

Spectral Gromov-Wasserstein Distances for Shape Matching

Facundo Mémoli
Stanford University, Mathematics
Stanford, CA 94305

memoli@math.stanford.edu

<http://math.stanford.edu/~memoli>

Abstract

We introduce a spectral notion of distance between shapes and study its theoretical properties. We show that our distance satisfies the properties of a metric on the class of isometric shapes, which means, in particular, that two shapes are at 0 distance if and only if they are isometric. Our construction is similar to the recently proposed Gromov-Wasserstein distance, but rather than viewing shapes merely as metric spaces, we define our distance via the comparison of heat kernels. This allows us to relate our distance to previously proposed spectral invariants used for shape comparison, such as the spectrum of the Laplace-Beltrami operator. In addition, the heat kernel provides a natural notion of scale, which is useful for multi-scale shape comparison. We also prove a hierarchy of lower bounds for our distance, which provide increasing discriminative power at the cost of increase in computational complexity.

1. Introduction

Many approaches have been proposed in the context of (pose invariant) shape classification and recognition, including the pioneering work on *size theory* by Frosini and collaborators [11], the *shape contexts* of Belongie et al. [2], the integral invariants of [18], the eccentricity functions of [15], the *shape distributions* of [25], the *canonical forms* of [10], and the *Shape DNA* methods in [26]. The common idea of these methods is to compute certain metric invariants, or *signatures* of the shapes. These signatures are then embedded into a common metric space to facilitate comparison, and shapes whose signatures have a small distance are considered similar.

Unfortunately, the question of proving that a given family of signatures captures proximity or similarity of shapes in a reasonable way has hardly been 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 easier question of whether the similarity between two shapes forces their signatures to be

similar. These questions cannot be well formulated until one agrees on notions of (1) *equality* and (2) *similarity* between shapes, such that the most similar shapes to a given shape, are those that are considered equal to it.

One way to address these issues was suggested by Mémoli and Sapiro [21, 22], who consider shapes as metric spaces, and (1) define two shapes to be equal when they are *isometric*, and (2) use the Gromov-Hausdorff distance [14] as a measure of dissimilarity between shapes. Once this framework is adopted, a natural question arises: whether a given family of signatures is *stable* under perturbations of the shape in the Gromov-Hausdorff sense. Unfortunately, despite its generality, it has been difficult to use the Gromov-Hausdorff framework in a natural way for explaining most of the existing shape matching procedures. In addition, the computation of the GH distance in practice leads to NP-hard combinatorial optimization problems.

These problems were addressed by Mémoli in [19, 20], who defines the *Gromov-Wasserstein distances* based on ideas from *mass transportation*. This family of distances (1) exhibits a number of desirable theoretical features, (2) directly yields continuous variable quadratic optimization problems with linear constraints, and (3) provides a dissimilarity measure under which a large number of shape signatures can be shown to be *stable*. The framework of [19] assumes that in addition to a metric, shapes are also endowed with a certain notion of *weight* associated to each point of the shape. In this context, *stability* of an invariant signature means that one can write precise lower bounds for the GW distance which encode a certain comparison of the signatures in question. The signatures of [15, 18, 2, 25] have all been shown to be stable under perturbations in the Gromov-Wasserstein sense in [19]. More recently, certain persistence topology based signatures have also been shown to be GW stable [8].

Spectral Methods In this work, we aim to extend the ideas of [19] to a different class of techniques for shape classification and shape comparison, called the *spectral methods*. These methods are generally based on constructions

that use the eigenvalues and eigenfunctions of the Laplace-Beltrami operator defined on the shape.

Perhaps the best known spectral invariant was introduced to the shape matching community by Reuter et al. in the remarkable [26], where the authors propose using the a subset of the collection of all eigenvalues (spectrum) of the Laplace-Beltrami operator of a shape as its signature for shape retrieval and comparison. The invariance of the spectrum of Laplace-Beltrami operator to isometric deformations (deformations that leave the geodesic distance unchanged) ensures that this signature can be used to recognize the same shape in different poses. From a theoretical point of view, however, it is not possible to fully classify shapes using this signature, since there exist compact non-isometric shapes whose Laplace-Beltrami operators have the same spectra.

The work of Rustamov [27] is based on the observation that the eigenvalues of the Laplace-Beltrami operator together with the corresponding eigenfunctions characterize the shape up to isometry. The author introduces the Global Point Signature (GPS) of a point on the shape, which encodes both the eigenvalues and the eigenfunctions of the Laplace-Beltrami operator evaluated at that point. The histogram of distances [25] between the signatures of all points is then used for shape comparison.

More recently, Sun et al. [28] and Gebal et al. [12] introduced a robust and multi-scale invariant of a shape based on the *heat kernel*, which arises by solving a certain partial differential equation involving the Laplace-Beltrami operator. Their signature, called the Heat Kernel Signature (or HKS), is defined for every point on the shape and it is also inherently multi-scale. Interestingly, Sun et al. proved that the set of all HKS on the shape “almost always” characterizes it up to isometry, see [28] for details.

Methods based on Heat Diffusion have also been used to analyze graphs (See [1] and references therein), where signatures similar to the HKS have been proven useful in graph classification and comparison.

One of the major challenges of signature-based methods is the choice of metric between signatures. Indeed, defining an appropriate metric between signatures is intimately related to the appropriate notion of distance between shapes themselves. For example, based on the work Reuter et al. [26], one could propose (among other options) using the ℓ^2 metric between the two sets of ordered eigenvalues. This metric, however, will tend to give more weight to larger eigenvalues, which correspond to “high frequency” eigenfunctions. Even more dramatically, in the continuous setting, this metric is not guaranteed to converge when considering the full spectrum of the Laplace-Beltrami operator of the continuous shape.

Following [21, 22], a conceptually different approach was recently adopted by Bronstein et al. [5] who endow

the shapes with the spectrum-based *diffusion distance* (introduced to the applied literature by Lafon in [17]) and then estimate the Gromov-Hausdorff distance between the resulting metric spaces. The primary motivation of [5] is to exploit the apparent stability of diffusion distances to local changes in the topology of the shape. As a result, their method can potentially be used to classify shapes undergoing such changes making it robust in many applications. One of the limitations of [5] is the difficulty of establishing the relation between their distance and different notions of invariance. Indeed, it is unclear whether two objects whose diffusion based GH distance is 0 are necessarily isometric with respect to geodesic distances. Similarly, it is unclear what is the precise notion of similarity encoded by the fact that two shapes (endowed with the diffusion distance) are at small GH distance.

Contributions In this paper, we define a spectral notion of distance between shapes, called the *Spectral Gromov-Wasserstein distance* and formally in Theorem 3.1 show that our definition satisfies the properties of a metric on the collection of all isometry classes of shapes. This means, in particular, that two shapes with 0 distance are necessarily isometric with respect to the geodesic distance. We also encode a *scale parameter* into our definition and argue how our notion of similarity is *foliated* with respect to this parameter.

We also address the question of similarity between shapes by proving, in Theorem 3.2 and Theorem 3.3, a series of *lower bounds* on our metric that involve previously proposed spectral signatures. These lower bounds imply that two shapes such that a suitably chosen distance between their signatures is large, have to be far in terms of our spectral metric. In particular, in Theorem 3.2 we prove that two (interrelated) invariants: the HKS and the *heat trace*, are both stable with respect to the metric we construct. One of the main observations is that the heat trace contains *exactly the same information as the spectrum* of [26]. A second and third set of previously proposed ideas that we address with our construction and relate to are those of [27] and [5]. Again, we exhibit lower bounds for our metric that establish explicit links (via Theorem 3.3 to (suitably reinterpreted version of) those extremely interesting proposals.

At a high level, our construction is based on substituting the heat kernels in the definition of Gromov-Wasserstein distance for the geodesic distances. This is motivated by classical the result by Varadhan (Lemma 2.2) which relates these two quantities: for small $t \simeq 0$, $-4t \ln k_X(t, x, x') \simeq d_X^2(x, x')$. Using the heat kernel, however, has the advantage of directly encoding a scale parameter (t), which allows for multi-scale comparison. We then use *measure couplings*, in the same way as done in [19] to compare heat kernels on different shapes. Thus, our proposed dissimilarity measure (see Definition 3.1) is $\inf_{\mu} \sup_{t>0} c(t)$.

$\|k_X(t, \cdot, \cdot) - k_Y(t, \cdot, \cdot)\|_{L^p(\mu \otimes \mu)}$ where μ is a coupling between the normalized area measures on the shapes X and Y ; $p \geq 1$; and $c(t)$ is a certain function that prevents the blow up of the heat kernels. Note that for a fixed coupling μ we are taking $\sup_{t>0}$ which is to be interpreted as choosing the *most discriminative scale* that tells X apart from Y .

The primary motivations for our work are (1) to provide a formal, continuous (as opposed to discrete) treatment to spectral-based shape comparison, (2) to inter-relate the different invariants introduced so far, and (3) to provide a rigorous framework inside which potential future spectral invariants can be analyzed. The hierarchy of lower bounds that we provide further allows to compute estimates of the spectral distance that we introduce, by increasing accuracy at the cost of an increase in computation time. This is especially useful in shape retrieval where inexpensive shape comparison methods can be used to quickly reject dissimilar shapes.

2. Background

In this section, we review some standard concepts of metric geometry and measure theory that will be used in our presentation. A good reference of the former is [6]. A reference of the latter is [9]. We also review the definitions of the GH and GW distances and related concepts.

Definition 2.1 (Correspondence). *For non-empty sets A and B , a subset $R \subset A \times B$ is a correspondence (between A and B) if and only if (1) $\forall a \in A$ there exists $b \in B$ s.t. $(a, b) \in R$, and (2) $\forall b \in B$ there exists $a \in A$ s.t. $(a, b) \in R$. Let $\mathcal{C}(A, B)$ denote the set of all possible correspondences between sets A and B .*

Definition 2.2. *The support of a Borel measure ν on a metric space (Z, d) , denoted by $\text{supp}[\nu]$, is the minimal closed subset $Z_0 \subset Z$ such that $\nu(Z \setminus Z_0) = 0$.*

Given a metric space (X, d) , a Borel measure ν on X , a function $f : X \rightarrow \mathbb{R}$, and $p \in [1, \infty]$ we denote by $\|f\|_{L^p(\nu)}$ the L^p norm of f w.r.t. the measure ν .

Remark 2.1. *When ν is a probability measure, i.e. $\nu(X) = 1$, $\|f\|_{L^p(\nu)} \geq \|f\|_{L^q(\nu)}$ for $p, q \in [1, \infty]$ and $p \geq q$.*

Definition 2.3. *An isometry between metric spaces (X, d_X) and (Y, d_Y) is any bijective map $\psi : X \rightarrow Y$ s.t. $d_X(x, x') = d_Y(\psi(x), \psi(x'))$ for all $x, x' \in X$. We say that X and Y are isometric whenever there exists an isometry between these spaces.*

We will denote by \mathfrak{R} the collection of all compact and connected Riemannian manifolds without boundary. For $(M, g) \in \mathfrak{R}$, where g is the metric tensor on M , we will denote by vol_M the Riemannian volume measure on M .

Finally, we let ℓ^2 denote the Hilbert space of all square summable sequences and given two such sequences $A = \{a_i\}$ and $B = \{b_i\}$ we denote by $A \bullet_{\ell^2} B = \sum_i a_i b_i$ their inner product.

2.1. Gromov-Wasserstein distances

We now review the main features of the *Gromov-Wasserstein* distance. For metric spaces (X, d_X) and (Y, d_Y) let $\Gamma_{X,Y} : X \times Y \times X \times Y \rightarrow \mathbb{R}^+$ be given by

$$\Gamma_{X,Y}(x, y, x', y') := |d_X(x, x') - d_Y(y, y')|. \quad (1)$$

Recall the definition of the Gromov-Hausdorff distance between (compact) metric spaces X and Y :

$$d_{\mathcal{GH}}(X, Y) = \frac{1}{2} \inf_R \|\Gamma_{X,Y}\|_{L^\infty(R \times R)} \quad (2)$$

where R is a correspondence between X and Y as defined above. This defines a metric on the collection of all isometry classes of metric compact spaces [6, Chapter 7].

As discussed in [19, 20], there is not a unique notion of Gromov-Wasserstein distance – different constructions lead to different properties. In this paper we stick to the construction proposed in [19] which we recall below. Following [19], we view a given shape not just as a set of points with a metric on them, but also assume, in addition, that a probability measure on the (sets of) points is specified. The resulting structure is called an *mm-space*.

Definition 2.4 ([14]). *A metric measure space (mm-space for short) will always be a triple (X, d_X, μ_X) where (a) (X, d_X) is a compact metric space, and (b) μ_X is a Borel probability measure on X i.e. $\mu_X(X) = 1$.*

Two mm-spaces (X, d_X, μ_X) and (Y, d_Y, μ_Y) are called isomorphic (or equal) iff there exists an isometry $\psi : X \rightarrow Y$ such that $\mu_X(\psi^{-1}(B)) = \mu_Y(B)$ for all $B \subset Y$ measurable. We will denote by \mathcal{G}_w the collection of all mm-spaces.

When it is clear from the context, we will denote the triple (X, d_X, μ_X) by only X . In the finite case, $\mu_X(x)$ provides a collection of weights that signal the relative importance or trustworthiness of the point $x \in X$ in relation to the other points on the shape. Moreover, we will assume w.l.o.g. that for all our mm-spaces $X = \text{supp}[\mu_X]$.

Example 2.1 (Riemannian manifolds as mm-spaces). *Let (M, g) be a compact Riemannian manifold. Consider the metric d_M on M induced by the metric tensor g and the normalized measure μ_M , that is, for all measurable $C \subset M$, $\mu_M(C) = \frac{\text{vol}_M(C)}{\text{vol}(M)}$. Then (M, d_M, μ_M) is a mm-space.*

Definition 2.5. *Given two metric measure spaces (X, d_X, μ_X) and (Y, d_Y, μ_Y) we say that a measure μ on the product space $X \times Y$ is a coupling of μ_X and μ_Y iff (1) $\mu(A \times Y) = \mu_X(A)$, and (2) $\mu(X \times A') = \mu_Y(A')$ for all measurable sets $A \subset X$, $A' \subset Y$. We denote by $\mathcal{M}(\mu_X, \mu_Y)$ the set of all couplings of μ_X and μ_Y .*

In words, a coupling of μ_X and μ_Y is a probability measure with marginals μ_X and μ_Y .

Lemma 2.1 ([19]). *Let μ_X and μ_Y be Borel probability measures on compact metric spaces X and Y . If*

$\mu \in \mathcal{M}(\mu_X, \mu_Y)$, then $R(\mu) := \text{supp}[\mu]$ belongs to $\mathcal{C}(\text{supp}[\mu_X], \text{supp}[\mu_Y])$.

Definition 2.6 (Gromov-Wasserstein distance, [19]). For $p \in [1, \infty]$ we define the distance $d_{\mathcal{GW},p}$ between two mm-spaces X and Y by

$$d_{\mathcal{GW},p}(X, Y) := \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \frac{1}{2} \|\Gamma_{X,Y}\|_{L^p(\mu \otimes \mu)}. \quad (3)$$

Example 2.2. For finite mm-spaces X and Y , $\mathcal{M}(\mu_X, \mu_Y)$ can be regarded as the set of all matrices $M = ((m_{ij}))$ with non-negative elements such that $\sum_i m_{ij} = \mu_Y(j)$ and $\sum_j m_{ij} = \mu_X(i)$ for all $i = 1, \dots, \#X$ and $j = 1, \dots, \#Y$. Therefore, for finite $p \geq 1$,

$$d_{\mathcal{GW},p}(X, Y) := \frac{1}{2} \left(\inf_M \sum_{i,j} \sum_{i',j'} (\Gamma_{X,Y}(x_i, y_j, x_{i'}, y_{j'}))^p m_{ij} m_{i'j'} \right)^{1/p},$$

which leads to a quadratic optimization problem with continuous variables and linear constraints.

The GW distance satisfies many useful properties, see [19] for a more complete list.

Theorem 2.1 ([19]). Let $p \in [1, \infty]$, then

- (a) $d_{\mathcal{GW},p}$ defines a metric on the set of all (isomorphism classes of) mm-spaces.
- (b) For any two mm-spaces X and Y , $d_{\mathcal{GW},p}(X, Y) \geq d_{\mathcal{GW},q}(X, Y)$ whenever $\infty \geq p \geq q \geq 1$.

2.2. Heat kernels on compact manifolds.

Our primary goal is to extend the Gromov-Wasserstein distance to the spectral setting. The definition we give in the following section is based on the fundamental solution of the heat equation, also known as the *heat kernel*. In this section we introduce the heat kernels on compact manifolds and list some of their key properties. For detailed exposition of the material presented here, we refer the reader to excellent survey [13].

Let X be a compact Riemannian manifold, the Laplace-Beltrami operator Δ_X is a generalisation of the Laplacian to non-euclidean domains and maps functions defined on X to other such functions. If X is compact then Δ_X has a discrete set of eigenvalues λ_i (called the *spectrum*) and eigenfunctions ζ_i , which can be chosen to define an orthonormal basis to the set of square-integrable functions on X —we make this assumption henceforth.

Remark 2.2 (Weyl’s formula). One has the following asymptotic formula [7] for the eigenvalues of the Laplace-Beltrami operator on a compact and connected d -dimensional Riemannian manifold M :

$$\lambda_\ell \sim 4\pi^2 \left(\frac{\ell}{\omega_d \text{Vol}(M)} \right)^{2/d} \quad \text{as } \ell \rightarrow \infty.$$

Remark 2.3 (Scaling). Given a d -dimensional Riemannian manifold (M, g) with eigenvalues λ_i and eigenfunctions ζ_i , and $a > 0$, the eigenvalues and eigenfunctions of $(M, a^2 g)$ are λ_i/a^2 and $\zeta_i/a^{d/2}$.

Remark 2.4. As mentioned in the introduction, based on Reuter et al. [26] one could consider the spectrum of the LB operator on a manifold as its signature and compare signatures of two manifolds by computing the ℓ^2 norm of the difference between the two sets of ordered eigenvalues. Although the practical associated signature (which uses cropped versions of the spectra) shows very good discrimination power, the previous remarks tells us that the resulting theoretical notion of distance would not be well defined when comparing, for example, (M, g) with $(M, a^{-2} g)$ for some $a \neq 1$.¹ Therefore, as was observed in [26], in practice, one has to be careful to only consider a subset of eigenvalues or by doing appropriate scaling.

The Heat Operator on X is defined as: $H_t = e^{-t\Delta_X}$, where Δ_X is the Laplace-Beltrami operator on X . The operators H_t and Δ_X have the same eigenfunctions and if λ is an eigenvalue of Δ_X , then $e^{-\lambda t}$ is an eigenvalue of H_t corresponding to the same eigenfunction. Similarly to the Laplace-Beltrami operator, the Heat Operator is invariant under isometric changes to the manifold. It is well-known [7] that for any $X \in \mathfrak{X}$, there exists a continuous function $k_X : \mathbb{R}^+ \times X \times X \rightarrow \mathbb{R}$ such that $H_t f(x) = \int_X k_X(t, x, x') f(x') \text{vol}_X(dx')$ where $\text{vol}_X(dx')$ is the volume form at $x' \in X$. The minimal function k_X that satisfies the equation above, is called the **heat kernel** on X . For compact X , the heat kernel has the following eigen-decomposition:

$$k_X(t, x, x') = \sum_{i=0}^{\infty} e^{-\lambda_i t} \zeta_i(x) \zeta_i(x'), \quad (4)$$

where λ_i is the i^{th} eigenvalue (each counted the number of times equal to its multiplicity) and ζ_i is the corresponding eigenfunction of the Laplace-Beltrami operator Δ_X .

Example 2.3. For all $x, x' \in \mathbb{R}^d$ and $t > 0$, the heat kernel on \mathbb{R}^d is given by $k_{\mathbb{R}^d}(t, x, x') = \frac{1}{(4\pi t)^{d/2}} \exp\left(-\frac{\|x-x'\|^2}{4t}\right)$.

The heat kernel is also related to the geodesic distances on the manifold for small values of t :

Lemma 2.2 ([24]). For any $X \in \mathfrak{X}$,

$$\lim_{t \downarrow 0} (-4t \ln k_X(t, x, x')) = d_X^2(x, x'),$$

for all $x, x' \in X$. Here $d_X(x, x')$ is the geodesic distance between x and x' on X .

For compact manifolds, the long-term behavior of the heat-kernel is given explicitly: $\lim_{t \rightarrow \infty} k_X(t, x, x') =$

¹In this case, $\|\Lambda - \Lambda'\|_{\ell^2} = |a^2 - 1| \cdot \|\Lambda\|_{\ell^2}$, where $\Lambda = \{\lambda_i\}_{i=0}^{\infty}$, $\Lambda' = \{a^2 \lambda_i\}$, and λ_i is the i th eigenvalue of (M, g) . Now, Weyl’s asymptotic expansion guarantees that $\|\Lambda\|_{\ell^2} = \infty$ and hence the proposed distance between spectra is infinity as well.

$\frac{1}{\mathbf{Vol}(X)}$. In other words, as t goes to infinity, the heat distribution on X converges to a constant,² regardless of the initial distribution.

Heat Kernel Signature and Heat Trace

From Lemma 2.2 it is clear that two manifolds X and Y are isometric if and only if there exists a surjective map $\phi : X \rightarrow Y$ such that $k_X(t, x, x') = k_Y(t, \phi(x), \phi(x'))$ for all $x, x' \in X$ and $t \in \mathbb{R}^+$. This means that the set of all heat kernel functions uniquely characterizes the shape up to isometry. However, in practice, this set can be rather large. The Heat Kernel Signature, recently introduced to computer graphics by Sun et al. [28] is defined as a restriction of the heat kernel to the diagonal times the volume of the shape:³

$$\mathbf{hks}_X : X \times \mathbb{R}^+ \rightarrow \mathbb{R}^+, \quad (x, t) \mapsto \mathbf{Vol}(X) \cdot k_X(t, x, x).$$

It is a well known that the HKS of a point x is related to the scalar curvature $s_X(x)$ for small values of t [23]: $\mathbf{hks}_X(t, x) \sim (4\pi t)^{-d/2} \sum_{i=0}^{\infty} a_i t^i$ as $t \rightarrow 0$ where d is the dimension of X , $a_0 = \mathbf{Vol}(X)$ and $a_1 = \mathbf{Vol}(X) \cdot \frac{1}{6} s_X(x)$.

For a given $X \in \mathfrak{X}$ let $\mathcal{H}_X : \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow [0, 1]$ be defined by $\mathcal{H}_X(t, s) = \mu_X\{x \in X \mid \mathbf{hks}_X(x, t) \leq s\}$. In other words, for a fixed $t > 0$ and $s > 0$, $\mathcal{H}_X(t, s)$ gives the (normalized) area of the set of points on X whose HKS at scale t is below the threshold s . Here, μ_X denotes the *normalized area measure* (recall Example 2.1).

One can further aggregate (average, actually) Heat Kernel Signatures at all points on the manifold to obtain the Heat Trace:

$$K_X(t) := \frac{1}{\mathbf{Vol}(X)} \int_X \mathbf{hks}_X(t, x) \mathbf{vol}_X(dx) := \sum_{i=0}^{\infty} e^{-\lambda_i t}, \quad (5)$$

where λ_i is, again, the i^{th} eigenvalue of Δ_X . Similarly to the Heat Kernel Signature, the heat trace contains a lot of geometric information about the manifold, as can be seen, from the following well known expansion: $K_X(t) \sim (4\pi t)^{-d/2} \sum_{i=0}^{\infty} u_i t^i$ as $t \rightarrow 0$, where $u_0 = \mathbf{Vol}(X)$, $u_1 = \frac{1}{6} \int_X s(x) \mathbf{vol}_X(dx)$. Note that for 2-dimensional manifolds (surfaces) without boundary, $u_1 = \frac{1}{3} \pi \chi(X)$ by the Gauss-Bonnet theorem, where $\chi(X)$ is the Euler characteristic of X . See [26] for an application of these expansions to shape analysis.

Remark 2.5 (The heat trace and the shape DNA). *It is well known and easy to show that the spectrum $\{\lambda_i\}$ can be deduced from the heat trace $K_X(t) : \mathbb{R}^+ \rightarrow \mathbb{R}$. Indeed, $\lambda_0 = \inf\{a > 0 \text{ s.t. } \lim_{t \rightarrow \infty} e^{at} K_X(t) \neq 0\}$ and the multiplicity of λ_0 , $N_0 = \lim_{t \rightarrow \infty} e^{\lambda_0 t} K_X(t)$. Both λ_1 and its multiplicity can be obtained from: $K'_X(t) = K_X(t) - N_0 e^{-t\lambda_0}$, and this process can be iterated to the whole spectrum.*

²Or alternatively, $k_X(t, x, x')$ converges to the uniform density.

³We are slightly modifying the definition in [28] up to a multiplicative factor. This invariant is a very well known subject in the literature dealing with the heat kernel [13].

Thus, knowledge of the Heat Trace is **equivalent** to that of the spectrum, which suggests a way to formally analyze the Shape DNA of [26] from the point of view of $K_X(t)$.

Diffusion distances and Rustamov's invariant

In 1994 Berard et al [4] introduced the idea of embedding a Riemannian manifold into a Banach space via a spectral type of embedding that uses the heat kernel. This idea is deeply connected to the proposal of diffusion distances introduced in the applied literature in [17]. In more detail, consider $C_0([0, \infty], \ell^2)$ endowed with the norm $\Theta(\gamma, \sigma) := \sup\{\|\gamma(t) - \sigma(t)\|_{\ell^2}, t \in [0, \infty]\}$. For each $t > 0$, let $c(t) := e^{-t^{-1}}$. Consider now the map from $X \in \mathfrak{X}$ (with orthonormal base of eigenfunctions of the Laplace-Beltrami operator $\{\zeta_i\}$ and eigenfunctions $\{\lambda_i\}$) to $C_0([0, \infty], \ell^2)$ that assigns to $x \in X$ the element

$$I_X[x](t) = c(t) \left\{ e^{-\frac{\lambda_i}{2} t} \zeta_i(x) \right\}; \quad t \in [0, \infty]. \quad (6)$$

It can be seen that this embedding is continuous, and it is clear that

$$\Theta(I_X[x], I_X[x']) = \sup_{t>0} c(t) \cdot d_{X;t}^{\text{spec}}(x, x')$$

where $d_{X;t}^{\text{spec}}(x, x') := (k_X(t, x, x) + k_X(t, x', x') - 2k_X(t, x, x'))^{1/2}$ will be called the *diffusion distance on X at scale t* (which coincides with the definition of the diffusion distance given in [17]). Note that the use of $c(t)$ here avoids the blow up as $t \downarrow 0$ of the usual definition of the diffusion distance introduced in [17].⁴

We consider another useful invariant of any $X \in \mathfrak{X}$: for each $t > 0$ define $\mathcal{G}_X(\cdot, t) : [0, \infty) \rightarrow [0, 1]$ to be total (normalized) mass of pairs of points (x, x') in $X \times X$ s.t. $d_{X;t}^{\text{spec}}(x, x') \leq s$. Of course this is mathematical language for the histogram of inter-point distances between pairs of points. Precisely,

$$\mathcal{G}_X(s, t) = (\mu_X \otimes \mu_X) \left((x, x') \mid d_{X;t}^{\text{spec}}(x, x') \leq s \right).$$

Remark 2.6. *In [5], the authors propose to compute the GH distance between shapes X, Y which are endowed with the diffusion metric at a fixed scale t . We show that if instead one computes the GW distance between such shapes, one obtains a lower bound for the spectral notion of metric between shapes that we construct in this paper.*

Remark 2.7 (The GPS embedding of Rustamov). *Similarly to the spectral embedding of Berard et al, Rustamov [27] proposes embedding $X \in \mathfrak{X}$ into ℓ^2 via the map $R_X : X \rightarrow \ell^2$ defined by*

$$X \ni x \mapsto \left\{ \frac{1}{\sqrt{\lambda_i}} \zeta_i(x) \right\}$$

His shape matching proposal is to consider a certain version of the shape distributions idea [25] in the embedded space. Namely, for a given shape X , he proposes to

⁴Roughly speaking, for a d -dimensional Riemannian manifold X , the heat kernel behaves like $t^{-d/2}$ for $t \simeq 0$ (and therefore it blows up). Hence $c(t) = e^{-t^{-1}}$ is 'strong' enough to keep $c(t) \cdot k_X(t, \cdot, \cdot)$ bounded as t approaches 0 for any $d \in \mathbb{N}$.

compute the histogram of distances $\|R_X[x] - R_X[x']\|_{\ell^2}$, $x, x' \in X$. We will show below, that a certain reformulation of the GPS+Shape distributions procedure of Rustamov can be expressed using the invariant \mathcal{G}_X defined above, and that this yields is a lower bound for the spectral notion of distance we construct in this paper. The proposed reinterpretation is to look at $I_X[x]$ instead of $R_X(x)$ and to use \mathcal{G}_X as a proxy for the shape distributions invariant he proposed to use. Note that this is correct since by definition of I_X , $\|I_X[x](t) - I_X[x'](t)\|_{\ell^2} = c(t) \cdot d_{X;t}^{\text{spec}}(x, x')$. Also, note that for all $t > 0$ and $x, x' \in X$, by (4) we have $I_X[x](t) \bullet_{\ell^2} I_X[x'](t) = (c(t))^2 \cdot k_X(t, x, x')$ which can be compared to Rustamov's motivation for the embedding he proposed: namely, the fact that $R_X[x] \bullet_{\ell^2} R_X[x'] = G_X(x, x')$ for all $x, x' \in X$ where G_X is the Green function on X , see [27, Section 4].

2.3. The notion of scale

As mentioned earlier, the time parameter t in the heat kernel can be naturally interpreted as a certain notion of scale. For example, the Heat Kernel Signature of a point x reflects differential properties of the surface at x (such as curvature) for small value of t , whereas $\lim_{t \rightarrow \infty} k_X(t, x, x) = \frac{1}{\text{Vol}(X)}$ independent of x .

Here we provide a point of view based on homogenization of partial differential equations [3] which allows to interpret t as scale. Consider the real line with metric given by the periodic C^2 function $g : \mathbb{R} \rightarrow \mathbb{R}^+$ with period 1 such that $\gamma^{-1} \leq g(x) \leq \gamma$ for some fixed $\gamma > 0$. One can regard $M_g = (\mathbb{R}, g)$ as a weighted Riemannian manifold [13] under the standard Lebesgue measure. The geodesic distance on M_g admits an explicit expression: $d_g(x, x') = \int_x^{x'} g^{1/2}(s) ds$, $x \geq x'$. For any metric g satisfying the conditions above, one can consider the resulting Laplace-Beltrami operator on M_g to be $\Delta_g = \frac{d}{dx} \left(\frac{1}{g(x)} \frac{d}{dx} \right)$. Let k_g denote the heat kernel associated to M_g . Note in particular that when $g = g_0 > 0$ constant, $d_{g_0}(x, x') = (g_0)^{1/2} |x - x'|$, and $k_{g_0}(t, x, x') = \sqrt{\frac{g_0}{4\pi t}} e^{-g_0 \frac{(x-x')^2}{4t}}$, $x, x' \in \mathbb{R}$.

One would expect that for $t \rightarrow \infty$, the heat kernel $k_g(t, \cdot, \cdot)$ looks like the heat kernel $k_{\bar{g}}(t, \cdot, \cdot)$ corresponding to a certain constant metric \bar{g} . This would be in agreement with the intuition that large values of t offer a coarse scale view of the underlying metric structure. This intuition can be made precise alluding to a result due to Tsuchida:

Theorem 2.2 ([30]). *There exists a positive constant C such that $\sup_{x, x' \in \mathbb{R}} |k_g(t, x, x') - k_{\bar{g}}(t, x, x')| \leq \frac{C}{t}$, for all $t > 0$, where $\bar{g} = \int_0^1 g(x) dx$.*

Example 2.4. *Pick $0 < \varepsilon < 1$ and $m \in \mathbb{N}$ and let $g(x) := 1 + \varepsilon \cdot \sin(2\pi m \cdot x)$. Then $\bar{g} = 1$. Tsuchida's theorem then guarantees that as t approaches infinity, $|k_g(t, x, x') - \frac{1}{\sqrt{4\pi t}} \exp\{-(x - x')^2/4t\}| < \frac{C}{t}$, for all $x, x' \in \mathbb{R}$, that is, k_g looks like the heat kernel corresponding to a flat one dimensional profile.*

Using this intuition, we can argue that the parameter t can be interpreted as a notion of scale. This means that intuitively, for a fixed x , $k_X(t, x, \cdot)$ is a function that reflects the the properties of a geodesic neighborhood of x . As t grows, the so does the size of the neighborhood, but the information becomes more and more smoothed out.

3. Spectral Gromov Wasserstein metric

In this section we carry out an adaptation of the Gromov-Wasserstein distance to the class of Riemannian manifolds. A similar construction is possible for the Gromov-Hausdorff distance.

Our construction is similar to the one proposed by Kasue and Kumura [16], but different in that we utilize the formalism of mm-spaces and coupling measures. This allows us to obtain explicit bounds that encode several interesting invariants, see §3.1.

The heat kernel provides a naturally multi-scale decomposition of the geometry of a Riemannian manifold. We want to make this appear explicitly in the definition of a metric on the collection of isometry classes of \mathfrak{R} . For $X, Y \in \mathfrak{R}$ and $t \in \mathbb{R}^+$ define $\Gamma_{X, Y, t}^{\text{spec}} : \mathbb{R}^+ \times X \times Y \times X \times Y \rightarrow \mathbb{R}^+$ by

$$(t, x, y, x', y') \mapsto |\text{Vol}(X) \cdot k_X(t, x, x') - \text{Vol}(Y) \cdot k_Y(t, y, y')|. \quad (7)$$

Definition 3.1. *For $X, Y \in \mathfrak{R}$ and $p \in [1, \infty]$ let*

$$d_{\mathcal{GW}, p}^{\text{spec}}(X, Y) := \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \sup_{t > 0} c(t) \cdot \|\Gamma_{X, Y, t}^{\text{spec}}\|_{L^p(\mu \otimes \mu)},$$

where $c(t) = e^{-at^{-1}}$ for some $a > 0$.⁵

Remark 3.1 (Multi-scale aspect of the definition). *Our definition exploits the scale parameter t of the heat kernel in order to define a spectral version of the GW distance. This means that two Riemannian manifolds will be considered to be similar in the spectral GW sense if and only if they are similar at all scales t . This is encoded in the definition of our metric by first taking the supremum over all $t > 0$, and then choosing the best coupling. The use of the parameter t in this fashion provides a natural **foliation** of the notion of approximate isometry between Riemannian manifolds.*

Theorem 3.1. *For all $p \in [1, \infty]$, $d_{\mathcal{GW}, p}^{\text{spec}}(\cdot, \cdot)$ defines a metric on the collection of all isometry classes of \mathfrak{R} . Moreover, for any $X, Y \in \mathfrak{R}$, $d_{\mathcal{GW}, p}^{\text{spec}}(X, Y) \geq d_{\mathcal{GW}, q}^{\text{spec}}(X, Y)$ for all $1 \leq q \leq p \leq \infty$.*

Proof. The proof of the triangle inequality is easy. Let $X, Y, Z \in \mathfrak{R}$ be s.t. $d_{\mathcal{GW}, p}^{\text{spec}}(X, Y) < \varepsilon_1$ and $d_{\mathcal{GW}, p}^{\text{spec}}(Y, Z) < \varepsilon_2$. Let $\mu_1 \in \mathcal{M}(\mu_X, \mu_Y)$ and $\mu_2 \in \mathcal{M}(\mu_Y, \mu_Z)$ be s.t. $c(t) \cdot \|\Gamma_{X, Y, t}^{\text{spec}}\|_{L^p(\mu_1 \otimes \mu_1)} < \varepsilon_1$ and $c(t) \cdot \|\Gamma_{Y, Z, t}^{\text{spec}}\|_{L^p(\mu_2 \otimes \mu_2)} < \varepsilon_2$ for all $t \in \mathbb{R}^+$. For fixed $t \in \mathbb{R}^+$, by the triangle inequality for the absolute value:

$$\Gamma_{X, Y, t}^{\text{spec}}(x, y, x', y') \leq \Gamma_{X, Z, t}^{\text{spec}}(x, z, x', z') + \Gamma_{Z, Y, t}^{\text{spec}}(z, y, z', y'), \quad (8)$$

⁵Compare with Definition 2.6. The reason for the use of $c(t)$ is to prevent the blow up of the quantities involved as $t \downarrow 0$.

for all $x, x' \in X, y, y' \in Y$ and $z, z' \in Z$. Now, by the Gluing Lemma [31, Lemma 7.6], there exists a probability measure $\mu \in \mathcal{P}(X \times Y \times Z)$ with marginals μ_1 on $X \times Z$ and μ_2 on $Z \times Y$. Let μ_3 be the marginal of μ on $X \times Y$. Using the fact that μ has marginal $\mu_Z \in \mathcal{P}(Z)$ on Z and the triangle inequality for the L^p norm (i.e Minkowski's inequality) and (8), we obtain

$$\begin{aligned} \|\Gamma_{X,Y,t}\|_{L^p(\mu_3 \otimes \mu_3)} &= \|\Gamma_{X,Y,t}\|_{L^p(\mu \otimes \mu)} \\ &\leq \|\Gamma_{X,Z,t} + \Gamma_{Z,Y,t}\|_{L^p(\mu \otimes \mu)} \\ &\leq \|\Gamma_{X,Z,t}\|_{L^p(\mu \otimes \mu)} + \|\Gamma_{Z,Y,t}\|_{L^p(\mu \otimes \mu)} \\ &= \|\Gamma_{X,Z,t}\|_{L^p(\mu_1 \otimes \mu_1)} + \|\Gamma_{Z,Y,t}\|_{L^p(\mu_2 \otimes \mu_2)} \\ &\leq (\varepsilon_1 + \varepsilon_2)/c(t). \end{aligned}$$

Hence $d_{\mathcal{GW},p}^{\text{spec}}(X, Y) \leq \sup_{t>0} c(t) \cdot \|\Gamma_{X,Y,t}\|_{L^p(\mu_3 \otimes \mu_3)} < \varepsilon_1 + \varepsilon_2$. The conclusion follows now by taking $\varepsilon_1 \rightarrow d_{\mathcal{GW},p}^{\text{spec}}(X, Y)$ and $\varepsilon_2 \rightarrow d_{\mathcal{GW},p}^{\text{spec}}(Y, Z)$.

We prove that $d_{\mathcal{GW},p}^{\text{spec}}(X, Y) = 0$ implies that X and Y are isometric. Assume first that $p \in [1, \infty)$. Let $(\varepsilon_n) \subset \mathbb{R}^+$ be s.t. $\varepsilon_n \rightarrow 0$ and $(\mu_n) \subset \mathcal{M}(\mu_X, \mu_Y)$ be s.t.

$$\|\Gamma_{X,Y,t}^{\text{spec}}\|_{L^p(\mu_n \otimes \mu_n)} < \varepsilon_n/c(t) \quad \text{for all } n \in \mathbb{N} \text{ and } t \in \mathbb{R}^+. \quad (9)$$

Since $\mathcal{M}(\mu_X, \mu_Y)$ is compact for the weak topology on $\mathcal{P}(X \times Y)$ (see [31, pp. 49]), we can assume that up to extraction of a sub-sequence, μ_n converges to some $\mu_0 \in \mathcal{M}(\mu_X, \mu_Y)$. We assume that $\mu_n \rightarrow \mu_0$ weakly. Then, $\mu_n \otimes \mu_n \rightarrow \mu_0 \otimes \mu_0$ weakly as well. Since for fixed $t \in \mathbb{R}^+$, $\Gamma_{X,Y,t}^{\text{spec}}$ is continuous on $X \times Y \times X \times Y$ and hence bounded (since we are considering only compact manifolds), one has that

$$\iint_{X \times Y} \iint_{X \times Y} (\Gamma_{X,Y,t}^{\text{spec}})^p d\mu_n \otimes \mu_n \rightarrow \iint_{X \times Y} \iint_{X \times Y} (\Gamma_{X,Y,t}^{\text{spec}})^p d\mu_0 \otimes \mu_0$$

as $n \uparrow \infty$. By (9) we obtain that for all $t > 0$ $\|\Gamma_{X,Y,t}^{\text{spec}}\|_{L^p(\mu_0 \otimes \mu_0)} = 0$. It follows that $\Gamma_{X,Y,t}^{\text{spec}}(x, y, x', y') = 0$ for all $(x, y), (x', y') \in R(\mu_0)$ which is equivalent to

$$\mathbf{Vol}(X) \cdot k_X(t, x, x') = \mathbf{Vol}(Y) \cdot k_Y(t, y, y'),$$

for all $(x, y), (x', y') \in R(\mu_0)$. By Lemma 2.1, $R(\mu_0) \in \mathcal{C}(X, Y)$. Consider a map $\phi : X \rightarrow Y$ s.t. $(x, \phi(x)) \in R(\mu_0)$ for all $x \in X$. Then by the above we find that $\mathbf{Vol}(X) \cdot k_X(t, x, x') = \mathbf{Vol}(Y) \cdot k_Y(t, \phi(x), \phi(x'))$ for all $t > 0$ and $x, x' \in X$. By Lemma 2.2, it follows then that $d_X(x, x') = d_Y(\phi(x), \phi(x'))$ for all $x, x' \in X$. Similarly, we can find a map $\psi : Y \rightarrow X$ s.t. $d_Y(y, y') = d_X(\psi(y), \psi(y'))$ for all $y, y' \in Y$. It follows that $\zeta := \phi \circ \psi$ is an isometry from Y into itself and since Y is compact, ζ has to be surjective. It follows that ϕ (and also ψ) is an isometry.

Now, for the case $p = \infty$, pick $(\varepsilon_n) \subset \mathbb{R}^+$ be s.t. $\varepsilon_n \rightarrow 0$ and let $(\mu_n) \subset \mathcal{M}(\mu_X, \mu_Y)$ be s.t. $\|\Gamma_{X,Y,t}^{\text{spec}}\|_{L^\infty(\mu_n \otimes \mu_n)} < \varepsilon_n/c(t)$ for all $n \in \mathbb{N}$ and $t \in \mathbb{R}^+$. Then, by Remark 2.1, (9) holds as well for finite p and the argument above applies. \square

3.1. Lower bounds

In the two theorems below, we establish **two hierarchies** of different lower bounds for the spectral GW distance.

Theorem 3.2. For all $X, Y \in \mathfrak{R}$ and $p \geq 1$,

$$\begin{aligned} d_{\mathcal{GW},\infty}^{\text{spec}}(X, Y) &\stackrel{(A)}{\geq} \\ \sup_{t>0} c(t) \cdot \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \|hks_X(\cdot, t) - hks_Y(\cdot, t)\|_{L^p(\mu)} &\stackrel{(B)}{\geq} \\ \sup_{t>0} c(t) \cdot \int_0^\infty |\mathcal{H}_X(t, s) - \mathcal{H}_Y(t, s)| ds &\stackrel{(C)}{\geq} \\ \sup_{t>0} c(t) \cdot |K_X(t) - K_Y(t)|. & \end{aligned}$$

Remark 3.2. Observe that in the context of standard GH/GW distances, the same bound as in (A) would be trivial since the restriction of $\Gamma_{X,Y}$ to $\{(x, y, x, y), x \in X, y \in Y\}$ is 0.

Remark 3.3. Note that in the lower bounds above the order of the $\sup_{t>0}$ and the \inf_μ are inverted with respect to the order that appears in the definition of $d_{\mathcal{GW},p}^{\text{spec}}(\cdot, \cdot)$. This will allow us to obtain lower bounds at different scales and then consider the most discriminative scale. It is obvious that we can take the sup over a smaller, possibly finite, collection of interesting and/or computable scales T and we will still obtain a lower bound for $d_{\mathcal{GW},p}^{\text{spec}}(\cdot, \cdot)$.

Remark 3.4 (About bound (C) and Shape-DNA). Since knowledge of the heat trace is equivalent to knowledge of the spectrum, lower bound (C) can be interpreted as a version of the Shape DNA signature of Reuter et al. [26] that is compatible with the spectral GW distance. This was one of the goals of our project.

Remark 3.5. The question of the quality of the discrimination provided by the heat trace is of course very important from both the theoretical and the practical points of view. It is known that there exist isospectral Riemannian manifolds that are not isometric. Examples of these constructions are the spheres of Szabo [29]. An interesting theoretical problem is that of finding non-isometric $X, Y \in \mathfrak{R}$ s.t. they have (1) the same HKS, (2) they have the same distribution of HKSs (but different HKSs), and (3) have the same heat traces (but different distributions of HKSs).

Remark 3.6. Notice that in practical applications, computing the lower bound given by Theorem 3.2 (A) involves solving a Linear Optimization Problem with $n_X \times n_Y$ variables (and $n_X + n_Y$ constraints) where n_X (resp. n_Y) is the number of vertices in X (resp. Y). This may be expensive for large models. Therefore, lower bounds (C) which is based on the distribution functions associated to the heat kernel signature seems more suitable in practice.

Theorem 3.3. For all $X, Y \in \mathfrak{R}$,

$$\begin{aligned} d_{\mathcal{GW},\infty}^{\text{spec}}(X, Y) &\stackrel{(A')}{\geq} \\ \sup_{t>0} c(t) \cdot \left(d_{\mathcal{GW},\infty}(X_t, Y_t)\right)^2 &\stackrel{(B')}{\geq} \\ \sup_{t>0} c(t) \cdot \left(\frac{1}{2} \int_0^\infty |\mathcal{G}_X(s, t) - \mathcal{G}_Y(s, t)| ds\right)^2 & \end{aligned}$$

where $X_t = (X, d_{X;t}^{\text{spec}}, \mu_X)$ and $Y_t = (Y, d_{Y;t}^{\text{spec}}, \mu_Y)$.

Remark 3.7. Note that the lower bound (A') establishes a link with the proposal of [5] whereas lower bound (B') embodies the computation of a procedure similar to the one proposed by Rustamov, see Remark 2.7. In particular, it follows that for any $t > 0$, $\frac{1}{c(t)}d_{\mathcal{GW},\infty}^{spec}(X, Y) \geq (d_{\mathcal{GW},\infty}(X_t, Y_t))^2$. It is not known whether for a fixed $t > 0$, $d_{\mathcal{GW},\infty}(X_t, Y_t) = 0$ implies that X and Y are isometric. Furthermore, it would be interesting to investigate the properties of the RHS of (A') as to whether it provides a (pseudo) metric on \mathfrak{R} and whether it is equivalent to $d_{\mathcal{GW},\infty}^{spec}(\cdot, \cdot)$. Also, it is of interest to understand the relationship between the two hierarchies of lower bounds we have established.

4. Discussion

Our proposal hinges on a specialization of the original Gromov-Wasserstein notion of distance between mm-spaces. Which we call **spectral Gromov-Wasserstein distance**. This distance incorporates spectral information directly via the use of heat kernels. At this stage we are able to prove that our proposed metric satisfies all the properties of a metric on the collection \mathfrak{R} of compact Riemannian manifolds without boundary, and the extension to a larger class of shapes will be subject of future efforts. The fact that the GW-spectral metric encodes scale in a natural way makes it suitable for multi-scale matching of shapes. Besides the lower bounds we have presented, others are possible and in the future we plan to tackle this. The practical performance of our ideas remains to be tested.

5. Acknowledgements

We thank Maks Ovsjanikov and Jian Sun for their input and feedback about this work. We would like to acknowledge support for this project by ONR grant N00014-09-1-0783.

References

- [1] X. Bai. *Heat Kernel Analysis on Graphs*. PhD thesis, The University of York, 2007. 2
- [2] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(4):509–522, April 2002. 1
- [3] A. Bensoussan, J.-L. Lions, and G. Papanicolaou. *Asymptotic analysis for periodic structures*, volume 5 of *Studies in Mathematics and its Applications*. North-Holland Publishing Co., Amsterdam, 1978. 6
- [4] P. Berard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994. 5
- [5] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *Submitted to the Intl. Journal of Computer Vision (IJCV)*, 2009. 2, 5, 8
- [6] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry*, volume 33 of *AMS Graduate Studies in Math*. American Mathematical Society, 2001. 3
- [7] I. Chavel. *Eigenvalues in Riemannian Geometry*. Academic Press, 1984. 4
- [8] F. Chazal, D. Cohen-Steiner, L. Guibas, F. Mémoli, and S. Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Submitted to SGP*, 2009. 1
- [9] R. M. Dudley. *Real analysis and probability*, volume 74 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2002. Revised reprint of the 1989 original. 3
- [10] A. E. (Elbaz) and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(10):1285–1295, 2003. 1
- [11] P. Frosini. A distance for similarity classes of submanifolds of Euclidean space. *Bull. Austral. Math. Soc.*, 42:3:407–416, 1990. 1
- [12] K. Gebal, A. Baerentzen, H. Aans, and R. Larsen. Shape analysis using the auto diffusion function. In *SGP*, 2009. 2
- [13] A. Grigor'yan. Heat kernels on weighted manifolds and applications. *Cont. Math.*, 398:93–191, 2006. 4, 5, 6
- [14] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*, volume 152 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1999. 1, 3
- [15] A. B. Hamza and H. Krim. Geodesic object representation and recognition. In *Lecture Notes in Computer Science*, volume 2886, p. 378–387. Springer, 2003. 1
- [16] A. Kasue and H. Kumura. Spectral convergence of Riemannian manifolds. *Tohoku Math. J.*, 46:147–179, 1994. 6
- [17] S. Lafon. *Diffusion Maps and Geometric Harmonics*. PhD thesis, Yale University, 2004. 2, 5
- [18] S. Manay, D. Cremers, B. Hong, A. Yezzi, and S. Soatto. Integral invariants for shape matching. *PAMI*, 28(10):1602–1618, October 2006. 1
- [19] F. Mémoli. On the use of Gromov-Hausdorff distances for shape comparison. In *Proceedings of Point Based Graphics 2007*, Prague, Czech Republic, 2007. 1, 2, 3, 4
- [20] F. Mémoli. Gromov-Hausdorff distances in Euclidean spaces. *Computer Vision and Pattern Recognition Workshops, 2008. CVPR Workshops 2008. IEEE Computer Society Conference on*, p. 1–8, June 2008. 1, 3
- [21] F. Mémoli and G. Sapiro. Comparing point clouds. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, p. 32–40, New York, NY, USA, 2004. ACM. 1, 2
- [22] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005. 1, 2
- [23] S. Minakshisundaram and A. Pleijel. Some properties of the eigenfunctions of the laplace operator on riemannian manifolds. *Can. J. Math.*, 1949. 5
- [24] J. R. Norris. Heat kernel asymptotics and the distance function in Lipschitz Riemannian manifolds. *Acta Math.*, 179(1):79–103, 1997. 4
- [25] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM Trans. Graph.*, 21(4):807–832, 2002. 1, 2, 5
- [26] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-Beltrami spectra as "shape-dna" of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006. 1, 2, 4, 5, 7
- [27] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Symposium on Geometry Processing*, p. 225–233, 2007. 2, 5, 6
- [28] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *SGP*, 2009. 2, 5
- [29] Z. I. Szabó. A cornucopia of isospectral pairs of metrics on spheres with different local geometries. *Ann. of Math. (2)*, 161(1):343–395, 2005. 7
- [30] T. Tsuchida. Long-time asymptotics of heat kernels for one-dimensional elliptic operators with periodic coefficients. *Proc. Lond. Math. Soc. (3)*, 97(2):450–476, 2008. 6
- [31] C. Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003. 7