from the database reduced the probability of error to 2%.

8. Perspectives

In this paper we only used Euclidean and geodesic distances to illustrate our approach, however many other metrics may be considered, such as the diffusion distance [BBK*].

The generality of our approach makes it applicable in a large variety of contexts beyond 3d shape classification. Our signatures are indeed well-defined, stable and computable for any arbitrary finite metric space that is given in the form of a distance matrix or equivalent representation. This makes our approach well-suited for classification problems in general, even in cases where the data are high-dimensional or lying in some non-Euclidean space without any information on the embedding beside the inter-point distances.

At a more theoretical level, let us point out that Theorem 3.1 enables us to define *limit Rips persistence diagrams* for compact metric spaces. Indeed, given a compact metric space (Z, d_Z) , consider a sequence of finite ϵ -samples of (Z, d_Z) such that ϵ converges to zero. The k -th Rips persistence diagram of (Z, d_Z) can then be defined as the limit of the k -th persistence diagrams of the Rips filtrations of the sequence. This limit is independent of the choice of the sequence of ϵ -samples, since any two finite ϵ -samples of (Z, d_Z) lie at Gromov-Hausdorff distance 2ϵ of each other, which by Theorem 3.1 implies that the persistence diagrams of their Rips filtrations are 2ϵ -close in the bottleneck distance. Thus, the limit Rips persistence diagram of (Z, d_Z) is well-defined, and by Theorem 3.1 it is stable under small perturbations of the space in the Gromov-Hausdorff distance. Similarly, for any real-valued function $f : X \rightarrow \mathbb{R}$, Theorem 3.2 enables to define limit Rips persistence diagrams for (X, d_X, f) and to guarantee their stability. This concept of limit persistence diagram of a compact metric space (X, d_X) is reminiscent of the approach to topological persistence developed in [CCSG*09]. However, it may not be equivalent to the persistence diagram of the limit filtration $\mathcal{R}(X, d_X)$ or $\mathcal{R}(X, d_X, f)$ in general. Finding the exact relationships that exist between these diagrams is still open.

Acknowledgements. This work was carried out within the Associate Team *TGDA: Topological and Geometric Data Analysis* involving the Geometrica group at INRIA and the Geometric Computing group at Stanford University. The authors wish to thank the referees for their suggestions to make

the paper more accessible to non-specialists. They acknowledge the support of NSF through grants ITR 0205671, FRG 0354543 and FODAVA 808515, of NIH through grant GM-072970, of ONR through grant N00014-09-1-0783, and of DARPA through grant HR0011-05-1-0007.

References

- [BBI01] BURAGO D., BURAGO Y., IVANOV S.: *A Course in Metric Geometry*, vol. 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001.
- [BBK*] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. Submitted.
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing* 28, 5 (2006), 1812–1836.
- [BBK09] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Topology-invariant similarity of nonrigid shapes. *Intl. Journal of Computer Vision (IJCV)* 81, 3 (March 2009), 281–301.
- [BC99] BURKARD R., CELA E.: Linear assignment problems and extensions. In *Handbook of Combinatorial Optimization, supplement vol. A*, Du D.-Z., Pardalos P. M., (Eds.). Kluwer Academic Publishers, 1999, pp. 75–149.
- [BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 4 (April 2002), 509–522.
- [CCS07] CHAZAL F., COHEN-STEINER D.: *Geometric Inference*. to appear as a book chapter, Springer, 2007.
- [CCSG*09] CHAZAL F., COHEN-STEINER D., GUIBAS L. J., GLISSE M., OUDOT S. Y.: Proximity of persistence modules and their diagrams. In *Proc. 25th ACM Symp. Comp. Geom.* (2009).
- [CFL06a] CERRI A., FROSINI P., LANDI C.: Stability in multi-dimensional size theory, 2006.
- [CFL06b] CERRI A., FROSINI P., LANDI C.: Stability in multi-dimensional size theory, 2006.
- [CM09] CARLSSON G., MÉMOLI F.: Characterization and stability of hierarchical clustering methods. <http://comptop.stanford.edu/preprints/>, 2009.
- [CO08] CHAZAL F., OUDOT S. Y.: Towards persistence-based reconstruction in Euclidean spaces. In *Proc. 24th ACM Sympos. Comput. Geom.* (2008), pp. 232–241.
- [CSEH05] COHEN-STEINER D., EDELSBRUNNER H., HARER J.: Stability of persistence diagrams. In *Proc. 21st ACM Sympos. Comput. Geom.* (2005), pp. 263–271.
- [CZCG05] CARLSSON G., ZOMORODIAN A., COLLINS A., GUIBAS L. J.: Persistence barcodes for shapes. *International Journal of Shape Modeling* 11, 2 (2005), 149–187.
- [dFL05] D’AMICO M., FROSINI P., LANDI C.: *Natural pseudo-distance and optimal matching between reduced size functions*. Tech. Rep. 66, DISMI, Univ. degli Studi di Modena e Reggio Emilia, Italy., 2005.
- [dFL06] D’AMICO M., FROSINI P., LANDI C.: Using matching distance in size theory: A survey. *IJIST* 16, 5 (2006), 154–161.
- [EH] EDELSBRUNNER H., HARER J.: *Persistent Homology - A survey*. eds J. E. Goodman, J. Pach and R. Pollack, AMS, to appear.
- [EK03] (ELBAZ) A. E., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 10 (2003), 1285–1295.
- [ELZ00] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. In *Proc. 41st Annu. IEEE Sympos. Found. Comput. Sci.* (2000), pp. 454–463.
- [Fro90] FROSINI P.: A distance for similarity classes of submanifolds of euclidean space. *Bull. Austral. Math. Soc.* 42:3 (1990), 407–416.
- [GM05] GHRIST R., MUHAMMAD A.: Coverage and hole-detection in sensor networks via homology. In *Proc. the 4th International Symposium on Information Processing in Sensor Networks (IPSN’05)* (2005), pp. 254–260.
- [Gro99] GROMOV M.: *Metric structures for Riemannian and non-Riemannian spaces*, vol. 152 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1999.
- [Hat01] HATCHER A.: *Algebraic Topology*. Camb. Univ. Press, 2001.
- [HK03] HAMZA A. B., KRIM H.: Geodesic object representation and recognition. In *Lecture Notes in Computer Science*, vol. 2886. Springer, 2003, pp. 378–387.
- [HSKK01] HILAGA M., SHINAGAWA Y., KOHMURA T., KUNII T. L.: Topology matching for fully automatic similarity estimation of 3d shapes. In *SIGGRAPH ’01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM, pp. 203–212.
- [JPI] JPlex: Persistent Homology Computations Library. <http://comptop.stanford.edu/programs/jplex/index.html>.
- [MCH*06] MANAY S., CREMERS D., HONG B., YEZZI A., SOATTO S.: Integral invariants for shape matching. *PAMI* 28, 10 (October 2006), 1602–1618.
- [Mém07] MÉMOLI F.: On the use of Gromov-Hausdorff Distances for Shape Comparison. Botsch M., Pajarola R., Chen B., Zwicker M., (Eds.), Eurographics Association, pp. 81–90.
- [Mém08] MÉMOLI F.: Gromov-Hausdorff distances in Euclidean spaces. In *CVPR Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (NORDIA’08)* (2008).
- [MS04] MÉMOLI F., SAPIRO G.: Comparing point clouds. In *SGP ’04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing* (New York, NY, USA, 2004), ACM, pp. 32–40.
- [MS05] MÉMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.* 5, 3 (2005), 313–347.
- [Mun84] MUNKERS J. R.: *Elements of algebraic topology*. Addison-Wesley, Redwood City, California, 1984.
- [OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Trans. Graph.* 21, 4 (2002), 807–832.
- [Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Symposium on Geometry Processing* (2007), pp. 225–233.
- [RWP05] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-spectra as fingerprints for shape matching. In *SPM ’05: Proceedings of the 2005 ACM symposium on Solid and physical modeling* (New York, NY, USA, 2005), ACM Press, pp. 101–106.
- [SP04] SUMNER R. W., POPOVIĆ J.: Deformation transfer for triangle meshes. In *SIGGRAPH ’04: ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), ACM, pp. 399–405.
- [ZC05] ZOMORODIAN A., CARLSSON G.: Computing persistent homology. *Discrete Comput. Geom.* 33, 2 (2005), 249–274.

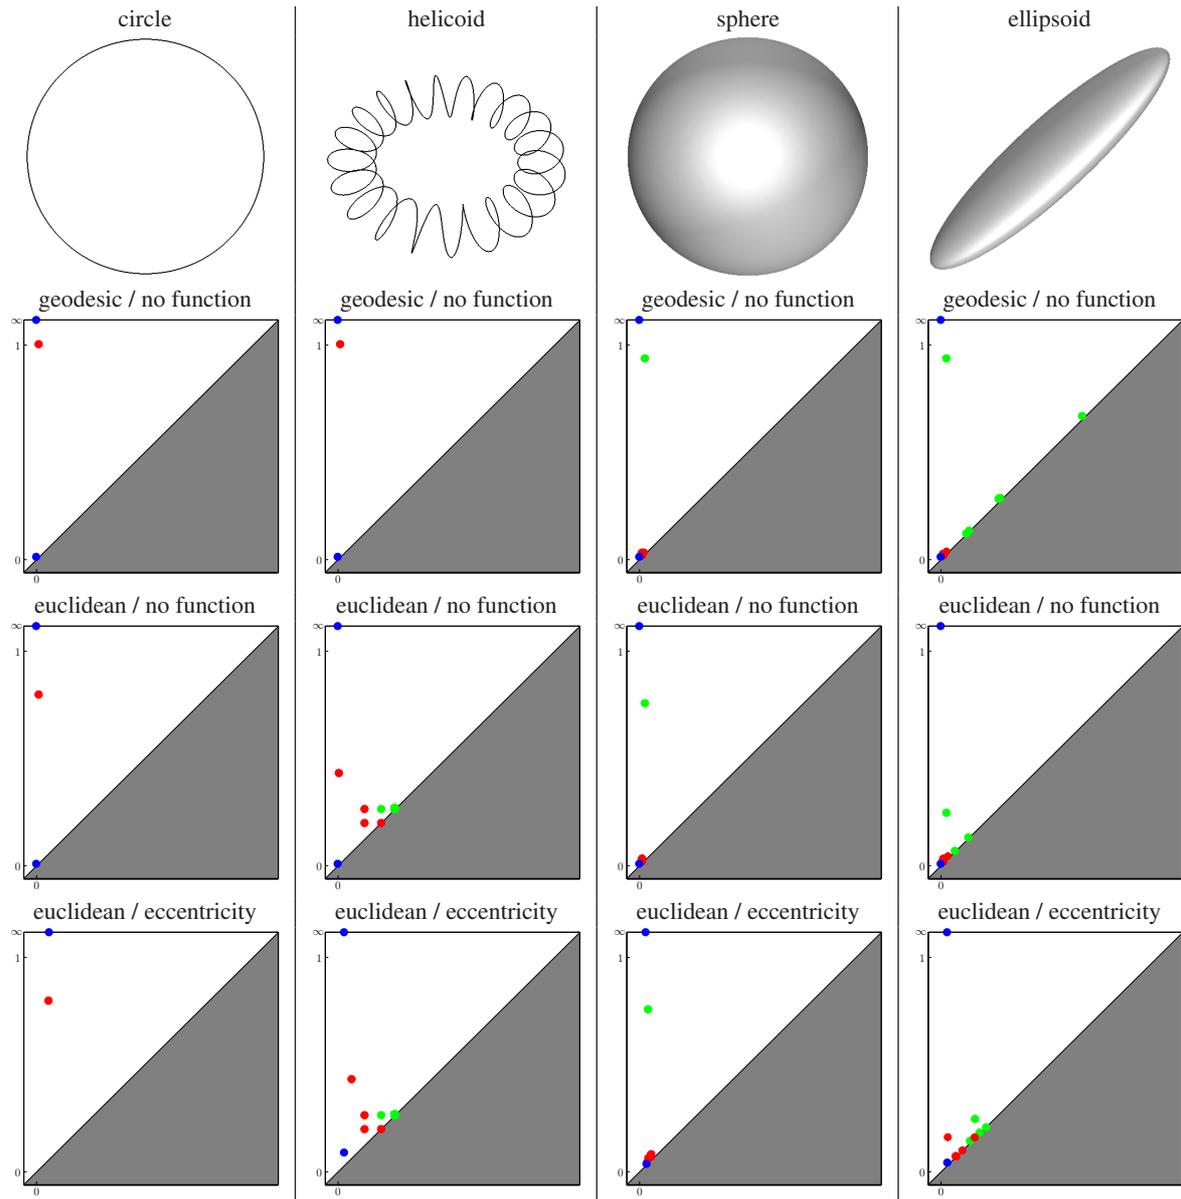


Figure 6: Some toy examples, from left to right: the unit circle, a helical curve of same length 2π drawn on a torus, the unit sphere, an ellipsoid whose smallest equatorial ellipse has same length 2π as the equator of the sphere. On each shape, a uniform 0.0125-sample has been generated, on top of which various Rips filtrations have been constructed using different metrics and functions: geodesic distance and no function (second row), Euclidean distance and no function (third row), Euclidean distance and 0.2 times the eccentricity (fourth row). The corresponding persistence diagrams are presented in rows 2 through 4: each picture shows the 0-dimensional (blue), 1-dimensional (red), and 2-dimensional (green) diagrams super-imposed.