

On the use of Gromov-Hausdorff Distances for Shape Comparison

Facundo Mémoli^{1†}

¹Department of Mathematics, Stanford University, California, USA.

Abstract

It is the purpose of this paper to propose and discuss certain modifications of the ideas concerning Gromov-Hausdorff distances in order to tackle the problems of shape matching and comparison. These reformulations render these distances more amenable to practical computations without sacrificing theoretical underpinnings. A second goal of this paper is to establish links to several other practical methods proposed in the literature for comparing/matching shapes in precise terms. Connections with the Quadratic Assignment Problem (QAP) are also established, and computational examples are presented.

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

1. Introduction

Given the great advances in recent years in the fields of shape acquisition and modelling, and the resulting huge collections of digital models that have been obtained it is of great importance to be able to define and compute meaningful notions of similarity between shapes which exhibit invariance to different deformations and or poses of the objects represented by those shapes. It is the case that similar problems arise in different disciplines such as molecular biology, databases of objects, face recognition, matching of articulated objects and pattern recognition in general.

There have been many approaches to the problem of (pose invariant) shape matching and recognition in the literature, for example [HK03, LH05, OFCD02, RWP05, EK03, Fro90, RTG00, CG99, MS05]. In many cases the underlying idea revolved around the comparison of certain metric invariants of the shapes so as to ascertain whether they were in fact the *same shape* (up to a certain notion of invariance).

The concept of Gromov-Hausdorff distances [Gro99] was first proposed as a tool for formalizing shape comparison ideas in [MS05]. This distance is able to detect the *metric* similarity between the shapes as it operates on their metric

structure, that is, shapes are viewed as *metric spaces*. This notion of distance compares the full metric information contained in the shapes, as opposed to other notions that may only compare simple (incomplete) invariants. Therefore two shapes will be declared *equal* if and only if they are *isometric*. This means that the invariance properties are to be encoded by the metrics one chooses to endow the shapes with. For example, if the shapes are endowed with Euclidean metrics, the underlying invariance is to *rigid* isometries.

The ideas presented in this paper are not restricted to 3D shapes as they can be applied to any *point clouds* (sets of points) which are endowed with metric structures.

This paper presents new results that extend the original definition of Gromov-Hausdorff distances in a way such that the associated discrete problems one needs to solve in practical applications are of an easier nature than yielded by previous related approaches, [MS05, BBK06]. The practical approach put forward in [MS05] was inherently combinatorial and hard to optimize. Another feature of this approach is that the approximation bounds between the discrete and continuous entities were probabilistic, and the proof of these required the assumption that the shapes were actually smooth embedded manifolds.

In [BBK06], the authors, also under the assumption that the underlying shapes are smooth surfaces, proposed a con-

[†] Supported by DARPA grant number HR0011-05-1-0007.

tinuous optimization setting which seemed to alleviate some of the impracticalities of the original proposal. They however had to resort to local interpolation of the metrics in order to make their computations possible (this is where the smoothness assumption is used). The underlying ideas were still of a somewhat combinatorial nature in the sense that their *matchings* were *maps* between the shapes. In this paper we completely remove this feature from the framework and consider a different way of *pairing* two shapes which directly leads to *standard* optimization problems where no additional assumptions are needed, i.e. no assumption about the smoothness of the underlying shape. We show that our approach leads to Quadratic Optimization Problems (QOPs) with *linear* constraints. In addition, this new set of ideas allows connecting the Gromov-Hausdorff framework to pre-existing approaches which have proven useful in the Shape Comparison/Matching arena. Examples of these methods are those proposed by [OFCD02, HK03, EK03, CG99, KV05]. In more detail, our modified notion of Gromov-Hausdorff distance will admit these notions of similarity as lower and/or upper bounds.

Throughout our presentation we use some simple concepts from measure theory and point set topology which can be consulted for example in [Dud02].

The paper is organized as follows: Section §2 introduces the problem of Shape Comparison in a general setting and presents basic elements such as notions of shape similarity upon which the rest of the paper is based. Section §3 briefly discusses the idea of introducing invariances into our notions of similarity. Section §4 reviews the notion of Gromov-Hausdorff distance and its main properties. In that section we also discuss connections with the Quadratic Assignment Problem. Section §5 delves into the core of the paper where a new notion of proximity between metric spaces is introduced and its connections with the Gromov-Hausdorff distance and other notions are established. Some basic lower and upper bounds are presented there. Section §6 presents other more interesting lower and upper bounds for the proposed notion of similarity. The aim is twofold, on one hand doing this makes apparent the connection to other approaches found in the literature, and on the other it provides lower bounds which are easily computable and consequently of practical value. Section §7 discusses the computational aspect of our ideas, establishing that the problems one needs to solve in practice are either Linear or Quadratic Optimization Problems (with linear constraints). We present computational examples in Section §8 and conclusions in Section §9. Due to space limitations, long technical proofs are not given in this paper and will be presented elsewhere.

2. Comparing objects

An *object* in a compact metric space (Z, d) will be a compact subset of Z . Let $\mathcal{C}(Z)$ denote the set of all compact subsets of Z (objects).

Assume that inside the metric space (Z, d) we are trying

to compare objects A and B . One possibility is to use the Hausdorff distance:

$$d_{\mathcal{H}}^Z(A, B) := \max \left(\sup_{a \in A} \inf_{b \in B} d(a, b), \sup_{b \in B} \inf_{a \in A} d(a, b) \right) \quad (1)$$

In general, whenever one intends to compare two objects, a correspondence/alignment is established for this purpose. The following definition and Proposition make this apparent for the case of the Hausdorff distance.

Definition 1 (Correspondence) For sets A and B , a subset $R \subset A \times B$ is a correspondence (between A and B) if and only if

- $\forall a \in A$, there exists $b \in B$ s.t. $(a, b) \in R$
- $\forall b \in B$, there exists $a \in A$ s.t. $(a, b) \in R$

Let $\mathcal{R}(A, B)$ denote the set of all possible correspondences between sets A and B .

Proposition 1 Let (Z, d) be a compact metric space. Then the Hausdorff distance between any two sets $A, B \subset Z$ can be expressed as:

$$d_{\mathcal{H}}^Z(A, B) = \inf_R \sup_{(a, b) \in R} d(a, b) \quad (2)$$

where the infimum is taken over all $R \in \mathcal{R}(A, B)$.

The Hausdorff distance is indeed a metric on the set of compact subsets of the (compact) metric space (Z, d) .

Proposition 2 [BBI01]

1. Let $A, B, C \in \mathcal{C}(Z)$ then

$$d_{\mathcal{H}}^Z(A, B) \leq d_{\mathcal{H}}^Z(A, C) + d_{\mathcal{H}}^Z(B, C).$$

2. If $d_{\mathcal{H}}^Z(A, B) = 0$ for $A, B \in \mathcal{C}(Z)$ then $A = B$.

In practice, these two properties are desirable, and we will insist on having them for whichever notion of similarity between shapes we chose to work with. These properties imply in particular that if one is interested in comparing objects A and B , and if $\mathbb{A}_n \subset Z$ and $\mathbb{B}_m \subset Z$ are finite (maybe “noisy”) samples of A and B respectively, then

$$|d_{\mathcal{H}}^Z(A, B) - d_{\mathcal{H}}^Z(\mathbb{A}_n, \mathbb{B}_m)| \leq d_{\mathcal{H}}^Z(A, \mathbb{A}_n) + d_{\mathcal{H}}^Z(B, \mathbb{B}_m) \quad (3)$$

In practice we always have to rely on finite samples of an object. The quality of the approximation of an object A by such a finite set \mathbb{A}_n can obviously be described by $d_{\mathcal{H}}^Z(A, \mathbb{A}_n)$.

Therefore (3) tells us that comparing these discrete samples gives us an answer as good as the approximation of the underlying objects by these discrete sets.

The new ideas we propose rely on the idea of relaxing the notion of correspondence as given by Definition 1. In order to do this we need to introduce a new class of objects which

we call $\mathcal{C}_w(Z)$ (and define precisely below). Roughly speaking, an object in this class will be specified by not only the set of points that form it but one is also required to specify a distribution of importance over these points.

This relaxed notion of correspondence between objects is called *matching measure* (or coupling) and is made precise in Definition 3 below.

There is a very well known family of distances which makes use of this different way of pairing objects. These are the so called, Mass Transportation distances [Vil03] (a.k.a. Wasserstein-Kantorovich-Rubinstein distances as known in the Math community or Earth Mover's distance [Mum91, RTG00] in the Shape Recognition arena). We review those concepts next.

2.1. Transporation Distances

Assume as before that $A, B \in \mathcal{C}(Z)$ and let μ_A and μ_B be Borel probability measures with supports A and B , respectively. Informally, this means that if a set $C \subset Z$ is s.t. $C \cap A = \emptyset$ then $\mu_A(C) = 0$. A precise definition is given below:

Definition 2 *The support of a measure μ on a metric space (Z, d) , denoted by $\text{supp}[\mu]$, is the minimal closed subset $Z_0 \subset Z$ such that $\mu(Z \setminus Z_0) = 0$.*

These probability measures can be thought of as acting as weights for each point in each of the sets. A simple interpretation is that for each $a \in A$, $r > 0$, $\mu_A(B(a, r))$ quantifies the (relative) importance of the point a at scale r (in the discrete case this measure can clearly be interpreted as signaling how much we trust the sample point). In other words, if a' is another point in A and if $\mu_A(B(a, r)) \leq \mu_A(B(a', r))$ we would say that a' is more important than a at scale r . Note that since μ_A (μ_B) is a probability measure, $\mu_A(A) = 1$ ($\mu_B(B) = 1$).

We naturally require that $A = \text{supp}[\mu_A]$ and $B = \text{supp}[\mu_B]$. By taking μ_A and μ_B into account we will therefore be comparing not only the geometry of the sets, but also, the distribution of “importance” over the sets. We introduce

$$\mathcal{C}_w(Z) := \{(A, \mu_A), A \in \mathcal{C}(Z)\}$$

where μ_A is a Borel probability measure with $\text{supp}[\mu_A] = A$.

Definition 3 (Matching Measure) *Let $A, B \in \mathcal{C}_w(Z)$. We say that a measure μ on the product space $A \times B$ is a coupling of μ_A and μ_B iff*

$$\mu(A_0 \times B) = \mu_A(A_0), \quad \mu(A \times B_0) = \mu_B(B_0) \quad (4)$$

for all Borel sets $A_0 \subset A$, $B_0 \subset B$. We denote by $\mathcal{M}(\mu_A, \mu_B)$ the set of all couplings of μ_A and μ_B .

It turns out that for each $\mu \in \mathcal{M}(\mu_A, \mu_B)$, $\text{supp}[\mu] \subset A \times B$ is a *correspondence* which we denote by $R(\mu)$.

Lemma 1 *Given $\mu \in \mathcal{M}(\mu_A, \mu_B)$, then $R(\mu) := \text{supp}[\mu]$ belongs to $\mathcal{R}(A, B)$.*

For each $p \geq 1$ we consider the following family of distances on $\mathcal{C}_w(Z)$:

$$d_{W,p}^Z(A, B) := \inf_{\mu \in \mathcal{M}(\mu_A, \mu_B)} \left(\int_{A \times B} d^p(a, b) d\mu(a, b) \right)^{1/p} \quad (5)$$

for $1 \leq p < \infty$, and

$$d_{W,\infty}^Z(A, B) := \inf_{\mu \in \mathcal{M}(\mu_A, \mu_B)} \sup_{(a,b) \in R(\mu)} d(a, b) \quad (6)$$

These distances are none other than the Wasserstein-Kantorovich-Rubinstein distances between measures, [Vil03, Dud02]. These distances have been considered for Shape Comparison/Matching applications several times (for some values of p , typically $p = 1$ or 2), see for example [RTG00, CG99, KV05].

As we did in the definition of $d_{W,p}^Z$, in the sequel we will abuse notation by sometimes representing an object $(A, \mu_A) \in \mathcal{C}_w(Z)$ also by either A or μ_A . The reader should keep in mind, however, that a measurable set $A \subset Z$ can be represented by many probability measures, all that is required is that those probability measures have support A .

An initial question, which is now easy to answer is how do these distances relate to $d_{\mathcal{H}}^Z(\cdot, \cdot)$. Upon noting that (1) and (6) are essentially the same expression and that $R(\mu) \in \mathcal{R}(A, B)$ we obtain:

Proposition 3 *For (A, μ_A) and (B, μ_B) in $\mathcal{C}_w(Z)$*

$$d_{\mathcal{H}}^Z(A, B) \leq d_{W,\infty}^Z((A, \mu_A), (B, \mu_B))$$

for all choices of μ_A and μ_B such that $A = \text{supp}[\mu_A]$ and $B = \text{supp}[\mu_B]$.

This connection between Hausdorff and Mass Transportation distances has already been pointed out in the robotics literature, [HMO4].

We review the main properties of this family of distances next.

Proposition 4 [Vil03]

1. *For each $1 \leq p \leq \infty$ $d_{W,p}^Z$ defines a metric on $\mathcal{C}_w(Z)$.*
2. *For any $1 \leq p \leq q \leq \infty$ and $A, B \in \mathcal{C}_w(Z)$*

$$d_{W,p}^Z(A, B) \leq d_{W,q}^Z(A, B).$$

Distances $d_{W,p}^Z$, for finite p offer an interesting alternative to the Hausdorff distance. Note that in the finite discrete case, computing them involves solving a Linear Optimization Problem (LOP).

Assume \mathbb{A}_n and \mathbb{B}_m are finite (possibly noisy) samples from A and B , respectively. Assume further that each of them

is given together with a discrete probability measure μ_n and ν_n , respectively. Then Property 1 above implies that (cf. (3))

$$|d_{W,p}^Z(A, B) - d_{W,p}^Z(\mathbb{A}_n, \mathbb{B}_m)| \leq d_{W,p}^Z(A, \mathbb{A}_n) + d_{W,p}^Z(B, \mathbb{B}_m). \quad (7)$$

This is a trivial observation, the point one must make is that now the quality of the approximation of A by its discrete representative \mathbb{A}_n is governed by the Wasserstein-Kantorovich-Rubinstein distance between μ_A and μ_n . For objects (A, μ_A) and (B, μ_B) we sometimes abuse the notation by writing $d_{W,p}^Z(\mu_A, \mu_B)$ instead of $d_{W,p}^Z(A, B)$.

3. Introducing Invariances

Let $\mathcal{O}(Z)$ denote our class of objects. We have mentioned two choices: $\mathcal{C}(Z)$ and $\mathcal{C}_w(Z)$.

For the case of any well behaved notion of distance D on $\mathcal{O}(Z)$ one can decide to study the following problem: Let \mathcal{T} be the group of isometries on Z (i.e. $\forall T \in \mathcal{T}, z, z' \in Z, d(T(z), T(z')) = d(z, z')$) and $A, B \in \mathcal{O}(Z)$, then consider

$$D^{\mathcal{T}}(A, B) := \inf_{T \in \mathcal{T}} D(A, T(B)). \quad (8)$$

The most common case, $(Z, d) = (\mathbb{R}^k, \|\cdot\|)$, has been approached by several authors in the past, with $D = d_{\mathcal{H}}(\cdot)$ (and $\mathcal{O}(Z) = \mathcal{C}(Z)$) in [HKR93], and $D = d_{W,p}^Z(\cdot)$ (and $\mathcal{O}(Z) = \mathcal{C}_w(Z)$) [CG99, RTG00] and references therein.

A different idea, used to certain extent in [HK03, OFCD02, EK03, BK04] consists of comparing the invariants to $T \in \mathcal{T}$. This leads to comparing the metric information of A and B more directly. More precisely, if $A = \{a_1, \dots, a_m\}$ and $B = \{b_1, \dots, b_n\}$ are finite sets of points, ideally one would like to meaningfully compare the distance matrices $D_A = ((\|a_i - a_j\|))$ and $D_B = ((\|b_i - b_j\|))$ in a fashion compatible with the particular choice of D (and $\mathcal{O}(Z)$) that we have made. In the past, people have been resorting to comparisons between simple invariants constructed from the distance matrices. For example, in [HK03] the authors use (essentially) the row sums of the distance matrices as the invariants they compare.

At any rate, comparing these distances matrices amounts to considering $(A, \|\cdot\|)$ and $(B, \|\cdot\|)$ as metric spaces without any reference to Z .

This is the basic idea of the so called Gromov-Hausdorff distances (which arises from the choice $D = d_{\mathcal{H}}^Z(\cdot)$ and $\mathcal{O}(Z) = \mathcal{C}(Z)$) and Measured Gromov-Hausdorff distances (arising from the choice $D = d_{W,p}^Z$ and $\mathcal{O}(Z) = \mathcal{C}_w(Z)$). We discuss these ideas next.

4. The Gromov-Hausdorff Distance

Following [Gro99], we introduce the Gromov-Hausdorff distance between (compact) metric spaces X and Y :

$$d_{\mathcal{GH}}(X, Y) := \inf_{Z, f, g} d_{\mathcal{H}}^Z(f(X), g(Y)) \quad (9)$$

where $f : X \rightarrow Z$ and $g : Y \rightarrow Z$ are isometric embeddings (distance preserving) into the metric space Z . This expression seems daunting from the computational point of view. We will recall equivalent tamer expressions below. Nevertheless, this expression helps framing the procedure of [EK03] inside the Gromov-Hausdorff realm, see [MS05].

Definition 4 (Metric Coupling) From now on let $\mathcal{D}(d_X, d_Y)$ denote the set of all possible metrics on the disjoint union of X and Y , $X \sqcup Y$. This means that besides satisfying all triangle inequalities, it also holds that if $d \in \mathcal{D}(d_X, d_Y)$ then $d(x, x') = d_X(x, x')$ and $d(y, y') = d_Y(y, y')$ for all $x, x' \in X$ and $y, y' \in Y$.

Remark 1 One can equivalently (in the sense of equality) define the Gromov-Hausdorff distance between metric spaces (X, d_X) and (Y, d_Y) as ([BBI01])

$$d_{\mathcal{GH}}(X, Y) = \inf_{R, d} \sup_{(x, y) \in R} d(x, y) \quad (10)$$

where the infimum is taken over $R \in \mathcal{R}(X, Y)$ and $d \in \mathcal{D}(d_X, d_Y)$.

Next, we state some well known properties of the Gromov-Hausdorff distance $d_{\mathcal{GH}}(\cdot)$ which will be essential for our presentation.

From now on, for metric spaces (X, d_X) and (Y, d_Y) let $\Gamma : X \times X \times Y \times Y \rightarrow \mathbb{R}^+$ be given by

$$\Gamma(x, x'; y, y') := |d_X(x, x') - d_Y(y, y')| \quad (11)$$

Proposition 5 1. Let (X, d_X) , (Y, d_Y) and (Z, d_Z) be metric spaces then

$$d_{\mathcal{GH}}(X, Y) \leq d_{\mathcal{GH}}(X, Z) + d_{\mathcal{GH}}(Y, Z).$$

2. If $d_{\mathcal{GH}}(X, Y) = 0$ and (X, d_X) , (Y, d_Y) are compact metric spaces, then (X, d_X) and (Y, d_Y) are isometric.
3. Let $\{x_1, \dots, x_n\} \subset X$ be a R -covering of the compact metric space (X, d_X) . Then $d_{\mathcal{GH}}(X, \{x_1, \dots, x_n\}) \leq R$.
4. For compact metric spaces (X, d_X) and (Y, d_Y) :

$$\begin{aligned} \frac{1}{2} |\text{diam}(X) - \text{diam}(Y)| &\leq d_{\mathcal{GH}}(X, Y) \\ &\leq \frac{1}{2} \max(\text{diam}(X), \text{diam}(Y)) \end{aligned} \quad (12)$$

where $\text{diam}(X) := \max_{x, x' \in X} d_X(x, x')$ stands for the Diameter of the metric space (X, d_X) .

5. For bounded metric spaces (X, d_X) and (Y, d_Y) ,

$$d_{\mathcal{GH}}(X, Y) = \frac{1}{2} \inf_{R \in \mathcal{R}(X, Y)} \sup_{\substack{x_1, x_2 \in X \\ y_1, y_2 \in Y \\ s.t. (x_i, y_i) \in R}} \Gamma(x_1, x_2; y_1, y_2) \quad (13)$$

Proofs of Properties 1 to 5 can be found in [BBI01].

Remark 2 It is possible to use Gromov-Hausdorff ideas to define a notion of partial similarity between two objects, see [MS05].

Remark 3 We want to argue that expression (13) is reminiscent of the QAP (Quadratic Assignment Problem). This will let us loosely infer something about the inherent complexity of computing the Gromov-Hausdorff distance. Let's restrict ourselves to the case of finite metric spaces, $\mathbb{X} = \{x_1, \dots, x_n\}$ and $\mathbb{Y} = \{y_1, \dots, y_m\}$. For $R \in \mathcal{R}(\mathbb{X}, \mathbb{Y})$ let δ_{ij}^R equal 1 if $(i, j) \in R$ and 0 otherwise. Then we have:

$$d_{\mathcal{GH}}(X, Y) = \frac{1}{2} \min_R \max_{i,k,j,l} \Gamma_{ikjl} \delta_{ij}^R \delta_{kl}^R$$

where $\Gamma_{ikjl} := |d_X(x_i, x_k) - d_Y(y_j, y_l)|$. Now, one can obtain a family of related problems by relaxing the max to a sum as follows: Fix $p \geq 1$ and let $\Gamma_{ikjl}^{(p)} := |d_X(x_i, x_k) - d_Y(y_j, y_l)|^p$, then one can also consider the problem:

$$(\mathcal{P}_p) \quad \min_R \sum_{ij} \sum_{kl} \Gamma_{ikjl}^{(p)} \delta_{ij}^R \delta_{kl}^R.$$

Note that one can recast the above problem as follows. Let Δ denote the set of matrices defined by the constraints below:

1. $\delta_{ij} \in \{0, 1\}$ for all i, j
2. $\sum_i \delta_{ij} \geq 1$ for all j
3. $\sum_j \delta_{ij} \geq 1$ for all i

and let $L_p(\delta) := \sum_{ij} \sum_{kl} \Gamma_{ikjl}^{(p)} \delta_{ij} \delta_{kl}$. then (\mathcal{P}_p) is equivalent to

$$\min_{\delta \in \Delta} L_p(\delta)$$

which can be regarded as a generalized version of the QAP. In the standard QAP ([PW94]) $n = m$ and the inequalities 2. and 3. defining Δ above are actually equalities, what forces each δ to be a permutation matrix.

Actually, we prove next that, when $n = m$, (\mathcal{P}_p) reduces to a QAP. It is known that the QAP is an NP-hard problem [PW94].

In fact, it is clear that for any $\delta \in \Delta$ there exist $\pi \in \Pi_n$ ($n \times n$ permutations matrices) such that $\delta_{ij} \geq \pi_{ij}$ for all $1 \leq i, j \leq n$. Then, since $\Gamma_{ikjl}^{(p)}$ is non negative for all $1 \leq i, j, k, l \leq n$, it follows that $L_p(\delta) \geq L_p(\pi)$. Therefore the minimal value of $L_p(\delta)$ is attained at some $\delta \in \Pi_n$.

Remark 4 It was pointed out in [MS05] that Property 5 above can be recast in a somewhat clearer form: For functions $\phi : X \rightarrow Y$ and $\psi : Y \rightarrow X$ consider the numbers $A(\phi) := \sup_{x_1, x_2 \in X} |d_X(x_1, x_2) - d_Y(\phi(x_1), \phi(x_2))|$, $B(\psi) := \sup_{y_1, y_2 \in Y} |d_X(\psi(y_1), \psi(y_2)) - d_Y(y_1, y_2)|$ and $C(\phi, \psi) := \sup_{x \in X, y \in Y} |d_X(x, \psi(y)) - d_Y(\phi(x), y)|$, then

$$d_{\mathcal{GH}}(X, Y) = \inf_{\substack{\phi : X \rightarrow Y \\ \psi : Y \rightarrow X}} \frac{1}{2} \max(A(\phi), B(\psi), C(\phi, \psi)) \quad (14)$$

Formula (14) is suggestive from the computational point of view and leads to considering certain algorithmic procedures such as those in [MS05, BBK06].

In this paper we carry out a modification of the original formulation of [MS05], namely, we propose to substitute the underlying Hausdorff distance by a *relaxed* notion of proximity between objects (more precisely by the Wasserstein-Kantorovich-Rubinstein distance) and then find what the equivalent version of Property 5 in Proposition 5 would be.

The basic idea is to consider *three* out of the (four) different expressions we have for the Gromov-Hausdorff distance and try to pick the one that will provide the most computationally tractable framework without sacrificing the theoretical underpinnings. These three expressions are (13), (10) and (14). The path starting at (14) has been explored first in [MS05] and later in [BBK06]. In this paper we concentrate therefore on (13) and (10). We argue below that these two options are more natural and finally single out one of them based on computational cost considerations of the associated discrete problem. Interestingly, we will also show that the two final expressions, call them (13)* and (10)* are not related by an equality, in contrast with the fact that (13) = (10), see Remark 8 below.

In contrast with the line followed in [MS05, BBK06] which is expounded and justified for points sampled from smooth surfaces, the formalism of measure metric spaces used here (and explained below) allows our approach to work in more (theoretical and practical) generality.

One initial observation is that both expressions (13) and (10) make use of the notion of Relation/Correspondence. We have seen in previous sections that in fact, at the level of Hausdorff distances the formal substitution of correspondences for measure couplings, and of max for L_p norms ($p \geq 1$) leads to Wasserstein-Kantorovich-Rubinstein distances. We now carry out the same program on the Gromov-Hausdorff distance. One more word about this program, that directly alludes to the promised computational advantage of foregoing correspondences in favor of matching measures, is that the former objects are essentially of combinatorial nature whereas the latter objects can take continuous values, even in the case of discrete spaces, this point is further discussed in §7.

5. L_p Gromov-Hausdorff Distances

It is the purpose of this paper to present a modification of the ideas underlying the definition of the Gromov-Hausdorff distance given above that is better suited for practical applications. For this we consider more structure than just a set of points with a metric on them: We also assume a probability measure is given on the (sets of) points, as was the case in §2.1. Again, this probability measure can be thought of as indicating the importance of the difference points in the dataset.

5.1. Measure Metric Spaces

Definition 5 [Gro99] A metric measure space (mm-space for short) will always be a triple (X, d_X, μ_X) where

- (X, d_X) is a compact metric space.
- μ_X is a Borel probability measure on X i.e. $\mu_X(X) = 1$.

When it is clear from the context, we will denote the triple (X, d_X, μ_X) by only X . The reason for imposing $\mu_X(X) = 1$ is that we think of μ_X as a modelization of the acquisition process or sampling procedure. Moreover, we will assume w.l.o.g. that for all our mm-spaces $X = \text{supp}[X]$.

Two mm-spaces (X, d_X, μ_X) and (Y, d_Y, μ_Y) are called *isomorphic* iff there exists an isometry $\psi : \text{supp}[\mu_X] \rightarrow \text{supp}[\mu_Y]$ such that $\mu_X(\psi^{-1}(B)) = \mu_Y(B)$ for all $B \subset Y$ measurable.

5.2. The distance

As before we need to introduce a notion of correspondence/coupling between the mm-spaces involved in the comparison. The definition below is essentially the same as Definition 3.

Definition 6 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

$$\mu(A \times Y) = \mu_X(A), \quad \mu(X \times A') = \mu_Y(A') \quad (15)$$

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 .

Starting from (13) we now construct a new, tentative notion of distance between metric spaces. We use this expression as our starting point because we want the new distance to directly compare the metrics of X and Y (in a meaningful way). Roughly, we will substitute the maxs in (13) by L_p norms and correspondences by coupling measures.

For $p \in [1, \infty)$ and $\mu \in \mathcal{M}(\mu_X, \mu_Y)$ let

$$J_p(\mu) := \frac{1}{2} \left(\int_{X \times Y} \int_{X \times Y} \Gamma(x, x'; y, y')^p d\mu(x, y) d\mu(x', y') \right)^{1/p} \quad (16)$$

and also let

$$J_\infty(\mu) := \frac{1}{2} \sup_{\substack{x, x' \in X \\ y, y' \in Y \\ \text{s.t. } (x, y), (x', y') \in R(\mu)}} \Gamma(x, x'; y, y') \quad (17)$$

Remark 5 Note that given the definitions above one has that (under suitable regularity assumptions)

$$J_p(\mu) \xrightarrow{p \uparrow \infty} J_\infty(\mu)$$

Definition 7 For $\infty \geq p \geq 1$ we define the distance D_p between two mm-spaces X and Y by

$$D_p(X, Y) := \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} J_p(\mu) \quad (18)$$

One actually needs to prove that expression (18) in fact defines a metric on the set of all isomorphism classes of mm-spaces. This is an interesting technical step in itself. These and other properties of D_p , of similar spirit to those reported for $d_{\mathcal{GH}}(\cdot, \cdot)$ in Proposition 5 are listed in Proposition 6 below.

Remark 6 In [Stu06] Sturm introduced another distance for mm-spaces (for each $p \geq 1$) as follows: (he presented the case $p = 2$)

$$S_p(X, Y) := \inf_{\mu, d} \left(\int_{X \times Y} d(x, y)^p \mu(dx, dy) \right)^{1/p} \quad (19)$$

where the infimum is taken over all $d \in \mathcal{D}(d_X, d_Y)$ (recall Definition 4) and $\mu \in \mathcal{M}(\mu_X, \mu_Y)$.

The corresponding definition for $p = \infty$ is

$$S_\infty(X, Y) = \inf_{\mu, d} \sup_{(x, y) \in R(\mu)} d(x, y) \quad (20)$$

This proposal corresponds to what we called (10)*.

Remark 7 Note the similarity between S_∞ and $d_{\mathcal{GH}}(\cdot, \cdot)$ as given by formula (14). Since $\{R(\mu) \mid \mu \in \mathcal{M}(\mu_X, \mu_Y)\} \subset \mathcal{R}(X, Y)$ it is obvious that $S_\infty(X, Y) \geq d_{\mathcal{GH}}(X, Y)$, cf. Proposition 2.

Remark 8 At this point it is clear that in our construction, $(13)^* = (18)$ and $(10)^* = (19)$. Note that since (13) and (10) are equal, one could conjecture (18) and (19) to be equal as well. In this respect, one can prove that $S_p \geq D_p$ for $1 \leq p \leq \infty$ and that $S_\infty = D_\infty$. However, for $p < \infty$ the equality does not hold in general. One simple counterexample is the following: For each $n \geq 3$ let $X = X_n$ be the $(n-1)$ -simplex (which has n points) endowed with metric $d_{ij} = 1$ for $i \neq j$ and probability measure $\nu_i = \frac{1}{n}$. Let Y be a single point $\{y\}$. It is easy to verify that then $S_1(X_n, Y) \geq \frac{1}{2}$. On the other hand, from Proposition 6 below we know that $D_1(X_n, Y) = \frac{1}{2} \sum_i \sum_j \nu_i \nu_j d_{ij} = \frac{1 - \sum_i \nu_i^2}{2} = \frac{n-1}{2n}$. Hence we see that $S_1(X_n, Y) > D_1(X_n, Y)$ for $n \geq 2$.

As we argue in §7, the expression D_p is more amenable to numerical computations than that of S_p .

Remark 9 One may wonder what is the relationship between $d_{\mathcal{GH}}(X, Y)$ and (some of) the $D_p(X, Y)$'s. In this respect the proposition below asserts that $d_{\mathcal{GH}}(X, Y) \leq D_\infty(X, Y)$.

Definition 8 For a mm-space (X, d_X, μ_X) , for $p \in [1, \infty]$ we define its p -diameter as

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

for $1 \leq p < \infty$, and $\mathbf{diam}_\infty(X) := \frac{\mathbf{diam}(\text{supp}[\mu_X])}{2}$.

One can prove the following properties of D_p :

Proposition 6 (a) For each $p \geq 1$, D_p defines a metric on the set of all (isomorphism classes of) mm-spaces.
 (b) Whenever $\text{supp}[\mu_X] = X$ and $\text{supp}[\mu_Y] = Y$ we have

$$d_{\mathcal{GH}}(X, Y) \leq D_\infty(X, Y).$$

(c) What happens under two different probability measures on the same space (keeping the same metric)? Let (Z, d) be a compact metric space and α and β two different Borel probability measures on Z . Let $X = (Z, d, \alpha)$ and $Y = (Z, d, \beta)$ then

$$D_p(X, Y) \leq d_{W,p}^Z(\alpha, \beta).$$

(d) What happens under two different metrics on the same space (keeping the same probability measure)? Let Z be a compact metric space and α be a Borel probability measure on Z . Let $X = (Z, d, \alpha)$ and $Y = (Z, d', \alpha)$ then

$$D_p(X, Y) \leq \frac{1}{2} \|d - d'\|_{L_p(Z \times Z, \alpha \otimes \alpha)}.$$

(e) (What happens for a random sampling of the metric space?) Let $\mathbb{X}_m \subset X$ be a set of m random variables $\mathbf{x}_i : \Omega \rightarrow X$ defined on some probability space Ω with law μ_X . Let $\mu_m(\omega, \cdot) := \frac{1}{m} \sum_{i=1}^m \delta_{\mathbf{x}_i(\omega)}$ denote the empirical measure. For each $\omega \in \Omega$ consider the metric measure spaces (X, d_X, μ_X) and (X, d_X, μ_m) , then for μ_X -almost all $\omega \in \Omega$, $(X, d_X, \mu_m) \xrightarrow{D_p} (X, d_X, \mu_X)$ as $m \uparrow \infty$.

(f) When $Y = \{y\}$, for $p \in [1, \infty]$ then

$$D_p(X, Y) = \frac{\mathbf{diam}_p(X)}{2}$$

and from this and Property (a) (triangle inequality)

$$D_p(X, Y) \geq \left| \frac{\mathbf{diam}_p(X) - \mathbf{diam}_p(Y)}{2} \right| \quad (21)$$

(g) For mm-spaces X and Y and for $1 \leq p < \infty$ it holds that

$$S_p(X, Y) \geq D_p(X, Y).$$

Also, for mm-spaces X and Y such that $X = \text{supp}[\mu_X]$ and $Y = \text{supp}[\mu_Y]$ it holds that

$$S_\infty(X, Y) = D_\infty(X, Y)$$

(h) Upper bound in terms of p -diameters. For $p \geq 1$ one has:

$$D_p(X, Y) \leq \left(\frac{(\mathbf{diam}_p(X))^p + (\mathbf{diam}_p(Y))^p}{2} \right)^{1/p}.$$

(i) Ordering of the different distances: $D_p \geq D_q$ when $p \geq q \geq 1$

Remark 10 Writing the framework with p as a parameter is not gratuitous nor superfluous. In fact even the simple bound (21) will be useful for discriminating between certain metric spaces that the corresponding Gromov-Hausdorff bound (12) cannot. For example, consider the case when $X = (S^1, d_1, \mu_1)$ and $Y = (S^2, d_2, \mu_2)$ where d_1 and d_2 are the usual spherical distance metrics and μ_1 and μ_2 stand for normalized area on S^1 and S^2 , respectively. Then (12) vanishes as $(21)_{p=\infty}$ also does. However, since $\mathbf{diam}_2(S^1) = \pi/\sqrt{3}$ and $\mathbf{diam}_2(S^2) = \sqrt{\frac{\pi^2}{2} - 2}$, it follows that $(21)_{p=2}$ does permit telling S^1 and S^2 apart.

6. Lower and Upper Bounds and Connections to Other Approaches

In practice, having lower bounds that are easy to compute (in the sense that they are not computationally expensive) is very important as they facilitate classification tasks: If during a query the value of the lower bound is above a certain threshold one would say that the answer to the query is negative without incurring the potentially higher computational cost of evaluating the full metric.

Also, we intend to relate our framework to other proposals in the literature, we will do this by establishing that several notions of dissimilarity between shapes are essentially lower or upper bounds to our proposed D_p (18). We have already seen some lower and upper bounds in §5, in particular the relation to $d_{\mathcal{GH}}(\cdot, \cdot)$ and Sturm's proposal have been discussed there.

6.1. First Lower Bound

For $\infty > p \geq 1$ let $s_p^X(x) = \left(\int_X d_X(x, x')^p \mu_X(dx) \right)^{1/p}$ and $s_\infty^X(x) = \sup_{x' \in \text{supp}[\mu_X]} d_X(x, x')$. Then one obtains the following bounds for D_p :

- Case $\infty > p \geq 1$:

$$D_p(X, Y) \geq \frac{1}{2} \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \int_{X \times Y} |s_p^X(x) - s_p^Y(y)| \mu(dx, dy) \quad (22)$$

which is a Mass Transportation Problem ([Vil03]) for the cost $|s_p^X(x) - s_p^Y(y)|$.

- Case $p = \infty$

$$D_\infty(X, Y) \geq \frac{1}{2} \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \max_{(x, y) \in R(\mu)} |s_\infty^X(x) - s_\infty^Y(y)| \quad (23)$$

For $p \geq 1$ let FLB_p denote the right-hand-side of (22) (and (23)).

Equation (22) is a simple consequence of Minkowski's Inequality and (23) is a consequence of this simple fact:

for functions $f, g : Z \rightarrow \mathbb{R}$ it holds that $|\max f - \max g| \geq \max |f - g|$.

This lower bound is roughly what was computed by Hamza and Krim in [HK03]: They used the functions s_2^X for describing each space X .

6.2. Second Lower Bound

In this section we prove a lower bound for our family of distances which relies on some kind of comparison of the *distribution* of distances of X and Y . For $p \in [1, \infty)$ consider the following optimization problem:

$$\text{SLB}_p(X, Y) := \frac{1}{2} \inf_{\gamma \in \widehat{\mathcal{M}}} \left(\int_{X \times X} \int_{Y \times Y} |d_X(\mathcal{X}) - d_Y(\mathcal{Y})|^p d\gamma(\mathcal{X}, \mathcal{Y}) \right)^{1/p} \quad (24)$$

where $\widehat{\mathcal{M}} = \mathcal{M}(\mu_X \otimes \mu_X, \mu_Y \otimes \mu_Y)$ stands for the set of probability measures on $X \times X \times Y \times Y$ with marginals $\mu_X \otimes \mu_X$ and $\mu_Y \otimes \mu_Y$. Clearly, (24) provides a lower bound to $D_p(X, Y)$ since for any $\mu \in \mathcal{M}(\mu_X, \mu_Y)$, $\mu \otimes \mu \in \mathcal{M}(\mu_X \otimes \mu_X, \mu_Y \otimes \mu_Y)$. A similar claim is true for $p = \infty$.

It is clear now that this bound is nothing but a measure of distance between the distribution of interpoint distances in X and Y . In fact, one can prove the following:

Proposition 7 *For any mm-space (Z, d_Z, μ_Z) and $t \in [0, \text{diam}(Z)]$, let $F_Z(t) = \mu_Z \otimes \mu_Z(\{(z, z') | d_Z(z, z') \leq t\})$. Then for $p \in [1, \infty)$*

$$\text{SLB}_p(X, Y) \geq \frac{1}{2} \left(\int_0^1 |F_X^{-1}(u) - F_Y^{-1}(u)|^p du \right)^{1/p}$$

and for $p = 1$ the expression simplifies to

$$\text{SLB}_1(X, Y) \geq \frac{1}{2} \int_0^{\infty} |F_X(t) - F_Y(t)| dt.$$

Now, $F_Z(t)$ can be interpreted as follows: Assume that one randomly samples two points \mathbf{z} and \mathbf{z}' from Z independently, and each distributed according to the law μ_Z , then $F_Z(t)$ equals the probability that the distance between these two random samples is not greater than t , that is, $F_Z(t) = \mathbf{P}(d_Z(\mathbf{z}, \mathbf{z}') \leq t)$. This is exactly one of the signatures computed by Osada et al in the famous *Shape Distributions* approach to comparing shapes, [OFCD02]. This line was pursued in more theoretical terms, for the case of finite Euclidean metric sets, in [BK04].

6.3. Third Lower Bound

For $p \in [1, \infty)$ consider $F_p : X \times Y \rightarrow \mathbb{R}$ given by

$$F_p(x, y) := \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \int_{X \times Y} \Gamma(x, x'; y, y')^p \mu(dx', dy').$$

Then, it is clear that

$$D_p(X, Y) \geq \frac{1}{2} \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \left(\int_{X \times Y} F_p(x, y) \mu(dx, dy) \right)^{1/p} \quad (25)$$

This lower bound is reminiscent of Lawler's lower bound in the QAP literature, see Remark 3.

Let TLB_p denote the right-hand-side of (25). This lower bound is tighter than the one in §6.1 in the sense that $\text{TLB}_p \geq \text{FLB}_p$ for all $p \geq 1$. To the best of our knowledge this is the first time this bound is used in the context of Shape Matching/Comparison.

6.4. Upper Bounds

Assume we want to compare compact subsets of Euclidean space \mathbb{R}^k under invariance to rigid isometries the group of which we denote by \mathcal{T} . Then, following Section §3 one suitable notion of similarity between objects (X, μ_X) and (Y, μ_Y) in $\mathcal{O}_w(\mathbb{R}^k)$ is

$$E_{W,p}(X, Y) := \inf_{T \in \mathcal{T}} d_{W,p}^{\mathbb{R}^k}(X, T(Y)).$$

Consider the mm-spaces $X' = (X, \|\cdot\|, \mu_X)$ and $Y' = (Y, \|\cdot\|, \mu_Y)$, then by simple application of Minkowski's inequality (for the norm $\|\cdot\|$) one can easily prove that

$$D_p(X', Y') \leq E_{W,p}(X', Y').$$

This bound connects our work with [CG99, KV05] and references therein. A similar (stronger) claim holds true when we use $d_{\mathcal{H}}^Z$ instead of $d_{W,p}^Z$ and $d_{\mathcal{G}\mathcal{H}}$ instead of D_p .

7. Computational Technique

In this section we deal with the practical implementation of our ideas. We recast the discrete counterpart of the ideas we proposed as (continuous) optimization problems.

Assume we are given discrete mm-spaces $\mathbb{X} = \{x_1, \dots, x_n\}$ and $\mathbb{Y} = \{y_1, \dots, y_n\}$ with metrics d^X and d^Y , respectively, and probability measures $\lambda = \{\lambda_1, \dots, \lambda_n\}$ and $\nu = \{\nu_1, \dots, \nu_n\}$, respectively.

Let $\mathbf{M} := \{\mu \in \mathbb{R}_+^{n \times n} \mid 0 \leq \mu_{ij} \leq 1, \sum_i \mu_{ij} = \nu_j, \sum_j \mu_{ij} = \lambda_i, \text{ for all } 1 \leq i \leq n_X, 1 \leq j \leq n_Y\}$. Note that the number of constraints in \mathbf{M} (all of which are *linear*) is $(n_X + n_Y)$.

Let $p \in [1, \infty)$. Then the problem we intend to solve is:

$$(P_p) \begin{cases} \min_{\mu \in \mathbf{M}} \mathbf{J}_p(\mu) \\ \mathbf{J}_p(\mu) := \sum_{i,i'=1}^n \sum_{j,j'=1}^n \mu_{ij} \mu_{i'j'} |d_{ij}^X - d_{i'j'}^Y|^p \end{cases}$$

Problem (P_p) is a QOP (with linear constraints), albeit not necessarily convex. Nevertheless there exist a myriad of techniques in the literature for handling this kind of problems. For the computation of examples presented in §8 we implemented an alternate optimization procedure ([Lue03]) which relies on solving successive LOPs and which we initialize by solving the problem FLB_p (see below). We used a Matlab interface, [Gio] for the open source LOP solver `glpk` and `YALMIP` as an interpreter, [Löf04].

Let μ^* be the matching measure we obtain upon convergence of the method. We then estimate $D_p(\mathbb{X}, \mathbb{Y}) \simeq \frac{1}{2} (\mathbf{J}_p(\mu^*))^{1/p}$.

Remark 11 If one were to try to compute $S_p(\mathbb{X}, \mathbb{Y})$, the resulting optimization problem, we argue, would be significantly harder. In fact, the problem would read

$$(S_p) \begin{cases} \min_{(\mu, d) \in (\mathbf{M}, \mathbf{D})} \mathbf{I}_p(\mu, d) \\ \mathbf{I}_p(\mu, d) := \sum_{i=1}^n \sum_{j=1}^n \mu_{ij} d_{ij}^p \end{cases}$$

where $\mathbf{D} = \{d \in \mathbb{R}_+^{n \times n} \text{ s.t. } |d_{ij} - d_{i'j'}| \leq d_{ii'}^X \leq d_{ij} + d_{i'j'} \text{ and } |d_{ij} - d_{i'j'}| \leq d_{jj'}^Y \leq d_{ij} + d_{i'j'}, 1 \leq i, i' \leq n_X, 1 \leq j, j' \leq n_Y\}$. Note that the number of constraints in \mathbf{D} is $N_S := 2(n_Y \binom{n}{2} + n_X \binom{n}{2})$. Therefore for solving (S_p) one needs $2(n_X \times n_Y)$ variables and N_S constraints as opposed to $(n_X \times n_Y)$ variables and $(n_X + n_Y)$ constraints for solving (D_p) . Therefore, from the practical point of view, this justifies singling out D_p as a more convenient choice. Also, it is worth mentioning that problem (S_p) is a Bilinear Optimization problem, which can obviously be recast as a QOP.

It is useful to recast the lower bounds discussed in §6.1 in this optimization setting. We exemplify this for the first lower bound only:

$$(\text{FLB}_p) \begin{cases} \min_{\mu \in \mathbf{M}} \mathbf{L}_p(\mu) \\ \mathbf{L}_p(\mu) := \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \mu_{ij} |s_p^X(i) - s_p^Y(j)| \end{cases}$$

$$\text{where } s_p^X(i) = \left(\sum_{k=1}^n \lambda_k \left(d_{ik}^X \right)^p \right)^{1/p} \text{ and } s_p^Y(j) = \left(\sum_{k=1}^n \nu_k \left(d_{jk}^Y \right)^p \right)^{1/p} \text{ for } 1 \leq i \leq n_X \text{ and } 1 \leq j \leq n_Y.$$

While we do not it here explicitly, it should be clear that the discrete formulations of SLB_p and TLB_p also lead to LOPs. In the latter case, however, one needs to solve $n_X \times n_Y$ LOPs over the variable $\mu \in \mathbf{M}$.

8. Computational Examples

In this section we present some computational examples that exemplify the use of our framework. We used the publicly available (triangulated) shapes database [SP]. This database comprises shapes from seven different classes: camel, cat, elephant, faces, heads, horse and lion.

Each class contains several different *poses* of the same shape. These poses are richer than just rigid isometries, see Figure 1 for an example of what these shapes look like in the case of the camel models. The number of vertices in the models ranged from 7K to 30K. From each model X we selected 4000 points using the Euclidean *farthest point sampling* procedure. Briefly, one first randomly chooses a point from the dataset. Then one chooses 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 \hat{X} denote this reduced model. Then we defined an intrinsic distance using Dijkstra's algorithm on the graph $G(X)$ with vertex set X where each vertex is connected by an edge to those vertices with which it

shares a triangle. We further subsample \hat{X} again using the farthest point procedure (with the distance computed using $G(X)$) and retain only 50 points. Denote the resulting set by \mathbb{X} . We then endowed \mathbb{X} with the *normalized* distance metric inherited from the Dijkstra procedure described above, and a probability measure based on Voronoi partitions: The mass (measure) at point $x \in \mathbb{X}$ equals the proportion of points in \hat{X} which are closer to x than to any other point in \mathbb{X} . So from each model X_k we obtained a discrete mm-space $(\mathbb{X}_k, d^{(k)}, \nu^{(k)})$. Then we computed a matrix $\mathbf{D}_1 = ((d_{ij}))$ such that $d_{ij} = D_1(\mathbb{X}_i, \mathbb{X}_j)$ (that is, we fixed $p = 1$) where $1 \leq i < j \leq N$ and $N = 72$ (all models had between 9 and 11 poses). See Figure 1 for a graphical representation of \mathbf{D}_1 .

In order to evaluate the discriminative power contained in \mathbf{D}_1 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 \mathbf{D}_1) of the remaining shapes. By simple comparison between the class predicted by the classifier and the actual class to which the shape belongs we thus obtain an estimate for the probability of mis-classification $P_e(\mathbf{D}_1)$. We then repeat this procedure for 10000 random choices of the training set. Using the same randomized procedure we obtain an estimate of the *confusion matrix* C for this problem. That is, C_{ij} equals the probability that the classifier will assign class j to a shape when the actual class was i . We also evaluated the performance of FLB_1 using this method. We obtained $P_e(\mathbf{D}_1) = 0.025$ and $P_e(\text{FLB}_1) = 0.141$. Refer to Figure 1 for more details.

9. Conclusions

We have introduced a modification and expansion of the original Gromov-Hausdorff notion of distance between metric spaces which takes into account probability measures defined on measurable subsets of these metric spaces. This new definition allows for a discretization which is more natural and more general than previous approaches, [MS05, BBK06]. In addition to this, several previous approaches to the problem of Shape Matching/Comparison become inter-related when put into the common framework we introduce. Finally, computational experiments on a database of shapes were presented to exemplify the applicability of our ideas. Further developments such as the extension of the ideas here presented to *partial* shape matching will be reported elsewhere.

References

- [BBI01] BURAGO D., BURAGO Y., IVANOV S.: *A Course in Metric Geometry*, vol. 33 of *AMS Graduate Studies in Math.* American Mathematical Society, 2001.
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM Journal on Scientific Computing* 28, 5 (2006), 1812–1836.

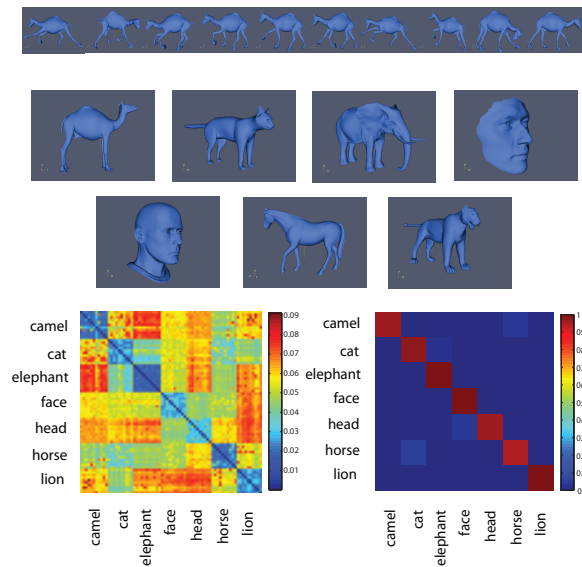


Figure 1: Top Row: As an example we show all the 10 poses of the camel model shape. Second and Third row show one pose of each of the shapes in the database. Last row, left: We show D_1 , see the text for more details. Right: Estimated confusion matrix for the 1-nearest neighbor classification problem described in the text.

- [BK04] BOUTIN M., KEMPER G.: On reconstructing n -point configurations from the distribution of distances or areas. *Adv. in Appl. Math.* 32, 4 (2004), 709–735.
- [CG99] COHEN S. D., GUIBAS L. J.: The earth mover's distance under transformation sets. In *ICCV (2)* (1999), pp. 1076–1083.
- [Dud02] DUDLEY R. M.: *Real analysis and probability*, vol. 74 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2002. Revised reprint of the 1989 original.
- [EK03] (ELBAZ) A. E., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 10 (2003), 1285–1295.
- [Fro90] FROSINI P.: A distance for similarity classes of submanifolds of euclidean space. *Bull. Austral. Math. Soc.* 42:3 (1990), 407–416.
- [Gio] GIORGETTI N.: Glpkmex: A matlab mex interface for the glpk library. <http://www.dii.unisi.it/cohes/tools/mex/downloads/glpkmex/index.html>.
- [Gro99] GROMOV M.: *Metric Structures for Riemannian and Non-Riemannian Spaces*, vol. 152 of *Progress in Mathematics*. Birkhäuser, Boston, US, 1999.
- [HK03] HAMZA A. B., KRIM H.: Geodesic object representation and recognition. In *Lecture Notes in Computer Science*, vol. 2886. Springer, 2003, pp. 378–387.

- [HKR93] HUTTENLOCHER D. P., KLANDERMAN G. A., RUCKLIDGE W. J.: Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15:9 (1993).
- [HM04] HOFFMAN J. R., MAHLER R. P. S.: Multitarget miss distance via optimal assignment. *IEEE Transactions on Systems, Man, and Cybernetics, Part A* 34, 3 (2004), 327–336.
- [KV05] KLEIN O., VELTKAMP R. C.: Approximation algorithms for computing the earth mover's distance under transformations. In *ISAAC (2005)*, pp. 1019–1028.
- [Löf04] LÖFBERG J.: Yalmip : A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference* (Taipei, Taiwan, 2004).
- [LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *International Conference of Computer Vision (ICCV)* (October 2005), vol. 2, pp. 1482 – 1489.
- [Lue03] LUENBERGER D. G.: *Linear and nonlinear programming*, second ed. Kluwer Academic Publishers, Boston, MA, 2003.
- [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.
- [Mum91] MUMFORD D.: Mathematical theories of shape: Do they model perception? *SPIE 1570* (1991), 2–10.
- [OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Trans. Graph.* 21, 4 (2002), 807–832.
- [PW94] PARDALOS P. M., WOLKOWICZ H. (Eds.): *Quadratic assignment and related problems*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 16. American Mathematical Society, Providence, RI, 1994. Papers from the workshop held at Rutgers University, New Brunswick, New Jersey, May 20–21, 1993.
- [RTG00] RUBNER Y., TOMASI C., GUIBAS L. J.: The earth mover's distance as a metric for image retrieval. *Int. J. Comput. Vision* 40, 2 (2000), 99–121.
- [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.
- [SP] SUMNER R. W., POPOVIC J.: Mesh data from deformation transfer for triangle meshes. <http://people.csail.mit.edu/sumner/research/deftransfe>.
- [Stu06] STURM K.-T.: On the geometry of metric measure spaces. I. *Acta Math.* 196, 1 (2006), 65–131.
- [Vil03] VILLANI C.: *Topics in optimal transportation*,

vol. 58 of *Graduate Studies in Mathematics*. American
Mathematical Society, Providence, RI, 2003.