

Comparing Point Clouds

Facundo Mémoli and Guillermo Sapiro

University of Minnesota
{memoli,guille}@ece.umn.edu

Abstract

Point clouds are one of the most primitive and fundamental surface representations. A popular source of point clouds are three dimensional shape acquisition devices such as laser range scanners. Another important field where point clouds are found is in the representation of high-dimensional manifolds by samples. With the increasing popularity and very broad applications of this source of data, it is natural and important to work directly with this representation, without having to go to the intermediate and sometimes impossible and distorting steps of surface reconstruction. A geometric framework for comparing manifolds given by point clouds is presented in this paper. The underlying theory is based on Gromov-Hausdorff distances, leading to isometry invariant and completely geometric comparisons. This theory is embedded in a probabilistic setting as derived from random sampling of manifolds, and then combined with results on matrices of pairwise geodesic distances to lead to a computational implementation of the framework. The theoretical and computational results here presented are complemented with experiments for real three dimensional shapes.

1. Introduction

Point clouds are one of the most primitive and fundamental manifold representations. One of the most popular sources of point clouds are 3D shapes acquisition devices, such as laser range scanners, with applications in many disciplines. These scanners provide in general raw data in the form of (noisy) unorganized point clouds representing surface samples. With the increasing popularity and very broad applications of this source of data, it is natural and important to work directly with this representation, without having to go to the intermediate step of fitting a surface to it (step that can add computational complexity and introduce errors). See for example [3, 10, 12, 19, 28, 29, 35] for a few of the recent works with this type of data, as well as all the papers in the recent June 2004 meeting on point clouds at ETH. Point clouds can also be used as primitives for visualization, e.g., [4, 19, 38], as well as for editing [43].

Another important field where point clouds are found is in the representation of high-dimensional manifolds by samples (see for example [2, 23, 39]). This type of high-dimensional and general co-dimension data appears in almost all disciplines, from computational biology to image analysis to financial data. Due to the extremely high dimen-

sionality in this case, it is impossible to perform manifold reconstruction, and the task needs to be performed directly on the raw data, meaning the point cloud.

The importance of this type of shape representation is leading to a recent increase in the fundamental study of point clouds [1, 2, 8, 11, 15, 31, 32, 39] (see also the papers mentioned in the first paragraph and references therein). The goal of this work, inspired in part by [13] and the tools developed in [31, 39], is to develop a theoretical and computational framework to compare shapes (sub-manifolds of \mathbb{R}^d) represented as point clouds.

As we have mentioned, a variety of objects can be represented as point clouds in \mathbb{R}^d . One is often presented with the problem of deciding whether two of those point clouds, and their corresponding underlying objects or manifolds, represent the same geometric structure or not (*object recognition and classification*). We are concerned with questions about the underlying unknown structures (objects), which need to be answered based on discrete and finite measures taken between their respective point clouds. In greater generality, we wonder what is the structural information we can gather about the object itself by exploring the point cloud which represents it.

Multidimensional scaling (MDS) for example has been used to approach in part this general problem of object recognition. Procedures based on MDS require that one first computes the interpoint distance matrix for all the members of the point cloud (or for a representative selected sub-set of them). If one is interested in comparing two different objects, the problem is reduced to a comparison between the corresponding interpoint distance matrices. If the distance we use is the Euclidean one, these matrices only provide information about their rigid similarity, and (assuming they have the same size) if they are equal (up to a permutations of the indices of all elements), we can only conclude that there exists a rigid isometry (rotation, reflection, translation) from one point cloud to the other. After adding compactness considerations, we can also say something about the true underlying (sampled) objects. Being a bit more rigorous, let the point clouds $\mathcal{P}_i \subset S_i$ be ε_i -coverings of the surfaces S_i in \mathbb{R}^3 , for $i = 1, 2$ (this will be formally defined below). Then assuming there exists a rigid isometry $\tau : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that $\tau(\mathcal{P}_1) = \mathcal{P}_2$, we can bound the Hausdorff distance (which we will also formally define below) between $\tau(S_1)$ and S_2 as follows: $d_{\mathcal{H}}(\tau(S_1), S_2) \leq d_{\mathcal{H}}(\tau(S_1), \tau(\mathcal{P}_1)) + d_{\mathcal{H}}(\tau(\mathcal{P}_1), \mathcal{P}_2) + d_{\mathcal{H}}(\mathcal{P}_2, S_2) = d_{\mathcal{H}}(S_1, \mathcal{P}_1) + d_{\mathcal{H}}(\tau(\mathcal{P}_1), \mathcal{P}_2) + d_{\mathcal{H}}(\mathcal{P}_2, S_2) \leq \varepsilon_1 + 0 + \varepsilon_2$. And of course the same kind of bound holds for the Hausdorff distance between the points clouds once we assume the underlying continuous objects are rigidly isometric, see §2.1 below, where we show that rigid isometries are also addressed with our approach.

If S_1 and S_2 happen to be isometric, thereby allowing for bends and not just rigid transformations, we wonder whether we will be able to detect this by looking at (finite) point clouds \mathcal{P}_i sampled from each S_i . This problem is much harder to tackle. We approach this problem through a probabilistic model, in part because in principle, there might exist even for the same object, two different samplings that look quite dissimilar (under the discrete measures we can cope with computationally), for arbitrarily fine scales (see below).

With the help of the theory presented here we recast these considerations in a rigorous framework and address the case where the distances considered to characterize each point cloud (object) are more general. We concentrate on the situation when we know the existence of an intrinsic notion of distance for each object we sample. For the applications of isometric invariant shape (surfaces) recognition, one must consider the distance as measured by paths constrained to travel on the surface of the objects, better referred to as *geodesic distance*. These have been used in [13] for bending invariant recognition in 3D (the theoretical foundations here developed include a justification of their work) and in [15, 39] to detect intrinsic surface dimensionality. This intrinsic framework not only has applications for the recognition of articulated objects for example, but also leads to comparing manifolds in a complete geometric way and without being influenced by the embedding space (and being as men-

tioned above, rigid isometrics just a particular case covered by our results).

In this paper, the fundamental approach used for isometric invariant recognition is derived then from the *Gromov-Hausdorff distance* [17], which we now present. If two sets (objects) X and Y are subsets of a common bigger metric space (Z, d_Z) , and we want to compare X to Y in order to decide whether they are/represent the same object or not, then an idea one might come up with very early on is that of computing the *Hausdorff distance* between them (see for example [9, 21] for an extensive use of this for shape statistics and image comparison):

$$d_{\mathcal{H}}^Z(X, Y) := \max(\sup_{x \in X} d_Z(x, Y), \sup_{y \in Y} d_Z(y, X))$$

But, what happens if we want to allow for certain deformations to occur and still decide that the manifolds are the same? More precisely, we are interested in being able to find a distance between metric spaces that is *blind* to isometric transformations (“bends”). This will permit a truly geometric comparison between the manifolds, independently of their embedding and bending position. Following [17], we introduce the *Gromov-Hausdorff distance* between Metric Spaces

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

where $f : X \rightarrow Z$ and $g : Y \rightarrow Z$ are *isometric embeddings* (distance preserving) into the metric space Z . It turns out that this measure of metric proximity between metric spaces is well suited for our problem at hand and will allow us to give a formal framework to address the isometric shape recognition problem (for point cloud data). However, this notion of distance between metric spaces encodes the “metric” disparity between the metric spaces, at first glance, in a computationally impractical way. We derive below new results that connect this notion of disparity with other more computationally appealing expressions.

Since we have in mind specific applications and scenarios such as those described above, and in particular surfaces and submanifolds of some Euclidean space \mathbb{R}^d , we assume that we are given as input points *densely* sampled from the metric space (surface, manifold). This will manifest itself in many places in the theory described below. We will present a way of computing a discrete approximation (or bound) to $d_{\mathcal{GH}}(\cdot, \cdot)$ based on the metric information provided by these point clouds. Due to space limitations, the proofs are omitted and are reported elsewhere (www.ima.umn.edu, June/July 2004 reports).

2. Theoretical Foundations

This section covers the fundamental theory behind the bending invariant recognition framework we develop. We use basic concepts of metric spaces, see for example [24] for a de-

tailed treatment of this and [5, 17, 20, 25, 36, 37] for proofs of Proposition 1 below.

Definition 1 (Metric Space) A set M is a metric space if for every pair of points $x, y \in M$ there is a well defined function $d_M(x, y)$ whose values are non-negative real numbers, such that (a) $d_M(x, y) = 0 \Leftrightarrow x = y$, and (b) $d_M(x, y) \leq d_M(y, z) + d_M(z, x)$ for any $x, y, z \in M$. We call $d_M : M \times M \rightarrow \mathbb{R}^+ \cup \{0\}$ the metric or distance. For clarity we will specify a metric space as the pair (M, d_M) .

Definition 2 (Covering) For a point x in the metric space (X, d_X) and $r > 0$, we will denote by $B_X(x, r)$ the set $\{z \in X : d_X(x, z) < r\}$. For a subset A of X , we use the notation $B_X(A, r) = \cup_{a \in A} B_X(a, r)$. We say that a set $C \subset X$ is an R -covering of X if $B_X(C, R) = X$. We will also frequently say that the set A is a n -covering of X if A constitutes, for some $r > 0$, a covering of X by n -balls with centers in points of A .

Definition 3 (Isometry) We say the metric spaces (X, d_X) and (Y, d_Y) are isometric when there exists a bijective mapping $\phi : X \rightarrow Y$ such that $d_X(x_1, x_2) = d_Y(\phi(x_1), \phi(x_2))$ for all $x_1, x_2 \in X$. Such a ϕ is an isometry between (X, d_X) and (Y, d_Y) .

Proposition 1

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

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

2. If $d_{\mathcal{G}\mathcal{H}}(X, Y) = 0$ 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 $d_{\mathcal{G}\mathcal{H}}(X, \{x_1, \dots, x_n\}) \leq R$.
4. For compact metric spaces (X, d_X) and (Y, d_Y) , $\frac{1}{2} |\mathbf{R}(X) - \mathbf{R}(Y)| \leq d_{\mathcal{G}\mathcal{H}}(X, Y) \leq \frac{1}{2} \max(\mathbf{D}(X), \mathbf{D}(Y))$, where $\mathbf{R}(X) := \min_{x \in X} \max_{x' \in X} d_X(x, x')$ and $\mathbf{D}(X) := \max_{x, x' \in X} d_X(x, x')$ stand for the Circumradius and Diameter of the metric space X , respectively.
5. For bounded metric spaces (X, d_X) and (Y, d_Y) ($x \in X, y \in Y$),

$$d_{\mathcal{G}\mathcal{H}}(X, Y) = \inf_{\substack{\phi : X \rightarrow Y \\ \psi : Y \rightarrow X}} \sup_{x, y} \frac{1}{2} |d_X(x, \psi(y)) - d_Y(y, \phi(x))|$$

From these properties, we can easily prove the following important result:

Corollary 1 Let X and Y be compact metric spaces. Let moreover \mathbb{X}_m be a r -covering of X (consisting of m points) and $\mathbb{Y}_{m'}$ be a r' -covering of Y (consisting of m' points). Then

$$|d_{\mathcal{G}\mathcal{H}}(X, Y) - d_{\mathcal{G}\mathcal{H}}(\mathbb{X}_m, \mathbb{Y}_{m'})| \leq r + r'$$

We can then say that if we could compute $d_{\mathcal{G}\mathcal{H}}(\cdot)$ for discrete metric spaces which are dense enough samplings of the continuous underlying ones, that number would be a good approximation to what happens between the continuous spaces. Currently, there is no computationally efficient way to directly compute $d_{\mathcal{G}\mathcal{H}}(\cdot)$ between discrete metric

spaces in general. This forces us to develop a roundabout path, see §2.2 ahead. Before going into the general case, we discuss next the application of our framework to a simpler but important case.

2.1. Intermezzo: The Rigid Isometries Case

When we are trying to compare two subsets X and Y of a larger metric space Z , the situation is less complex. The Gromov-Hausdorff distance boils down to a somewhat simpler Hausdorff distance between the sets. In more detail, one must compute $d_{\mathcal{G}\mathcal{H}}^{Z, rigid}(X, Y) := \inf_{\Phi} d_{\mathcal{H}}^Z(X, \Phi(Y))$, where $\Phi : Z \rightarrow Z$ ranges over all self-isometries of Z . If we know an efficient way of computing $\inf_{\Phi} d_{\mathcal{H}}^Z(X, \Phi(Y))$, then this particular shape recognition problem is well posed for Z , in view of Corollary 1, as soon as we can give guarantees of coverage. This can be done in the case of submanifolds of \mathbb{R}^d by imposing a probabilistic model on the samplings \mathbb{X}_m of the manifolds, and a bound on the curvatures of the family of manifolds. In more detail we can show that $\mathbf{P}(d_{\mathcal{H}}^{\mathbb{R}^d}(X, \mathbb{X}_m) > \delta_m) \simeq \frac{1}{\ln m}$ as $m \uparrow \infty$, for $\delta_m \geq \left(\frac{\ln m}{m}\right)^{1/k}$, where k is the dimension of X , see Section §3.2.

In the case of surfaces in $Z = \mathbb{R}^3$, Φ sweeps all rigid isometries, and there exist good algorithms which can actually solve the problem approximately. For example, in [16] the authors report an algorithm which for any given $0 < \alpha < 1$ can find $\widehat{\Phi}_\alpha$ such that $d_{\mathcal{H}}^{\mathbb{R}^3}(\mathbb{X}_m, \widehat{\Phi}_\alpha(\mathbb{Y}_{m'})) \leq (8 + \alpha) \inf_{\Phi} d_{\mathcal{H}}^{\mathbb{R}^3}(\mathbb{X}_m, \Phi(\mathbb{Y}_{m'}))$, with complexity $O(s^4 \log s)$ where $s = \max(m, m')$. This computational result, together with our theory, makes the problem of surface recognition (under rigid motions) well posed and well justified. In fact, just using (an appropriate version of) Corollary 1 and the triangle inequality, we obtain a bound between the distance we want to estimate $d_{\mathcal{H}}^{\mathbb{R}^3, rigid}(X, Y)$ and the observed (computable) value $d_{\mathcal{H}}^{\mathbb{R}^3}(\mathbb{X}_m, \widehat{\Phi}_\alpha(\mathbb{Y}_{m'}))$, having $d_{\mathcal{H}}^{\mathbb{R}^3, rigid}(X, Y) - (r + r') \leq d_{\mathcal{H}}^{\mathbb{R}^3}(\mathbb{X}_m, \widehat{\Phi}_\alpha(\mathbb{Y}_{m'})) \leq 10 \left(d_{\mathcal{H}}^{\mathbb{R}^3, rigid}(X, Y) + (r + r')\right)$. This bound gives a formal justification for the surface recognition problem from point samples, showing that it is well posed. To the best of our knowledge, this is the first time that such formality is shown for this very important problem, both in the particular case just shown and in the general one addressed next.

2.2. The General Recognition Case

The theory introduced by Gromov permits to address the concept of (metric) proximity between metric spaces. When dealing with discrete metric spaces, as those arising from samplings or coverings of continuous ones, it is convenient to introduce a distance between them, which ultimately is the one we compute for point clouds, see §3.4 ahead. For discrete metric spaces (both of cardinality n) ($\mathbb{X} = \{x_1, \dots, x_n\}, d_{\mathbb{X}}$) and ($\mathbb{Y} = \{y_1, \dots, y_n\}, d_{\mathbb{Y}}$) we define

the distance:

$$d_J(\mathbb{X}, \mathbb{Y}) := \min_{\pi \in \mathcal{P}_n} \max_{1 \leq i, j \leq n} \frac{1}{2} |d_{\mathbb{X}}(x_i, x_j) - d_{\mathbb{Y}}(y_{\pi_i}, y_{\pi_j})| \quad (1)$$

where \mathcal{P}_n stands for the set of all $n \times n$ permutations of $\{1, \dots, n\}$. A permutation π provides the correspondence between the points in the sets, and $|d_{\mathbb{X}}(x_i, x_j) - d_{\mathbb{Y}}(y_{\pi_i}, y_{\pi_j})|$ gives the pairwise distance/disparity once this correspondence has been assumed. It is evident that one has $d_{\mathcal{G}\mathcal{H}}(\mathbb{X}, \mathbb{Y}) \leq d_J(\mathbb{X}, \mathbb{Y})$, by virtue of Property 5 from Proposition 1. Moreover, we easily derive the following easy result, whose usefulness will be made evident in §3.

Corollary 2 Let (X, d_X) and (Y, d_Y) be compact metric spaces. Let $\mathbb{X} = \{x_1, \dots, x_n\} \subset X$ and $\mathbb{Y} = \{y_1, \dots, y_n\} \subset Y$, such that $B_X(\mathbb{X}, R_X) = X$ and $B_Y(\mathbb{Y}, R_Y) = Y$ (the point clouds provide R_X and R_Y coverings respectively). Then

$$d_{\mathcal{G}\mathcal{H}}(X, Y) \leq R_X + R_Y + d_J(\mathbb{X}, \mathbb{Y}) \quad (2)$$

with the understanding that $d_{\mathbb{X}} = d_X|_{\mathbb{X} \times \mathbb{X}}$ and $d_{\mathbb{Y}} = d_Y|_{\mathbb{Y} \times \mathbb{Y}}$.

Remark 1 This result tells us that if we manage to find coverings of X and Y for which the distance d_J is small, then if the radii defining those coverings are also small, the underlying manifolds X and Y sampled by these point clouds must be close in a metric sense. Another way of interpreting this is that we will never see a small value of $d_J(\mathbb{X}, \mathbb{Y})$ whenever $d_{\mathcal{G}\mathcal{H}}(X, Y)$ is big, a simple statement with practical value, since we will be looking at values of d_J , which depend on the point clouds. This is because, in contrast with $d_{\mathcal{G}\mathcal{H}}(\cdot, \cdot)$, the distance d_J is (approximately) computable from the point clouds, see §3.4.

regarding coverings of metric spaces. Given a metric space (X, d_X) , the discrete subset $N_{X,n}^{(r,s)}$ denotes a set of points $\{x_1, \dots, x_n\} \subset X$ such that (1) $B_X(N_{X,n}^{(r,s)}, r) = X$, and (2) $d_X(x_i, x_j) \geq s$ whenever $i \neq j$. In other words, the set provides a coverage and the points in the set are not too close to each other (the coverage is efficient). (Similar sampling conditions are common in the computational geometry literature, e.g., works by Amenta, Dey, Boissonnat, and others.)

Remark 2 For each $r > 0$ denote by $N(r, X)$ the minimum number of closed balls of radii r needed to cover X . Then, ([36], Chapter 10), one can actually show that the class $(\mathcal{M}, d_{\mathcal{G}\mathcal{H}}(\cdot, \cdot))$ of all compact metric spaces X whose covering number $N(r, X)$ are bounded for all (small) positive r by a function on the interval, $N : (0, C_1) \rightarrow (0, \infty)$, is *totally bounded*. This means that given $\rho > 0$, there exist a finite positive integer $k(\rho)$ and compact metric spaces $X_1, \dots, X_{k(\rho)} \in \mathcal{M}$ such that for any $X \in \mathcal{M}$ one can find $i \in \{1, \dots, k(\rho)\}$ such that $d_{\mathcal{G}\mathcal{H}}(X, X_i) \leq \rho$. This is very interesting from the point of view of applications since it allows to make the classification problem of metric spaces in a well posed and justified way. For example, in a system of storage/retrieval of faces/information manifolds, this concept permits the design of a clustering procedure for the objects.

The following Proposition will also be fundamental for our computational framework in §3, leading us to work with point clouds.

Proposition 2 ([17]) Let (X, d_X) and (Y, d_Y) be any pair of given compact metric spaces and let $\eta = d_{\mathcal{G}\mathcal{H}}(X, Y)$. Also, let $N_{X,n}^{(r,s)} = \{x_1, \dots, x_n\}$ be given. Then, given $\alpha > 0$ there exist points $\{y_1^\alpha, \dots, y_n^\alpha\} \subset Y$ such that

1. $d_J(N_{X,n}^{(r,s)}, \{y_1^\alpha, \dots, y_n^\alpha\}) \leq (\eta + \alpha)$
2. $B_Y(\{y_1^\alpha, \dots, y_n^\alpha\}, r + 2(\eta + \alpha)) = Y$
3. $d_Y(y_i^\alpha, y_j^\alpha) \geq s - 2(\eta + \alpha)$ for $i \neq j$.

Remark 3 This proposition first tells us that if the metric spaces happen to be sufficiently close in a metric sense, then given a s -separated covering on one of them, one can find a (s' -separated) covering in the other metric space such that d_J between those coverings (point clouds) is also small. This, in conjunction with Remark 1, proves that in fact our goal of trying to determine the metric similarity of metric spaces based on discrete observations of them is, so far, a (theoretically) well posed problem.

Since by Tychonoff's Theorem the n -fold product space $Y \times \dots \times Y$ is compact, if $s - 2\eta \geq c > 0$ for some constant c , by passing to the limit along the subsequences of $\{y_1^\alpha, \dots, y_n^\alpha\}_{\{\alpha > 0\}}$ as $\alpha \downarrow 0$ (if needed) above one can assume the existence of a set of different points $\{\bar{y}_1, \dots, \bar{y}_n\} \subset Y$ such that $d_J(\{\bar{y}_1, \dots, \bar{y}_n\}, N_{X,n}^{(r,s)}) \leq \eta$, $\min_{i \neq j} d_Y(\bar{y}_i, \bar{y}_j) \geq s - 2\eta > 0$, and $B_Y(\{\bar{y}_1, \dots, \bar{y}_n\}, r + 2\eta) = Y$.

Since we are given the finite sets of points sampled from each metric space, the existence of $\{\bar{y}_1, \dots, \bar{y}_n\}$ guaranteed by Proposition 2 doesn't seem to make our life a lot easier, those points could very well not be contained in our given finite datasets \mathbb{X}_m and \mathbb{Y}_m . The simple idea of using a triangle inequality (with metric d_J) to deal with this does not work in principle, since one can find, for the same underlying space, two nets whose d_J distance is not small, see [6, 30]. Let us explain this in more detail. Assume that as input we are given two finite sets of points \mathbb{X}_m and \mathbb{Y}_m on two metric spaces, X and Y , respectively, which we assume to be isometric. Then the results above ensure that for any given $N_{X,n}^{(r,s)} \subset \mathbb{X}_m$ there exists a $N_{Y,n}^{(r,s)} \subset Y$ such that $d_J(N_{X,n}^{(r,s)}, N_{Y,n}^{(r,s)}) = 0$. However, it is clear that this $N_{Y,n}^{(r,s)}$ has no reason to be contained in the given point cloud \mathbb{Y}_m . The obvious idea would be to try to rely on some kind of independence property on the sample which represents a given metric space, namely that for any two different covering nets N_1 and N_2 (of the same cardinality and with small covering radii) of X the distance $d_J(N_1, N_2)$ is also small. If this were granted, we could proceed as follows: $d_J(N_{X,n}^{(r,s)}, N_{Y,n}^{(\hat{r}, \hat{s})}) \leq d_J(N_{X,n}^{(r,s)}, N_{Y,n}^{(r,s)}) + d_J(N_{Y,n}^{(\hat{r}, \hat{s})}, N_{Y,n}^{(r,s)}) = 0 + \text{small}(r, \hat{r}, s, \hat{s})$, where $\text{small}(r, \hat{r}, s, \hat{s})$ is small number depending only on r, \hat{r}, s and \hat{s} . The property we fancy to rely upon was a conjecture proposed by Gromov in [18] (see also [40]) and disproved [6, 30]. Their counterexamples are for separated nets in \mathbb{Z}^2 . It is not known

whether we can construct counterexamples for compact metric spaces, or if there exists a characterization of a family of n -points separated nets of a given compact metric space such that any two of them are at a small d_j -distance which can be somehow controlled with n . A first step towards this is the density condition introduced in [7].

If counterexamples do not exist for compact metric spaces, then the above inequality should be sufficient. Without assuming this, we give below an argument which tackles the problem in a probabilistic way. In other words, we use a probabilistic approach to bound d_j for two different samples from a given metric space. For this, we pay the price, for some applications, of assuming the existence of a measure which comes with our metric space. On the other hand, probabilistic frameworks are natural for noisy random samples of manifolds as obtained in real applications.

2.3. A Probabilistic Setting for Submanifolds of \mathbb{R}^d

We now limit ourself to smooth submanifolds of any \mathbb{R}^d , although the work can be extended to more general metric spaces (once a notion of uniform probability measure is introduced).

Let Z be a smooth and compact submanifold of \mathbb{R}^d with intrinsic (geodesic) distance function $d_Z(\cdot, \cdot)$. We can now speak more freely about points $\{z_i\}_{i=1}^m$ sampled uniformly from X . For any measurable $C \subset Z$, $\mathbf{P}(z_i \in C) = \frac{\mathbf{a}(C)}{\mathbf{a}(Z)}$, where $\mathbf{a}(B)$ denotes the area of the measurable set $B \subset Z$. This uniform distribution can be replaced by other distributions, e.g., that adapt to the geometry of the underlying surface, and the framework here developed can be extended to those as well.

Let $\mathbb{Z} = \{z_1, \dots, z_n\}$ and $\mathbb{Z}' = \{z'_1, \dots, z'_n\}$ be two discrete subsets of Z (two point clouds). For any permutation $\pi \in \mathcal{P}_n$ and $i, j \in \{1, \dots, n\}$, $|d_Z(z_i, z_j) - d_Z(z'_{\pi_i}, z'_{\pi_j})| \leq d_Z(z_i, z'_{\pi_i}) + d_Z(z_j, z'_{\pi_j})$ and therefore we have

$$d_{\mathbb{B}}^Z(\mathbb{Z}, \mathbb{Z}') := \min_{\pi \in \mathcal{P}_n} \max_k d_Z(z_k, z'_{\pi_k}) \geq d_j(\mathbb{Z}, \mathbb{Z}') \quad (3)$$

This is known as the *Bottleneck Distance* between \mathbb{Z} and \mathbb{Z}' , both being subsets of Z . This is a possible way of measuring distance between two different samples of the same metric space, see the work [34]. For its application in Point Matching see [22] and references therein.

Instead of dealing with (3) deterministically, after imposing conditions on the underlying metric space Z , we derive probabilistic bounds for the left hand side. We also make evident that by suitable choices of the relations among the different parameters in the sampling process, this probability can be chosen at will. This result is then used to bound the distance d_j between two point cloud samples of a given metric space, thereby leading to the bound for (a quantity related to) $d_j(N_{X,n}^{(r,s)}, N_{Y,n}^{(r,s)})$ without assuming any kind of proximity of the nets (and from this, the bounds on the original Gromov-Hausdorff distance). We consider \mathbb{Z} to be fixed,

and we assume $\mathbb{Z}' = \{z'_1, \dots, z'_n\}$ to be chosen from a set $\mathbb{Z}_m \subset Z$ consisting of $m \gg n$ i.i.d. points sampled uniformly from Z . We introduce the *Voronoi diagram* $\mathcal{V}(\mathbb{Z})$ on Z , determined by the points in \mathbb{Z} (see for example [27]).

The i -th Voronoi cell of the Voronoi diagram defined by $\mathbb{Z} = \{z_1, \dots, z_n\} \subset Z$ is given by $A_i := \{z \in Z : d_Z(z_i, z) < \min_{j \neq i} d_Z(z_j, z)\}$. We then have $Z = \bigcup_{k=1}^n \overline{A_k}$.

We first want to find, amongst points in \mathbb{Z}_m , n different points $\{z_{i_1}, \dots, z_{i_n}\}$ such that each of them falls inside one Voronoi cell, $\{z_{i_k} \in A_k \text{ for } k = 1, \dots, n\}$. We provide lower bounds for $\mathbf{P}(\#(A_k \cap \mathbb{Z}_m) \geq 1, 1 \leq k \leq n)$, the probability of this happening. We can see the event as if we *collected* points inside all the Voronoi cells, a case of the *Coupon Collecting Problem*, see [14]: we buy merchandise at a coupon-giving store until we have collected all possible types of coupons. The next Lemma presents the basic results we need about this concept [41].

Lemma 1 (Coupon Collecting) If there are n different coupons one wishes to collect, such that the probability of seeing the k -th coupon is p_k (let $\vec{p} = (p_1, \dots, p_n)$), and one obtains samples of all of them in an independent way then: The probability $P_{\vec{p}}(n, r)$ of having collected all n coupons after r trials is given by

$$P_{\vec{p}}(n, r) = 1 - S_n \left(\sum_{j=2}^n (-1)^j \binom{n}{k=j} p_k \right)^n \quad (4)$$

where the symbol S_n means that we consider all possible combinations of the n indices in the expression being evaluated. (For example $S_3((p_1 + p_2)^r) = (p_1 + p_2)^r + (p_1 + p_3)^r + (p_2 + p_3)^r$.)

This result is used to prove the following fundamental probability bounds:

Theorem 1 Let (Z, d_Z) be a smooth compact submanifold of \mathbb{R}^d . Given a covering $N_{Z,n}^{(r,s)}$ of Z and a number $p \in (0, 1)$, there exists a positive integer $m = m_n(p)$ such that if $\mathbb{Z}_m = \{z_k\}_{k=1}^m$ is a sequence of *i.i.d.* points sampled uniformly from Z , with probability p one can find a set of n different indices $\{i_1, \dots, i_n\} \subset \{1, \dots, m\}$ with $d_{\mathbb{B}}^Z(N_{Z,n}^{(r,s)}, \{z_{i_1}, \dots, z_{i_n}\}) \leq r$.

This result can also be seen the other way around: for a given m , the probability of finding the aforementioned subset in \mathbb{Z}_m is given by (4), for \vec{p}_Z defined as follows: $p_Z^i = \mathbf{a}(A_i)/\mathbf{a}(Z)$, where A_i is the i -th Voronoi cell corresponding to $N_{Z,n}^{(r,s)}$, $1 \leq i \leq n$. Moreover, since for $\hat{z}_k \in N_{Z,n}^{(r,s)}$ $B_Z(\hat{z}_k, \frac{s}{2}) \subset A_k$ then one can lower bound all components of \vec{p}_Z . In practise one could use as a rule of thumb $m \simeq n \ln n$ which is the mean waiting time (in the equiprobable case) until all ‘‘coupons’’ are collected, [14].

Corollary 3 Let X and Y compact submanifolds of \mathbb{R}^d . Let $N_{X,n}^{(r,s)}$ be a covering of X with separation s such that for some positive constant c , $s - 2d_{\mathcal{G}\mathcal{H}}(X, Y) > c$. Then, given any

number $p \in (0, 1)$, there exists a positive integer $m' = m'_n(p)$ such that if $\mathbb{Y}_{m'} = \{y_k\}_{k=1}^{m'}$ is a sequence of i.i.d. points sampled uniformly from Y , we can find, with probability at least p , a set of n different indices $\{i_1, \dots, i_n\} \subset \{1, \dots, m'\}$ such that $d_{\mathcal{J}}(N_{X,n}^{(r,s)}, \{y_{i_1}, \dots, y_{i_n}\}) \leq 3d_{\mathcal{G}\mathcal{J}\mathcal{C}}(X, Y) + r$.

This concludes the main theoretical foundation of our proposed framework. We have shown that $d_{\mathcal{J}}$ is a good approximation of the Gromov-Hausdorff distance between the point clouds, in a probabilistic sense. Now, we must devise a computational procedure which allows us to actually find the subset $\{y_{i_1}, \dots, y_{i_n}\}$ inside the given point cloud \mathbb{Y}_m when it exists, or at least find it with a large probability. Note that in practise we can only access metric information, that is, interpoint distances. Point positions cannot be assumed to be accessible since that would imply knowing the (isometry) transformation that maps X into Y . A stronger result should take into account possible self-isometries of X (Y), which would increase the probability of finding a net which achieves small $d_{\mathcal{J}}$ distance to the fixed one. Next we present such a computational framework.

3. Computational Foundations

There are a number of issues that must be addressed in order to develop an algorithmic procedure from the theoretical results previously presented. These are now addressed.

3.1. Initial Considerations

In practise we have as input two independent point clouds \mathbb{X}_m and $\mathbb{Y}_{m'}$ each of them composed of i.i.d. points sampled uniformly from X and Y , respectively. We fix a number $n < \min(m, m')$ and construct good coverings $N_{X,n}^{(r,s)}$ of X and $N_{Y,n}^{(r',s')}$ of Y . Actually, r, s, r' and s' all depend on n , and we should choose n such that r and r' are small enough to make our bounds useful, see the additional computations below. Details on how we construct these coverings are provided in Section §3.3.

It is convenient to introduce the following additional notation. For $q \in \mathbb{N}$, let $\{1 : q\}$ denote the set $\{1, \dots, q\}$; also for a set of points $\mathbb{Z}_q = \{z_k\}_{k=1}^q$ and for a set of $1 \leq u \leq q$ indices $I_u = \{i_1, \dots, i_u\} \subset \{1 : q\}$, let $\mathbb{Z}_q[I_u]$ denote the set $\{z_{i_1}, \dots, z_{i_u}\}$.

Corollary 3 suggests that in practise we compute the symmetric expression

$$d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_m) := \max \left(\min_{J_n \subset \{1:m\}} d_{\mathcal{J}}(N_{X,n}^{(r,s)}, \mathbb{Y}_m[J_n]), \min_{I_n \subset \{1:m'\}} d_{\mathcal{J}}(N_{Y,n}^{(r',s')}, \mathbb{X}_m[I_n]) \right) \quad (5)$$

which depends not only on \mathbb{X}_m and $\mathbb{Y}_{m'}$ but also on specified covering nets $N_{X,n}^{(r,s)}$ and $N_{Y,n}^{(r',s')}$. However we prefer to omit the dependence in the list of arguments to keep the notation simpler.

Then, we know that with probability at least $P_{\overline{p}_X}(n, m) \times P_{\overline{p}_Y}(n, m')$ we have (we assume \mathbb{X}_m to be independent from $\mathbb{Y}_{m'}$) $d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_{m'}) \leq 3d_{\mathcal{G}\mathcal{J}\mathcal{C}}(X, Y) + \max(r, r')$. Moreover, in some precise sense $d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_{m'})$ upper bounds $d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m, \mathbb{Y}_{m'})$, something we need to require otherwise we would have solved one problem to gain another, and that implies (Corollary 1) a similar upper bound for $d_{\mathcal{G}\mathcal{J}\mathcal{C}}(X, Y)$.

In fact, for any $I_n \subset \{1 : m\}$

$$\begin{aligned} d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m, \mathbb{Y}_{m'}) &\leq d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m, \mathbb{X}_m[I_n]) + d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m[I_n], \mathbb{Y}_{m'}) \\ &\leq d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m, \mathbb{X}_m[I_n]) + d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m[I_n], N_{Y,n}^{(r',s')}) \\ &\quad + d_{\mathcal{G}\mathcal{J}\mathcal{C}}(N_{Y,n}^{(r',s')}, \mathbb{Y}_{m'}) \\ &\leq d_{\mathcal{J}\mathcal{C}}^X(\mathbb{X}_m, \mathbb{X}_m[I_n]) + d_{\mathcal{J}}(\mathbb{X}_m[I_n], N_{Y,n}^{(r',s')}) + r' \end{aligned}$$

Now, considering I_n such that $d_{\mathcal{J}}(\mathbb{X}_m[I_n], N_{Y,n}^{(r',s')}) = \min_{I_n \subset \{1:m\}} d_{\mathcal{J}}(N_{Y,n}^{(r',s')}, \mathbb{X}_m[I_n])$, we find $d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m, \mathbb{Y}_{m'}) \leq d_{\mathcal{J}\mathcal{C}}^X(\mathbb{X}_m, \mathbb{X}_m[I_n]) + d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_{m'}) + r'$.

Symmetrically, we also obtain for J_n such that $d_{\mathcal{J}}(\mathbb{Y}_m[J_n], N_{X,n}^{(r,s)}) = \min_{J_n \subset \{1:m'\}} d_{\mathcal{J}}(N_{X,n}^{(r,s)}, \mathbb{Y}_{m'}[J_n])$

$$\begin{aligned} d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m, \mathbb{Y}_{m'}) &\leq d_{\mathcal{J}\mathcal{C}}^Y(\mathbb{Y}_{m'}, \mathbb{Y}_{m'}[J_n]) + d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_{m'}) + r \\ \text{Hence,} \quad d_{\mathcal{G}\mathcal{J}\mathcal{C}}(\mathbb{X}_m, \mathbb{Y}_{m'}) &\leq d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_{m'}) + \min \left(d_{\mathcal{J}\mathcal{C}}^X(\mathbb{X}_m, \mathbb{X}_m[I_n]), d_{\mathcal{J}\mathcal{C}}^Y(\mathbb{Y}_{m'}, \mathbb{Y}_{m'}[J_n]) \right) + \max(r, r'). \end{aligned}$$

Let $\Delta_X = d_{\mathcal{J}\mathcal{C}}^X(\mathbb{X}_m, \mathbb{X}_m[I_n])$ and $\Delta_Y = d_{\mathcal{J}\mathcal{C}}^Y(\mathbb{Y}_{m'}, \mathbb{Y}_{m'}[J_n])$. The computational procedure we infer is: *If $d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_{m'})$ is “large” we then know that $d_{\mathcal{G}\mathcal{J}\mathcal{C}}(X, Y)$ must also be “large” with high probability. On the other hand, if $d_{\mathcal{J}}(\mathbb{X}_m, \mathbb{Y}_{m'})$ is “small” and $\min(\Delta_X, \Delta_Y)$ is also “small” then $d_{\mathcal{G}\mathcal{J}\mathcal{C}}(X, Y)$ must also be “small.”*

3.2. Working with Point Clouds

First, all we have is a finite sets of points, point clouds, sampled from each metric space, and all our computations must be based on *these observations* alone. Since we made the assumption of randomness in the sampling (and it also makes sense in general to model the problem in this way, given the way shapes are acquired by a scanner for example), we must relate the number of acquired points to the coverage properties we wish to have. In other words, and following our theory above, we would like to say that given a desired probability p and a radius r , there exists a finite m such that the probability of covering all the metric space with m balls (intrinsic or not) of radius r centered at those m random points is at least p . This kind of characterizations are easy to deal with in the case of submanifolds of \mathbb{R}^d , where the *tuning* comes from the curvature bounds available. For this we follow arguments from [31]. Let Z be a smooth and compact submanifold of \mathbb{R}^d of dimension k . Let $\mathbb{Z}_m \subset Z$ consist of m i.i.d. points uniformly sampled from Z . Let K be an upper bound for the sectional curvatures of Z . Then we can prove

that for a sequence $r_m \rightarrow 0$ such that $r_m \gtrsim \frac{\ln m}{m}$ for large m , $\mathbf{P}\left(d_{\mathcal{J}\zeta}^{\mathcal{R}^d}(Z, \mathbb{Z}_m) > r_m\right) \simeq \frac{1}{\ln m}$.

Then, since one can prove, [31], that for any $z \in Z$, $\delta > 0$ small, $B(z, \delta) \cap Z \subset B_Z(z, C_K \delta)$, for some constant $C_K > 1$ depending only on metric properties of Z (curvatures and diameter), we also find $\mathbf{P}\left(d_{\mathcal{J}\zeta}^Z(Z, \mathbb{Z}_m) > r_m\right) \simeq \frac{1}{\ln m}$.

This relation gives us some guidance regarding how many points we must sample in order to have a certain covering radius, or to estimate the covering radius in terms of m . More precise estimates can be found in the reference mentioned above. The important point to remark is that this kind of relations should hold for the family of shapes we want to work with, therefore, once given bounds on the curvatures and diameters which characterize the family, one can determine a precise probabilistic covering relation for it. We leave the exploitation of this idea for future work.

Given the natural number $n \ll m$ (or eventually $s > 0$), we use the procedure described in §3.3 below to find n -points from \mathbb{Z}_m which constitute a covering of \mathbb{Z}_m of the given cardinality n (or of the given separation s) and of a resulting radius r . We denote this set by $N_{\mathbb{Z}_m, n}^{(r, s)} \subseteq \mathbb{Z}_m$.

3.3. Finding Coverings

In order to find the coverings, we use the well known Farthest Point Sampling (**FPS**) procedure, which we describe next. Suppose we have a dense sampling \mathbb{Z}_m of the smooth and compact submanifold of \mathbb{R}^d (Z, d_Z) as interpreted by the discussion above. We want to simplify our sampling and obtain a well separated covering net of the space. We also want to estimate the covering radius and separation of our net. It is important to obtain subsets which retain as best as possible the metric information contained in the initial point cloud in order to make computational tasks more treatable without sacrificing precision.

We first show a procedure to sample the whole space Z . Fix n the number of points we want to have in our simplified point cloud \mathcal{P}_n . We build \mathcal{P}_n recursively. Given \mathcal{P}_{n-1} , we select $p \in Z$ such that $d_Z(p, \mathcal{P}_n) = \max_{z \in Z} d_Z(z, \mathcal{P}_{n-1})$ (here we consider of course, geodesic distances). There might exist more than one point which achieves the maximum, we either consider all of them or randomly select one and add it to \mathcal{P}_{n-1} . This subsampling procedure has been studied and efficiently implemented in [33] for the case of surfaces represented as point clouds.

Let us now assume that the discrete metric space (\mathbb{Z}_m, d_Z) is a good random sampling of the underlying (Z, d_Z) in the sense that $d_{\mathcal{J}\zeta}(Z, \mathbb{Z}_m) \leq r$ with probability $p_{r, m}$, as discussed in Section §3.2. We then want to simplify \mathbb{Z}_m in order to obtain a set \mathcal{P}_n with n points which is both a good subsampling and a well separated net of Z . We want to use our n sampled points in the best possible way. We are then led to using the construction discussed above. Choose randomly

one point $p_1 \in \mathbb{Z}_m$ and consider $\mathcal{P}_1 = \{p_1\}$. Run the procedure **FPS** until $n - 1$ other points have been added to the set of points. Compute now $r_n = \max_{q \in \mathbb{Z}_m} d_Z(q, \mathcal{P}_n)$. Then, also with probability $p_{r, m}$, \mathcal{P}_n is a $(r + r_n)$ -covering net of Z with separation s_n , the resulting separation of the net. Following this, we now use the notation $N_{\mathbb{Z}_m}^{(r+r_n, s_n)}$.

We use a graph based distance computation following [39], or the exact distance, which can be computed only for certain examples (spheres, planes). We could also use the techniques developed for triangular surfaces in [26], or, being this the optimal candidate, the work on geodesics on (noisy) point clouds developed in [31].

3.4. Additional Implementational Details

In this section we conclude the details on the implementation of the framework here proposed. The first step of the implementation is the computation of $d_{\mathcal{J}}$ and subsequently $d_{\mathcal{J}}$, which from the theory we described before, bounds the Gromov-Hausdorff distance.

We have implemented a simple algorithm. Considering the matrix of pairwise geodesic distances between points of \mathbb{X}_m , we need to determine whether there exists a submatrix of the whole distance matrix corresponding to \mathbb{X}_m which has a small $d_{\mathcal{J}}$ distance to the corresponding matrix of a given $N_{Y, n}^{(r', s')}$. We select this latter net as the result of applying the **FPS** procedure to obtain a subsample consisting of n points, where the first two points are selected to be at maximal distance from each other. To fix notation, let $\mathbb{X}_m = \{x_1, \dots, x_m\}$ and $N_{Y, n}^{(r', s')} = \{y_{j_1}, \dots, y_{j_n}\}$. We then use the following algorithm.

- ($k = 1, 2$) Choose x_{i_1} and x_{i_2} such that $|d_X(x_{i_1}, x_{i_2}) - d_Y(y_{j_1}, y_{j_2})|$ is minimized.
- ($k > 2$) Let $x_{i_{k+1}} \in \mathbb{X}_m$ be such that $e_{k+1}(x_{i_{k+1}}) = \min_{1 \leq i_l \leq m} e_{k+1}(x_{i_l})$ where $e_{k+1}(x_{i_l}) := \max_{1 \leq r \leq k} |d_X(x_{i_l}, x_{i_r}) - d_Y(y_{j_{k+1}}, y_{j_r})|$.

We stop when n points, $\{x_{i_1}, x_{i_2}, \dots, x_{i_n}\}$ have been selected, and therefore a distance submatrix $((d_X(x_{i_u}, x_{i_v})))_{u, v=1}^n$ is obtained. Since we can write $d_{\mathcal{J}}(\{x_{i_1}, \dots, x_{i_n}\}, N_{Y, n}^{(r', s')}) = \frac{1}{2} \max_{1 \leq k \leq n} \max_{1 \leq l \leq k-1} |d_X(x_{i_k}, x_{i_l}) - d_Y(y_{j_k}, y_{j_l})| = \frac{1}{2} \max_{1 \leq k \leq n} e_k(x_{i_r})$ we then see that with our algorithm we are minimizing the error file-wise.

Of course, we now use the same algorithm to compute the other half of $d_{\mathcal{J}}$. This algorithm is not computationally optimal. We are currently studying computational improvements along with error bounds for the results provided by the algorithms.

4. Examples

We now present experiments that confirm the validity of the theoretical and computational framework introduced in previous sections. In the future, we plan to make these experi-

ments more rigorous, including concepts of hypothesis testing. As a simplification, for our experiments we have only computed $d_{\mathcal{F}}$ neglecting the other terms (see §3.1) which would provide a estimative of the Gromov-Hausdorff proximity between the shapes.

We complemented the more complex data (as presented below) with simple shapes: (1) A plane, $P_{\pi} = [-\frac{\pi}{\sqrt{8}}, \frac{\pi}{\sqrt{8}}]^2$ and (2) A sphere, $S = \{x \in \mathbb{R}^d : \|x\| = 1\}$.

We first test our framework when X and Y are isometric. We first consider $X = Y$ and see whether we make the right decision based on the discrete (random) measurements. Let \mathbb{X}_m and \mathbb{Y}_m be two independent sets composed of m independent, uniformly distributed random points on X . In the case of the sphere we generated this uniformly distributed sample points using the method of Muller, see [42]. We consider X to be either the plane P_{π} or the sphere S as defined above. Given n , from \mathbb{X}_m and \mathbb{Y}_m , and using the FPS procedure, we construct $N_{\mathbb{X}_m, n}$ and $N_{\mathbb{Y}_m, n}$ (we omit the supraindices since we won't use the values of covering radius and separation), and look for a metric match inside \mathbb{X}_m and \mathbb{Y}_m , respectively, following the algorithm described in §3.4 for the computation of $d_{\mathcal{F}}(\mathbb{X}_m, \mathbb{Y}_m)$. (Recall that actually $d_{\mathcal{F}}(\mathbb{X}_m, \mathbb{Y}_m)$ depends on n , see its definition (5).) For each dataset we tested for values of $m \in M = \{500, 600, \dots, 2000\}$ and $n \in N = \{5, 10, 15, \dots, 100\}$, and obtained the results reported below. In Table 1 we show the values of $d_{\mathcal{F}}$ for selected values of m and n . As expected, the values of $d_{\mathcal{F}}$ are small compared to $\mathbf{D}(P_{\pi}) = \mathbf{D}(S) = \pi$ (see below for the corresponding values when comparing non-isometric shapes). In Figure 1 (first two figures) we show a pseudocolor representation of the results for $d_{\mathcal{F}}$.

We now proceed to compare shapes that are not isometric, starting with $X = P_{\pi}$ (a plane) and $Y = S$ (a sphere). In this case we expect to be able to detect, based on the finite point clouds, that $d_{\mathcal{F}}$ is large. Table 1 (see also last two figures of Figure 1), shows the results of a simulation in which we compared the sphere S and the plane P_{π} , varying the net sizes and the total number of points uniformly sampled from them. The experiments have been repeated 100 times to produce this table, and the reported values consist of the mean of these 100 tests, as well as their maximum (the corresponding deviation was 1.72×10^{-2}). As expected, the values are larger than when comparing plane against plane or sphere against sphere.

We conclude the experiments with real (more complex) data. We have 4 sets of shapes (the datasets were kindly provided to us by Prof. Kimmel and his group at the Technion), each one with their corresponding bends. We ran the algorithm $N = 6$ times with $n = 70$, $m = 2000$, using the 4 nearest neighbors to compute the geodesic distance using the `isomap` engine. The data description and results are reported in Table 2. We note not only that the technique is able to discriminate between different object, but as expected, it doesn't get confused by bends. Moreover, the distances be-

$n \setminus m$	500	900	1500	1900
5	0.036793	0.015786	0.018160	0.0074027
25	0.041845	0.050095	0.026821	0.031019
45	0.081975	0.042198	0.038990	0.036376
65	0.068935	0.052482	0.035718	0.031512
85	0.077863	0.038660	0.036009	0.036894

$n \setminus m$	500	900	1500	1900
5	0.013282	0.013855	0.010935	0.013558
25	0.082785	0.043617	0.033095	0.033592
45	0.074482	0.067096	0.057161	0.040727
65	0.079456	0.076762	0.049503	0.043405
85	0.083577	0.083344	0.058094	0.054144

$n \setminus m$	500	1000	1500	2000
10	1.839×10^{-1}	1.902×10^{-1}	1.931×10^{-1}	1.942×10^{-1}
25	1.834×10^{-1}	1.908×10^{-1}	1.920×10^{-1}	1.944×10^{-1}
50	1.818×10^{-1}	1.899×10^{-1}	1.925×10^{-1}	1.933×10^{-1}
75	1.873×10^{-1}	1.882×10^{-1}	1.936×10^{-1}	1.939×10^{-1}
100	1.846×10^{-1}	1.913×10^{-1}	1.924×10^{-1}	1.936×10^{-1}

Table 1: Table with values of $d_{\mathcal{F}}$ for a plane (top), a sphere (middle), and a plane against a sphere (bottom).

tween a given object and the possible bends of another one are very similar, as it should be for isometric invariant recognition.

5. Conclusions

A theoretical and computational framework for comparing manifolds (metric spaces) given by point clouds was introduced in this paper. The theoretical component is based on the Gromov-Hausdorff distance, which has been extended and embedded in a probabilistic framework to deal with point clouds and computable distances. Examples supporting this theory were provided. We are currently working on improving the computational efficiency of the algorithm and comparing high dimensional point clouds with data from image sciences and neuroscience.

Acknowledgments: Work supported in part by ONR, NSF, NIH, and CSIC-Uruguay.

References

- [1] E. Arias-Castro, D. Donoho, and X. Huo, "Near-optimal detection of geometric objects by fast multiscale methods," *Stanford Statistics Department TR*, 2003.
- [2] M. Belkin and P. Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *University of Chicago CS TR-2002-01*, 2002.
- [3] J-D. Boissonnat and F. Cazals, in A. Chalmers and T-M. Rhyne, Editors, *EUROGRAPHICS '01*, Manchester, 2001.

- [4] M. Botsch, A. Wiratanaya, and L. Kobbelt, *EUROGRAPHICS Workshop on Rendering*, 2002.
- [5] D. Burago, Y. Burago, and S. Ivanov, *A Course in Metric Geometry*, AMS Graduate Studies in Math., Vol. 33.
- [6] D. Burago and B. Kleiner, *Geom. Funct. Anal.* **8**, pp. 273–282, 1998.
- [7] D. Burago and B. Kleiner, *Geom. Funct. Anal.* **12**, pp. 80–92, 2002.
- [8] G. Carlsson, A. Collins, L. Guibas, and A. Zomorodian, “Persistent homology and shape description I, preprint, *Stanford Math Department*, September 2003.
- [9] G. Charpiat, O. Faugeras, and R. Keriven. *Proceedings of the International Conference on Image Processing*, IEEE Signal Processing Society, September 2003.
- [10] T. K. Dey, J. Giesen, and J. Hudson, *Proc. 13th Canadian Conf. on Comp. Geom.*, pp. 85-88, 2001.
- [11] D. L. Donoho and C. Grimes, “Hessian eigenmaps: New locally-linear embedding techniques for high-dimensional data,” *Technical Report Department of Statistics, Stanford University*, 2003.
- [12] N. Dyn, M. S. Floater, and A. Iske, *Journal of Comp. and App. Math.* **145(2)**, pp. 505-517, 2002.
- [13] A. Elad (Elbaz) and R. Kimmel, *In Proc. of CVPR’01*, Hawaii, Dec. 2001.
- [14] W. Feller, *An Introduction to Probability Theory and its Applications*, John Wiley & Sons, Inc., New York-London-Sydney, 1971.
- [15] J. Giesen and U. Wagner, *ACM Symposium on Computational Geometry*, San Diego, June 2003.
- [16] M.T. Goodrich, J.S.B. Mitchell and M.W. Orletsky, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21:4**, pp. 371 - 379, 1999.
- [17] M. Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Progress in Mathematics **152**, Birkhäuser Boston, Inc., Boston, MA, 1999.
- [18] M. Gromov, *Geometric group theory, Vol. 2: Asymptotic invariants of infinite groups*. Edited by A. Niblo and Martin A. Roller. London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, 1993.
- [19] M. Gross et al, *Point Based Computer Graphics*, EUROGRAPHICS Lecture Notes, 2002.
- [20] K. Grove, “Metric differential geometry,” *Differential geometry (Lyngby, 1985)*, 171–227, Lecture Notes in Math., 1263, Springer, Berlin, 1987.
- [21] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **15:9**, 1993.
- [22] P. Indyk and S. Venkatasubramanian, *Comp. Geom.: Theory and Applications* **24**, pp. 115-128, 2003.
- [23] P. W. Jones, *Invent. Math.* **102**, pp. 1-15, 1990.
- [24] D. W. Kahn, *Topology. An Introduction to the Point-Set and Algebraic Areas*, Williams & Wilkins Co., Baltimore, Md., 1975.
- [25] N. J. Kalton and M. I. Ostrovskii, *Forum Math.* **11:1**, pp. 17-48, 1999.
- [26] R. Kimmel and J. A. Sethian, *Proc. National Academy of Sciences* **95:15**, pp. 8431-8435, 1998.
- [27] G. Leibon and D. Letscher, *Computational Geometry 2000*, Hong Kong, 2000.
- [28] L. Linsen, “Point cloud representation,” *CS Technical Report*, University of Karlsruhe, 2001.
- [29] L. Linsen and H. Prautzsch, *EUROGRAPHICS ’01*, 2001.
- [30] C.T. McMullen *Geom. Funct. Anal.* **8**, pp. 304-314, 1998.
- [31] F. Mémoli and G. Sapiro, *IEEE Workshop on Variational and Level-Sets Methods in Computer Vision*, Nice, France, October 2003.
- [32] N. J. Mitra and A. Nguyen, *ACM Symposium on Computational Geometry*, San Diego, June 2003.
- [33] C. Moenning and N. A. Dodgson, “Fast marching farthest point sampling for implicit surfaces and point clouds,” *Computer Laboratory Technical Report* **565**.
- [34] V. V. Nekrashevych, *Dopov. Nats. Akad. Nauk Ukr. Mat. Prirodozn. Tekh. Nauki*, pp. 18-21, 1997.
- [35] M. Pauly, M. Gross, and L. Kobbelt, *IEEE Vis. 2002*.
- [36] P. Petersen, *Riemannian Geometry*, Springer-Verlag, New York, 1998.
- [37] P. Petersen, *Differential geometry: Riemannian geometry* (Los Angeles, CA, 1990), pp 489-504, *Proc. Sympos. Pure Math.*, 54, Part 3, Amer. Math. Soc., Providence, RI, 1993.
- [38] S. Rusinkiewicz and M. Levoy, *Computer Graphics (SIGGRAPH 2000 Proceedings)*, 2000.
- [39] J. B. Tenenbaum, V. de Silva, and J. C. Langford, *Science*, pp. 2319-2323, December 2000.
- [40] D. Toledo, *Bull. Amer. Math. Soc.* **33**, pp. 395-398, 1996.
- [41] H. Von Schelling, *Amer. Math. Monthly* **61**, pp. 306–311, 1954.
- [42] E. W. Weisstein, mathworld.wolfram.com/HyperspherePointPicking.html
- [43] M. Zwicker, M. Pauly, O. Knoll, and M. Gross, *SIGGRAPH*, 2002.