

Distance Functions and Geodesics on Point Clouds

Facundo Mémoli

Electrical and Computer Engineering Dept.

University of Minnesota

Minneapolis, MN 55455

and

Instituto de Ingeniería Eléctrica

Universidad de la República

Montevideo, Uruguay

memoli@ece.umn.edu

Guillermo Sapiro*

Electrical and Computer Engineering Dept.

University of Minnesota

Minneapolis, MN 55455

guille@ece.umn.edu

May 28, 2003

Abstract

An new paradigm for computing intrinsic distance functions and geodesics on sub-manifolds of \mathbb{R}^d given by point clouds is introduced in this paper. The basic idea is that, as shown here, intrinsic distance functions and geodesics on general co-dimension sub-manifolds of \mathbb{R}^d can be accurately approximated by extrinsic Euclidean ones computed inside a thin offset band surrounding the manifold. This permits the use of computationally optimal algorithms for computing distance functions in Cartesian grids. We use these algorithms, modified to deal with spaces with boundaries, and obtain also for the case of intrinsic distance functions on sub-manifolds of \mathbb{R}^d , a computationally optimal approach. For point clouds, the offset band is constructed without the need to explicitly find the underlying manifold, thereby computing intrinsic distance functions and geodesics on point clouds while skipping the manifold reconstruction step. The case of point clouds representing noisy samples of a sub-manifold of Euclidean space is studied as well. All the underlying theoretical results are presented along with experimental examples for diverse applications and comparisons to graph-based distance algorithms.

*Corresponding author

1 Introduction

One of the most popular sources of point clouds are 3D shape acquisition devices, such as laser range scanners, with applications in geoscience, art (e.g., archival), medicine (e.g., prosthodontics), manufacturing (from cars to clothes), and security (e.g., recognition), among other 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 [10, 18, 20, 24, 27, 37, 38, 46, 47] for a few of the recent works with this type of data. Note that point clouds can also be used as primitives for visualization, e.g., [11, 27, 48], as well as for editing [58].

Another important field where point clouds are found is in the representation of high-dimensional manifolds by examples (see for example [30, 36, 53]). 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 dimensions in this case, it is impossible to perform manifold reconstruction, and the work needs to be performed directly on the raw data, meaning the point cloud.

Note that in general a point cloud representation is dimension free, in contrast with other popular representations such as triangular meshes. Some operations, such as the union of point clouds acquired from multiple views, are much easier than when performed on the triangular meshes obtained from them. This paper addresses one of the most fundamental operations in the study and processing of sub-manifolds of Euclidean space, the computation of intrinsic distance functions and geodesics. We show that this can be done working directly with the point cloud, without the need for reconstructing the underlying manifold. We present the corresponding theoretical results, experimental examples, and basic comparisons to mesh-based distance algorithms. The results are valid for general dimensions and co-dimensions, and for manifolds with or without boundary. These results include the analysis of noisy point clouds obtained from sampling the manifold.

A number of key building blocks form part of the framework here introduced. The first one is based on the fact that distance functions intrinsic to a given sub-manifold of \mathbb{R}^d can be accurately approximated by Euclidean distance functions computed in a thin offset band that surrounds this manifold. This concept was first introduced in [41], where convergence results were given for co-dimension one hyper-surfaces without boundary. This result is reviewed in §2. In this paper, we first extend these results to general co-dimension and to deal with manifolds with or without boundary, §3. We also show that the approximation is true not only for the intrinsic distance function but also for the intrinsic geodesic. This is not a straightforward corollary, since geodesics are based on the gradient of the distance function,¹ which contains singularities at the cut locus [49, 57]. The approximation of intrinsic distance functions (and geodesics) by extrinsic Euclidean ones permits to compute them using computationally optimal algorithms in Cartesian grids (as long as the discretization operation is permitted, memory wise, see §7.1 and §8). These algorithms are based on the fact that the distance function satisfies a Hamilton-Jacobi partial differential equation (see §2), for which consistent and fast algorithms have been developed in Cartesian grids [29, 51, 52, 55]² (see [32] for extensions to triangular meshes and [54] for other Hamilton-Jacobi equations). That

¹Geodesics are the integral curves corresponding to the gradient directions of the intrinsic distance function, and are obtained back-propagating in this gradient direction from the target point to the source one.

²Tsitsiklis first described an optimal-control type of approach to solve the Hamilton-Jacobi equation, while independently Sethian and Helmsen both developed techniques based on upwind numerical schemes.

is, due to these results, we can use computationally optimal algorithms in Cartesian grids (with boundaries) also to compute distance functions, and from them geodesics, intrinsic to a given manifold, and in a computationally optimal fashion. Note that in contrast with the popular Dijkstra algorithm, these numerical techniques are consistent, they converge to the true distance when the grid is refined. Dijkstra’s algorithm suffers from digitization bias due to metrication error when implemented on a grid (if no new graph edges are added to account for the new diagonals in each successive level of refinement of the grid), see [43, 44].

Once these basic results are available, we can then proceed and work with point clouds. The basic idea here is to construct the offset band directly from the point cloud and without the intermediate step of manifold reconstruction. This is addressed in §4 and §5 for noise-free points and manifold samples, and in §6 for points considered to be noisy samples of the manifold. In this case, we explicitly compute the probability that the constructed offset band contains the underlying manifold. As we expect, this probability is a function of the number of point samples, the noise level, the size of the offset, and the basic geometric characteristics of the underlying manifold. In the experimental section, §7, we present a number of important applications. These applications are given to show the importance of this novel computational framework, and are by no means exhaustive. Concluding remarks are provided in §8 where we also report the directions our research is taking.

To conclude this introduction, we should note that to the best of our knowledge, the only additional work explicitly addressing the computation of distance functions and geodesics for point clouds is the one reported in [8, 53].³ The comparison of performances in presence of noise between our framework and the one proposed in [8, 53] is deferred to Appendix A.⁴

2 Preliminary Results and Notation

In this section we briefly review the main results in [41], where the idea of approximating intrinsic distances and geodesics by extrinsic ones was first introduced.

2.1 Notation

We first introduce some basic notation that will be used throughout the article. For a compact and connected set $\Omega \in \mathbb{R}^d$, $d_\Omega(\cdot, \cdot)$ denotes the intrinsic distance between any two points of Ω , measured by paths constrained to be in Ω . We will also assume the convention that if $A \subset \mathbb{R}^d$ is compact, and x, y are not both in A then $d_A(x, y) = D$, for some constant $D \gg \max_{x, y \in A} d_A(x, y)$. Given a k -dimensional sub-manifold \mathcal{M} of \mathbb{R}^d , $\Omega_{\mathcal{M}}^h$ denotes the set $\{x \in \mathbb{R}^d : d(\mathcal{M}, x) \leq h\}$ (here the distance $d(\cdot, \cdot)$ is the Euclidean one). This is basically an h -offset of \mathcal{M} . To state that the sequence of functions $\{f_n(\cdot)\}_{n \in \mathbb{Z}^+}$ uniformly converges to $f(\cdot)$ as $n \uparrow \infty$, we frequently write $f_n \xrightarrow{n} f$. For a given event \mathcal{E} , $\mathbb{P}(\mathcal{E})$ stands for its probability of occurring. For a random variable (R.V. from now on) X , its mean value is denoted by $\mathbb{E}(X)$. We denote by $X \sim \mathbf{U}[A]$ that the R.V. X is *uniformly distributed* in the set A . For a function $f : \Omega \rightarrow \mathbb{R}$, and a subset A of Ω , $f|_A : A \rightarrow \mathbb{R}$ denotes

³In addition to studying the computation of distance functions on point clouds, [8, 53] address the important combination of this with multidimensional scaling for manifold analysis. Prior work on using geodesics and multidimensional scaling can be found in [50].

⁴While concluding this paper, we learned of a recent extension to Isomap reported in [25]. This paper is also mesh based, and follows the geodesics approach in Isomap with a novel neighborhood/connectivity approach and a number of interesting theoretical results and novel dimensionality estimation contributions. Further analysis of Isomap, as a dimensionality reduction technique, can be found in [19].

the restriction of f to A . For a smooth function $f : \Omega \rightarrow \mathbb{R}$, Df , D^2f , D^3f stand for the first, second (Hessian matrix) and third differential of f , respectively. Given a point x on the complete manifold \mathcal{S} , $B_{\mathcal{S}}(x, r)$ will denote the (intrinsic) open ball of radius $r > 0$ centered at x , and $B(y, r)$ will denote the *Euclidean* ball centered at y of radius r .

2.2 Prelude

In [41], we presented a new approach for the computation of weighted intrinsic distance functions on hyper-surfaces. We proved convergence theorems and addressed the fast, computationally optimal, computation of such approximations, see comments after Theorem 1 below. The key starting idea is that distance functions satisfy the (intrinsic) Eikonal equation, a particular case of the general class of Hamilton-Jacobi partial differential equations. Given $p \in \mathcal{S}$ (an hyper-surface in \mathbb{R}^d), we want to compute $d_{\mathcal{S}}(p, \cdot) : \mathcal{S} \rightarrow \mathbb{R}^+ \cup \{0\}$, the intrinsic distance function from every point on \mathcal{S} to p . It is well known that the distance function $d_{\mathcal{S}}(p, \cdot)$ satisfies, in the viscosity sense (see [39]), the equation

$$\begin{cases} \|\nabla_{\mathcal{S}} d_{\mathcal{S}}(p, x)\| = 1 \quad \forall x \in \mathcal{S} \\ d_{\mathcal{S}}(p, p) = 0 \end{cases}$$

where $\nabla_{\mathcal{S}}$ is the intrinsic differentiation (gradient). Instead of solving this intrinsic Eikonal equation on \mathcal{S} , we solve the corresponding extrinsic one in the offset band $\Omega_{\mathcal{S}}^h$:

$$\begin{cases} \|\nabla_x d_{\Omega_{\mathcal{S}}^h}(p, x)\| = 1 \quad \forall x \in \Omega_{\mathcal{S}}^h \\ d_{\Omega_{\mathcal{S}}^h}(p, p) = 0 \end{cases}$$

where $d_{\Omega_{\mathcal{S}}^h}(p, \cdot)$ is the Euclidean distance and therefore now the differentiation is the usual one.

Theorem 1 ([41]) *Let p and q be any two points on the smooth (orientable, without boundary) hyper-surface \mathcal{S} , then $|d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{S}}^h}(p, q)| \leq C_{\mathcal{S}}\sqrt{h}$, for small enough h ,⁵ where $C_{\mathcal{S}}$ is a constant depending on the geometry of \mathcal{S} .*

This simplification of the intrinsic problem into an extrinsic one permits the use of the computationally optimal algorithms mentioned in the introduction. This makes computing intrinsic distances, and from them geodesics, as simple and computationally efficient as computing them in Euclidean spaces. Moreover, as detailed in [41], the approximation of the intrinsic distance $d_{\mathcal{S}}$ by the extrinsic Euclidean one $d_{\Omega_{\mathcal{S}}^h}$ is never less accurate than the numerical error of these algorithms.

In [41], the result above was limited to hyper-surfaces of \mathbb{R}^d (co-dimension one sub-manifolds of \mathbb{R}^d) without boundary, and the theory was applied to implicit surfaces, where computing the offset band is straightforward. It is the purpose of the present work to extend the aforementioned Theorem to deal with: (1) sub-manifolds of \mathbb{R}^d of any codimension and possibly with boundary,⁶ (2) sub-manifolds of \mathbb{R}^d represented as point clouds, (3) random sampling of sub-manifolds of \mathbb{R}^d in presence of noise, and (4) convergence of geodesic curves in addition to distance functions. We should note that Theorem 1 holds even when the metric is not the one inherited from \mathbb{R}^d , obtaining weighted distance functions, see [41]. Although we will not present these new results in such generality, this is a simple extension that will be reported elsewhere.

⁵“Small enough h ” means that $h < 1/\max_i \kappa_i(\mathcal{S})$, where $\kappa_i(\mathcal{S})$ is the i -th principal curvature of \mathcal{S} . This guarantees having smoothness in $\partial\Omega_{\mathcal{S}}^h$, see [41].

⁶We will later impose some conditions on the boundary in order to get rate of convergence estimates. However, the uniform convergence in itself doesn't require other hypotheses beyond smoothness.

3 Sub-Manifolds of \mathbb{R}^d with Boundary

We first extend Theorem 1 to more general manifolds (with boundary and higher co-dimension) and we deal not only with distance functions but also with geodesics. The first extension is important for the learning of high-dimensional manifolds from samples and for scanned open volumes. The extension to geodesics is important for path planning on surfaces and for finding special curves such as crests and valleys, see [7, 41].

First we need to recall some results that will be key ingredients in our proofs below. All our results rest upon a certain degree of smoothness of geodesics in manifolds with boundary. We use “shortest path” and “minimizing geodesic” interchangeably.

Theorem 2 ([1]) *Let \mathcal{M} be a C^3 Riemannian manifold with C^1 boundary $\partial\mathcal{M}$. Then, any shortest path of $\partial\mathcal{M}$ is C^1 .*

We will eventually need more regularity on the geodesics than simply C^1 . This is achieved by requiring more regularity of the boundary.

Theorem 3 ([40]) *Let $\mathcal{U} : \mathbb{R}^d \rightarrow \mathbb{R}$ be a C^3 function such that for some $h \in \mathbb{R}$*

- *the interior of $\{x \in \mathbb{R}^d \mid \mathcal{U}(x) = h\}$ is non-empty and there we have $D\mathcal{U}(x) \neq 0$.*
- *the “obstacle” $\{x \in \mathbb{R}^d \mid \mathcal{U}(x) \geq h\}$ is compact.*

Let p and q be any two points in the same connected component of $\{x \in \mathbb{R}^d \mid \mathcal{U}(x) \leq h\}$, then the shortest (constrained) path joining both points is C^1 and has absolutely continuous (therefore Lipschitz) first derivative.

We now present the usual definition of length:

Definition 1 *Let $\alpha : [a, b] \rightarrow \mathbb{R}^d$ be a curve, then we define its length $\mathbf{L}(\alpha)$ as*

$$\mathbf{L}(\alpha) \triangleq \sup_{a=t_0 < \dots < t_N=b} \sum_{k=0}^{N-1} \|\alpha(t_{k+1}) - \alpha(t_k)\|$$

Remark 1 *Note that if α is Lipschitz with constant \mathcal{L}_α , then $\mathbf{L}(\alpha) = \int_a^b \|\dot{\alpha}(t)\| dt$ and $\mathbf{L}(\alpha) \leq \mathcal{L}_\alpha(b-a)$.*

Proposition 1 *Let \mathcal{S} be a smooth compact sub-manifold of \mathbb{R}^d with boundary $\partial\mathcal{S}$. Let x, y be any two points in \mathcal{S} . Then, $d_{\Omega_{\mathcal{S}}^h}(x, y)$ converges pointwise as $h \downarrow 0$.*

Proof:

Since $\Omega_{\mathcal{S}}^h \subseteq \Omega_{\mathcal{S}}^{h'}$ if $h' \geq h$, we have that $d_{\Omega_{\mathcal{S}}^h}(x, y) \geq d_{\Omega_{\mathcal{S}}^{h'}}(x, y)$. Also, for any $h > 0$, $d_{\Omega_{\mathcal{S}}^h}(x, y) \leq d_{\mathcal{S}}(x, y) \leq \text{diam}(\mathcal{S}) < +\infty$. Hence, the sequence $\{d_{\Omega_{\mathcal{S}}^h}(x, y)\}_{h>0}$ (for fixed x and y over \mathcal{S}) is bounded and non-decreasing, therefore it converges to the supremum of its range. \square

Theorem 4 *Let \mathcal{S} be a compact C^2 sub-manifold of \mathbb{R}^d with (possibly empty) smooth boundary $\partial\mathcal{S}$. Let x, y be any two points in \mathcal{S} . Then we have*

1. *Uniform convergence of distances:*

$$d_{\Omega_{\mathbb{S}}^h} |_{\mathcal{S} \times \mathcal{S}}(\cdot, \cdot) \xrightarrow{h \downarrow 0} d_{\mathcal{S}}(\cdot, \cdot)$$

2. *Convergence of geodesics:* Let x and y be joined by a unique minimizing geodesic $\gamma_{\mathcal{S}} : [0, 1] \rightarrow \mathcal{S}$ over \mathcal{S} , and let $\gamma_h : [0, 1] \rightarrow \Omega_{\mathbb{S}}^h$ be a $\Omega_{\mathbb{S}}^h$ -minimizing geodesic, then

$$\gamma_h \xrightarrow{h \downarrow 0} \gamma_{\mathcal{S}}$$

Proof:

Given our hypothesis on \mathcal{S} , and according to [23], there exists $H > 0$ such that $\partial\Omega_{\mathbb{S}}^h$ is $C^{1,1}$ for all $0 < h \leq H$. Then Theorem 2 guarantees that for $0 < h \leq H$, $\gamma_h : [0, 1] \rightarrow \Omega_{\mathbb{S}}^h$, the $\Omega_{\mathbb{S}}^h$ length minimizing geodesic joining x and y is of class C^1 . Since $d_{\Omega_{\mathbb{S}}^h}(x, y) \leq d_{\mathcal{S}}(x, y) \leq \text{diam}(\mathcal{S}) < +\infty$ for any $h \in (0, H]$, we see that we can admit our $\Omega_{\mathbb{S}}^h$ -geodesics to have Lipschitz constant $\mathcal{L} \leq \text{diam}(\mathcal{S})$. Obviously, the set $\Omega_{\mathbb{S}}^H$ is bounded, and then the family $\{\gamma_h\}_{0 < h \leq H}$ is bounded and equicontinuous. Hence, by Ascoli-Arzelá's Theorem, there exist a subsequence $\{\gamma_{h_k}\}_{k \in \mathbb{N}}$ and a curve $\gamma_0 \in C^0([0, 1], \mathcal{S})$ such that $\max_{t \in [0, 1]} \|\gamma_{h_k}(t) - \gamma_0(t)\| \xrightarrow{h_k \downarrow 0} 0$.

Moreover, by writing $|\gamma_0(t) - \gamma_0(t')| \leq |\gamma_{h_k}(t) - \gamma_0(t)| + |\gamma_{h_k}(t') - \gamma_0(t')| + \mathcal{L}|t - t'|$ and using the (pointwise) convergence of γ_{h_k} towards γ_0 , we find that \mathcal{L} is also a Lipschitz constant for γ_0 . Then we have $\gamma_0 \in C^{0,1}([0, 1], \mathcal{S})$.

Now, since γ_0 lies on \mathcal{S} but may not be a shortest path, we have that its (finite) length is greater than or equal to $d_{\mathcal{S}}(x, y)$. We also have the trivial inequality $d_{\mathcal{S}}(x, y) \geq d_{\Omega_{\mathbb{S}}^h}(x, y)$. Putting all together we obtain

$$\mathbf{L}(\gamma_h) = d_{\Omega_{\mathbb{S}}^h}(x, y) \leq d_{\mathcal{S}}(x, y) \leq \mathbf{L}(\gamma_0)$$

Therefore

$$\limsup_{h \downarrow 0} \mathbf{L}(\gamma_h) = \limsup_{h \downarrow 0} d_{\Omega_{\mathbb{S}}^h}(x, y) \leq d_{\mathcal{S}}(x, y) \leq \mathbf{L}(\gamma_0)$$

Note that $\mathbf{L}(\gamma_0) = \mathbf{L}(\lim_{h_k \downarrow 0} \gamma_{h_k}) \leq \liminf_{h_k \downarrow 0} \mathbf{L}(\gamma_{h_k})$. This is the semicontinuity of length, an immediate consequence of its definition, see [33].

Since $\liminf_{h_k \downarrow 0}(\cdot) \leq \limsup_{h_k \downarrow 0}(\cdot) \leq \limsup_{h \downarrow 0}(\cdot)$, we see that $\limsup_{h \downarrow 0} d_{\Omega_{\mathbb{S}}^h}(x, y) = \limsup_{h \downarrow 0} \mathbf{L}(\gamma_h)$ equals $d_{\mathcal{S}}(x, y)$, for all x and y in \mathcal{S} . From Proposition 1, we find that in fact $\lim_{h \downarrow 0} d_{\Omega_{\mathbb{S}}^h}(x, y)$ exists and equals $d_{\mathcal{S}}(x, y)$.

Then, we have that the function $d_{\Omega_{\mathbb{S}}^h} |_{\mathcal{S} \times \mathcal{S}}(\cdot, \cdot)$ satisfies:

- $d_{\Omega_{\mathbb{S}}^h} |_{\mathcal{S} \times \mathcal{S}} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R} \cup \{0\}$ is continuous for each $H > h > 0$.
- for each $(x, y) \in \mathcal{S} \times \mathcal{S}$, $\{d_{\Omega_{\mathbb{S}}^h} |_{\mathcal{S} \times \mathcal{S}}(x, y)\}_h$ is non-decreasing.
- $d_{\Omega_{\mathbb{S}}^h} |_{\mathcal{S} \times \mathcal{S}}(\cdot, \cdot)$ converges pointwise towards $d_{\mathcal{S}}(\cdot, \cdot)$, which is continuous.

Then by Dini's Uniform Convergence Theorem (see [6]) we can conclude that the convergence is uniform.

We can also see that γ_0 must be a minimizing geodesic of \mathcal{S} since from the above chain of equalities $\mathbf{L}(\gamma_0) = d_{\mathcal{S}}(x, y)$. Then, if there were only one such curve joining x with y , we would have uniform convergence (along any subsequence!) of γ_h towards γ_0 . \square

Remark 2 *In the previous Theorem, the convergence (of distances) is uniform but we will have forfeited rate of convergence estimates unless we impose additional conditions on $\partial\mathcal{S}$, as we do in Corollary 3. Note that the new setting is wider than the one considered in Theorem 1 since the codimension of the underlying manifold is not necessarily 1. This is very important for applications such as dimensionality reduction, where the dimension of the underlying manifold is unknown beforehand.*

Corollary 1 *Let \mathcal{S} and $\partial\mathcal{S}$ satisfy the hypotheses of Theorem 4 . Let $\{\Sigma_i\}_{i \in \mathbb{N}}$ be a family compact of sets in \mathbb{R}^d such that $\mathcal{S} \subseteq \Sigma_i \forall i \in \mathbb{N}$ and $d_{\mathcal{H}}(\Sigma_i, \mathcal{S}) \xrightarrow{i \uparrow +\infty} 0$. Then,*

$$d_{\Sigma_i}(\cdot, \cdot)|_{\mathcal{S} \times \mathcal{S}} \xrightarrow{i \uparrow +\infty} d_{\mathcal{S}}(\cdot, \cdot)$$

where $d_{\mathcal{H}}$ stands for the Hausdorff distance between sets.

We now present a uniform rate of convergence result for the distance in the band in the case $\partial\mathcal{S} = \emptyset$, and from this we deduce Corollary 3 below, which deals with the case $\partial\mathcal{S} \neq \emptyset$. This result generalizes the one presented in [41] because it allows for any codimension.

Theorem 5 *Under the same hypotheses of the Theorem above, with $\partial\mathcal{S} = \emptyset$, we have that for small enough $h > 0$:*

$$\max_{(x,y) \in \mathcal{S} \times \mathcal{S}} \left| d_{\Omega_{\mathcal{S}}^h}|_{\mathcal{S} \times \mathcal{S}}(x, y) - d_{\mathcal{S}}(x, y) \right| \leq C_{\mathcal{S}} \sqrt{h} \quad (1)$$

where the constant $C_{\mathcal{S}}$ does not depend on h . Also, we have the “relative” rate of convergence bound:

$$1 \leq \sup_{\substack{x, y \in \mathcal{S} \\ x \neq y}} \frac{d_{\mathcal{S}}(x, y)}{d_{\Omega_{\mathcal{S}}^h}(x, y)} \leq 1 + C_{\mathcal{S}} \sqrt{h} \quad (2)$$

Proof:

This is a remake of our proof of the main theorem in [41], therefore we skip some technical details which can be found there. All along the proof we will sometimes write d_h instead of $d_{\Omega_{\mathcal{S}}^h}$ for the sake of notational simplicity. We will denote by $k (\leq n - 1)$ the dimension of \mathcal{S} .

Let γ_0 be the arc length parametrized \mathcal{S} -shortest path Joining the points $x, y \in \mathcal{S}$; clearly, we have $\text{trace}(\gamma_0) \subset \mathcal{S}$. Let γ_h be the $\Omega_{\mathcal{S}}^h$ arc length parametrized shortest path joining x and y , which, as we know from Theorem 4, uniformly converges toward γ_0 . For a number H as in the proof of Theorem 4, we have $\gamma_h \in C^{1,1}([0, d_h], \mathcal{S})$, and also $\eta : \Omega_{\mathcal{S}}^H \rightarrow \mathbb{R}$ defined by $\eta(x) \triangleq \frac{1}{2}d^2(x, \mathcal{S})$ is smooth, see Appendix B. We define the projection operator $\Pi_{\mathcal{S}} : \Omega_{\mathcal{S}}^H \rightarrow \mathcal{S}$ by $\Pi_{\mathcal{S}}(x) = x - D\eta(x)$. We refer the reader to Appendix B for properties of $\Pi_{\mathcal{S}}$ and η which we use below.

Now, $d_{\Omega_{\mathcal{S}}^h}(x, y) = \mathbf{L}(\gamma_h) \leq d_{\mathcal{S}}(x, y) \leq \mathbf{L}(\Pi_{\mathcal{S}}(\gamma_h))$, then

$$\begin{aligned} d_{\mathcal{S}}(x, y) - d_{\Omega_{\mathcal{S}}^h}(x, y) &\leq |\mathbf{L}(\Pi_{\mathcal{S}}(\gamma_h)) - \mathbf{L}(\gamma_h)| \\ &\leq \int_0^{d_h} \left\| \overline{\Pi_{\mathcal{S}}(\gamma_h(t)) - \gamma_h(t)} \right\| dt \\ &= \int_0^{d_h} \left\| \overline{D\eta(\gamma_h(t))} \right\| dt \\ &\leq \sqrt{d_h \int_0^{d_h} \dot{V}(t) \cdot \dot{V}(t) dt} \quad (\text{by Cauchy-Schwartz Ineq.}) \\ &\leq \sqrt{d_h \int_0^{d_h} V(t) \cdot \dot{V}(t) dt} \quad (\text{Integrating by parts, see below.}) \end{aligned}$$

where $V(t) \triangleq D\eta(\gamma_h(t))$ and $V(0) = V(1) = 0$, see Appendix B.

Also $\dot{V}(t) = D^2\eta(\gamma_h(t))\dot{\gamma}(t)$ and since γ_h is Lipschitz and η is smooth, $\dot{V}(t)$ exists almost everywhere and $\ddot{V}(t) = D^3\eta(\gamma_h(t))[\dot{\gamma}_h(t), \dot{\gamma}_h(t)] + D^2\eta(\gamma_h(t))\ddot{\gamma}(t)$ at points of existence. Then since $D^3\eta D\eta = D^2\eta(I - D^2\eta)$ and $D^2\eta D\eta = D\eta$ (see Appendix B),

$$\begin{aligned} V \cdot \ddot{V} &= D^3\eta(\gamma_h)[D\eta(\gamma_h), \dot{\gamma}_h, \dot{\gamma}_h] + D^2\eta[\ddot{\gamma}_h, D\eta(\gamma_h)] \\ &= (D^2\eta(\gamma_h)(I - D^2\eta(\gamma_h))) [\dot{\gamma}_h, \dot{\gamma}_h] + \ddot{\gamma}_h \cdot D\eta(\gamma_h) \end{aligned}$$

The matrix $\Lambda(t) \triangleq D^2\eta(\gamma_h(t))(I - D^2\eta(\gamma_h(t)))$ filters out normal components, and has eigenvalues associated with the tangential bundle given by (let $d(t) = d(\gamma_h(t), \mathcal{S})$),

$$\lambda_i(t) = \frac{d(t)\lambda_i(0)}{(1 + d(t)\lambda_i(0))^2}, \text{ for } 1 \leq i \leq k$$

Note that $\max_{1 \leq i \leq k} |\lambda_i(t)|$ can be bounded by $d(t)$ times a certain finite constant K' independent of h .

On the other hand, we can bound a.e. $|\ddot{\gamma}_h(t)|$ by a finite constant, say K , which takes into account the maximal curvature of all the boundaries $\partial\Omega_{\mathcal{S}}^h$, $0 < h < H$, but does not depend on h .

Putting all this together, we find (recall that $\|D\eta(x)\| = \sqrt{2\eta(x)} = d(x, \mathcal{S})$, see Appendix B):

$$\begin{aligned} \left(d_{\mathcal{S}}(x, y) - d_{\Omega_{\mathcal{S}}^h}(x, y)\right)^2 &\leq d_h \int_0^{d_h} \Lambda(t)[\dot{\gamma}_h, \dot{\gamma}_h] dt \\ &\quad + d_h \int_0^{d_h} \|\ddot{\gamma}_h\| \|D\eta(\gamma_h)\| dt \\ &\leq K' \max_{t \in [0, d_h]} d(t) d_h^2 + K \max_{t \in [0, d_h]} d(t) d_h^2 \end{aligned}$$

Now, remembering that d_h stands for $d_{\Omega_{\mathcal{S}}^h}(x, y)$, that $\text{trace}(\gamma_h) \subset \Omega_{\mathcal{S}}^h$, and defining $C = K + K'$, we arrive, with just simple additional work at the relations (1) or (2). \square

Remark 3 *Note that, as the simple case of a circle in the plane shows, the rate of convergence is at most $C \cdot h$.*

We immediately obtain the following Corollary which will be useful ahead.

Corollary 2 *Let $p \in \mathcal{S}$, and $r \leq H$, then $B(p, r) \cap \mathcal{S} \subseteq B_{\mathcal{S}}(p, r(1 + C_{\mathcal{S}}\sqrt{r}))$.*

Proof:

Let $q \in B(p, r) \cap \mathcal{S}$, then by (2), $d_{\mathcal{S}}(p, q) \leq d_{\Omega_{\mathcal{S}}^r}(p, q)(1 + C_{\mathcal{S}}\sqrt{r})$. But $q \in B(p, r) \subset \Omega_{\mathcal{S}}^r$, then $d_{\Omega_{\mathcal{S}}^h}(p, q) = \|p - q\| \leq r$, what completes the proof. \square

Definition 2 ([21]) *We say that the compact manifold \mathcal{S} with boundary $\partial\mathcal{S}$ is strongly convex if for every pair of points x and y in \mathcal{S} , there exists a unique minimizing geodesic joining them whose interior is contained in the interior of \mathcal{S} .*

Using basically the same procedure as in Theorem 5 with the convexity hypotheses above we can prove the following Corollary whose (sketch of) proof is presented in Appendix C.

Corollary 3 ($\partial\mathcal{S} \neq \emptyset$) *Under the hypotheses of Theorem 2, and assuming \mathcal{S} to be strongly convex, we have for small enough $h > 0$ the same conclusions of Theorem 5 (rate of convergence).*

To conclude, in this section we extended the results in [41] to geodesics and distance functions in general codimension manifolds with or without (smooth) boundary, thereby covering all possible manifolds in common shape, graphics, visualization, and learning applications.⁷ We are now ready to extend this to manifolds represented as point clouds.

⁷Although in this paper we only consider manifolds with constant co-dimension, many of the results are extendible to variable co-dimensions, and this will be reported elsewhere.

4 Distance Functions and Geodesics on Point Clouds

We are now interested in making computations on manifolds represented as point clouds, i.e. *sampled manifolds*. In the case of this paper we will restrict ourselves to the computation of intrinsic distances.⁸ Let $\mathcal{P}_n \triangleq \{p_1, \dots, p_n\}$ be a set of n different points sampled from the *compact* submanifold \mathcal{S} and define⁹

$$\Omega_{\mathcal{P}_n}^h \triangleq \bigcup_{i=1}^n B(p_i, h)$$

Let h and \mathcal{P}_n be such that $\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h$. We then have $(\mathcal{S} \subseteq) \Omega_{\mathcal{P}_n}^h \subseteq \Omega_{\mathcal{S}}^h$. We now want to consider $d_{\Omega_{\mathcal{P}_n}^h}(p, q)$ for any pair of points $p, q \in \mathcal{S}$ and prove some kind of proximity to the real distance $d_{\mathcal{S}}(p, q)$. The argument carries over easily since:

$$d_{\Omega_{\mathcal{S}}^h}(p, q) \leq d_{\Omega_{\mathcal{P}_n}^h}(p, q) \leq d_{\mathcal{S}}(p, q)$$

hence

$$0 \leq d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(p, q) \leq d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{S}}^h}(p, q) \quad (3)$$

and the rightmost quantity can be bounded by $C_{\mathcal{S}} h^{1/2}$ (see §3) in the case that $\partial\mathcal{S}$ is either convex or void. In general, without hypotheses on $\partial\mathcal{S}$ other than some degree of smoothness, we can also work out uniform convergence since by virtue of Theorem 4 the upper bound in (3) uniformly converges to 0. The key condition is $\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h$, something that can obviously be coped with using the compactness of \mathcal{S} .¹⁰ We can then state the following

Theorem 6 (Uniform Convergence for Point Clouds)

Let \mathcal{S} be a compact smooth submanifold of \mathbb{R}^d possibly with boundary $\partial\mathcal{S}$. Then

1. **General Case:** Given $\varepsilon > 0$, there exists $h_\varepsilon > 0$, such that $\forall 0 < h \leq h_\varepsilon$ one can find finite $n(h)$ and a set of points $\mathcal{P}_{n(h)}(h) = \{p_1(h), \dots, p_{n(h)}(h)\}$ sampled from \mathcal{S} such that

$$\max_{p, q \in \mathcal{S}} \left(d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_{n(h)}(h)}^h}(p, q) \right) \leq \varepsilon$$

2. **$\partial\mathcal{S}$ is either void or convex:** For every sufficiently small $h > 0$ one can find finite $n(h)$ and a set of points $\mathcal{P}_{n(h)}(h) = \{p_1(h), \dots, p_{n(h)}(h)\}$ sampled from \mathcal{S} such that

$$\max_{p, q \in \mathcal{S}} \left(d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_{n(h)}(h)}^h}(p, q) \right) \leq C_{\mathcal{S}} \sqrt{h}$$

In practise, one must worry about both the number of points and the radii of the balls. Obviously, there is a tradeoff between both quantities. If we want to use few points, in order to cover \mathcal{S} with the balls we have to increase the value of the radius. Clearly, there exists a value H such that for values of h smaller than H we don't change the topology, see [3, 4, 5]. This implies that the number of points must be larger than a certain lower bound. This result can be generalized to

⁸Note that having the intrinsic distance allows us to compute basic intrinsic properties of the manifold, e.g., [12].

⁹The balls now used are defined with respect to the metric of \mathbb{R}^d , they are not intrinsic.

¹⁰By compactness, given $h > 0$ we can find finite $N(h)$ and points $p_1, p_2, \dots, p_{N(h)} \in \mathcal{S}$ such that $\mathcal{S} = \bigcup_{i=1}^{N(h)} B_{\mathcal{S}}(p_i, h)$. But since for $p \in \mathcal{S}$, $B_{\mathcal{S}}(p, h) \subset B(p, h) \cap \mathcal{S}$, and we also get $\mathcal{S} \subset \bigcup_{i=1}^{N(h)} B(p_i, h)$.

ellipsoids which can be locally adapted to the geometry of the point cloud [14], or from minimal spanning trees. Note that we are interested in the smallest possible offset of the point cloud that covers \mathcal{S} . Further comments on this are presented below and are also subject of current efforts to be reported elsewhere.

The practical significance of the previous Theorem is clear. Part 1 says that in general, given a desired precision for the computation of the distance, we have a maximum *non-zero* value for the radius of all the balls below which we can always find a *finite* number of points sampled from the manifold for which the “ Ω -set” formed by those points achieves the desired accuracy,¹¹ this is, we can choose the radius at our convenience *within* a certain range which depends on this level of accuracy. Part 2 says more, since it actually links ε to h_ε . It basically says that the radius of the balls must be of the order of the square of the desired error.

5 Extension to Random Sampling of Manifolds

In practise, we really do not have too much control over the way in which points are sampled by the acquisition device (e.g. scanner), or given by the learned sampled data. Therefore it is more realistic to make probabilistic a model of the situation and then try to conveniently estimate the probability of achieving a prescribed level of accuracy as a function of the number of points and the radii of the balls. It will be interesting to see how geometric quantities of \mathcal{S} enter in those bounds we will establish. However, since the bounds are based in local volume computations, and all manifolds are locally Euclidean, those curvature dependent quantities will be asymptotically negligible.

We now present a simple model for the current setting, while results for other models can be developed from the derivations below. Here we assume that the points in \mathcal{P}_n are independently and identically sampled on the sub-manifold \mathcal{S} in a uniform fashion,¹² we will write this as $p_i \sim \mathbf{U}[\mathcal{S}]$. For simplicity of exposition, we will restrict ourselves to the case when \mathcal{S} has no boundary.¹³ Also, we only deal with uniform i.i.d. sampling, other models for the sampling will be reported elsewhere.

We have to define the way in which we are going to measure accuracy. A possibility for such a measure is (for each $\varepsilon > 0$)

$$\mathbb{P} \left(\max_{p,q \in \mathcal{S}} \left(d_{\mathcal{S}}(p,q) - d_{\Omega_{\mathcal{P}_n}^h}(p,q) \right) > \varepsilon \right) \quad (4)$$

There is a potential problem with this way of testing accuracy, since we are assuming that when we use the approximate distance, we will be evaluating it on \mathcal{S} . This might seem a bit awkward since we don't exactly know all the surface but just some points on it. Moreover, a more natural and real-problem-motivated approach would be to measure the discrepancy over \mathcal{P}_n itself (see §7 ahead), over part of this set, or over another *trial* set of points \mathcal{Q}_m .

However, since for any set of points $\mathcal{Q}_m \subset \mathcal{S}$ we have that the following inclusion of events

$$\left\{ \max_{p,q \in \mathcal{Q}_m} \left(d_{\mathcal{S}}(p,q) - d_{\Omega_{\mathcal{P}_n}^h}(p,q) \right) > \varepsilon \right\} \subseteq \left\{ \max_{p,q \in \mathcal{S}} \left(d_{\mathcal{S}}(p,q) - d_{\Omega_{\mathcal{P}_n}^h}(p,q) \right) > \varepsilon \right\}$$

holds, bounding (4) suffices for dealing with any of the possibilities mentioned above. Note that we

¹¹We are considering the case when all the balls have the same radii.

¹²This means that for any subset $A \subseteq \mathcal{S}$, and any $p_i \in \mathcal{P}_n$, $\mathbb{P}(p_i \in A) = \frac{\mu(A)}{\mu(\mathcal{S})}$

¹³Even if we elaborate on the modifications needed in our arguments we should say that the same corresponding considerations presented in [8] are still valid in our case.

are somehow considering $d_{\Omega_{\mathcal{P}_n}^h}$ defined for all pairs of points in $\mathcal{S} \times \mathcal{S}$, even if it might happen that $\mathcal{S} \cap \Omega_{\mathcal{P}_n}^h \neq \mathcal{S}$. In any case we extend $d_{\Omega_{\mathcal{P}_n}^h}$ to all $\Omega_{\mathcal{S}}^h \times \Omega_{\mathcal{S}}$ by a large constant say $k \text{ diam}(\mathcal{S})$, $k \gg 1$.

Let us make some definitions so as to avoid an overload of notation:

$$\mathcal{E}_\varepsilon \triangleq \left\{ \max_{p,q \in \mathcal{S}} \left(d_{\mathcal{S}}(p,q) - d_{\Omega_{\mathcal{P}_n}^h}(p,q) \right) > \varepsilon \right\} \quad (5)$$

$$\mathcal{J}_{h,n} \triangleq \left\{ \mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h \right\} \quad (6)$$

Now, since $\mathcal{E}_\varepsilon = (\mathcal{E}_\varepsilon \cap \mathcal{J}_{h,n}) \cup (\mathcal{E}_\varepsilon \cap \mathcal{J}_{h,n}^c)$, using the union bound and then Bayes rule we have

$$\begin{aligned} \mathbb{P}(\mathcal{E}_\varepsilon) &\leq \mathbb{P}(\mathcal{E}_\varepsilon \cap \mathcal{J}_{h,n}) + \mathbb{P}(\mathcal{E}_\varepsilon \cap \mathcal{J}_{h,n}^c) \\ &= \mathbb{P}(\mathcal{E}_\varepsilon \mid \mathcal{J}_{h,n}) \mathbb{P}(\mathcal{J}_{h,n}) + \mathbb{P}(\mathcal{E}_\varepsilon \mid \mathcal{J}_{h,n}^c) \mathbb{P}(\mathcal{J}_{h,n}^c) \\ &\quad \Downarrow \\ \mathbb{P}(\mathcal{E}_\varepsilon) &\leq \mathbb{P}(\mathcal{E}_\varepsilon \mid \mathcal{J}_{h,n}) + \mathbb{P}(\mathcal{J}_{h,n}^c) \end{aligned} \quad (7)$$

It is clear now that we must find a convenient lower bound for the second term in the previous expression, the probability of covering all \mathcal{S} with the union of balls. The first term will be dealt with using the convergence theorems presented in previous sections. for this, we need a few lemmas.

Lemma 1 *Let K be an upper bound for the sectional curvatures of \mathcal{S} ($\dim(\mathcal{S}) = k$) and $x \in \mathcal{S}$ be a fixed point. Then under the hypotheses on \mathcal{P}_n described above, there exists a constant $\omega_k > 0$ and a function $\theta_{\mathcal{S}}(\cdot)$ with $\lim_{h \downarrow 0} \frac{\theta_{\mathcal{S}}(h)}{h^{k+1}} = 0$ such that for small enough $h > 0$:*

$$\mathbb{P}\left(\left\{x \notin \Omega_{\mathcal{P}_n}^h \cap \mathcal{S}\right\}\right) \leq \left(1 - \frac{\omega_k h^k + \theta_{\mathcal{S}}(h)}{\mu(\mathcal{S})}\right)^n \quad (8)$$

Moreover, one can further expand the right hand side of (8) as $\left(1 - \frac{\omega_k h^k (1 - K c_k h^2) + \phi_{\mathcal{S}}(h)}{\mu(\mathcal{S})}\right)^n$ for some c_k depending only on the dimension k of \mathcal{S} and a function $\phi_{\mathcal{S}}$ such that $\frac{\phi_{\mathcal{S}}(h)}{h^{k+2}} \rightarrow 0$ as $h \downarrow 0$.

Proof.

$$\mathbb{P}\left(\left\{x \notin \Omega_{\mathcal{P}_n}^h \cap \mathcal{S}\right\}\right) = \mathbb{P}\left(\left\{\bigcap_{i=1}^n \left\{x \notin B(p_i, h) \cap \mathcal{S}\right\}\right\}\right) \quad (9)$$

$$= \mathbb{P}\left(\left\{\bigcap_{i=1}^n \left\{p_i \notin B(x, h) \cap \mathcal{S}\right\}\right\}\right) \quad (10)$$

$$= \prod_{i=1}^n \mathbb{P}\left(\left\{p_i \notin B(x, h) \cap \mathcal{S}\right\}\right) \quad (11)$$

$$= \prod_{i=1}^n \left(1 - \mathbb{P}\left(\left\{p_i \in B(x, h) \cap \mathcal{S}\right\}\right)\right) \quad (12)$$

Since $B_{\mathcal{S}}(x, h) \subseteq B(x, h) \cap \mathcal{S}$,¹⁴ then $\mu(\mathcal{S} \cap B(x, h)) \geq \mu(B_{\mathcal{S}}(x, h))$ On the other hand note that

$$\begin{aligned} \mathbb{P}\left(\left\{p_i \in B(x, h) \cap \mathcal{S}\right\}\right) &= \frac{\mu(\mathcal{S} \cap B(x, h))}{\mu(\mathcal{S})} \\ &\geq \frac{\mu(B_{\mathcal{S}}(x, h))}{\mu(\mathcal{S})} \end{aligned}$$

¹⁴Consider $z \in B_{\mathcal{S}}(x, h)$, then $d_{\mathcal{S}}(x, z) \leq h$, but always $d(x, z) \leq d_{\mathcal{S}}(x, z)$, thus $d(x, z) \leq h$, what implies $z \in B(x, h) \cap \mathcal{S}$.

Finally, as shown in Appendix D one can lower bound $\mu(B_S(x, h))$ using information on the curvatures of \mathcal{S} , by means of the Bishop-Günther Volume Comparison Theorem. More precisely, we can write

$$\mu(B_S(x, h)) \geq \min_{\zeta \in \mathcal{S}} \mu(B_S(\zeta, h)) \geq \omega_k h^k + \theta_S(h)$$

where $\frac{\theta_S(h)}{h^q} \rightarrow 0$ when $h \rightarrow 0$ for $q \leq k+1$. Therefore, from (9) we obtain

$$\mathbb{P}(\{x \notin \Omega_{\mathcal{P}_n}^h \cap \mathcal{S}\}) \leq \left(1 - \frac{\omega_k h^k + \theta_S(h)}{\mu(\mathcal{S})}\right)^n$$

The last assertion follows from the proposition is Appendix D.3. \square

Remark 1 Note that we cannot, however, from (8), conclude that $\mathbb{P}(\mathcal{S} \not\subseteq \Omega_{\mathcal{P}_n}^h) \leq \left(1 - \frac{\omega_k h^k + \theta_S(h)}{\mu(\mathcal{S})}\right)^n$. In order to upper bound $\mathbb{P}(\mathcal{S} \not\subseteq \Omega_{\mathcal{P}_n}^h)$ we will first estimate $\mathbb{P}(B_S(x, \delta) \not\subseteq \Omega_{\mathcal{P}_n}^h)$ for any $x \in \mathcal{S}$, and small $\delta > 0$. Then we will use the compactness of \mathcal{S} by covering it with a finite δ -net consisting of $\mathcal{N}(\mathcal{S}, \delta)$ points, and conclude by using the union bound. Yet another intermediate step will therefore be to estimate the covering number $\mathcal{N}(\mathcal{S}, \delta)$.

Lemma 2 Under the hypotheses of the previous Lemma, let $\delta \in (0, h)$, then

$$\mathbb{P}(B_S(x, \delta) \not\subseteq \Omega_{\mathcal{P}_n}^h) \leq \left(1 - \frac{\omega_k (h - \delta)^k + \theta_S(h - \delta)}{\mu(\mathcal{S})}\right)^n \quad (13)$$

Proof:

We find α and β such that $\{B_S(q, \delta) \subseteq \Omega_{\mathcal{P}_n}^h\} \supseteq \{q \in \Omega_{\mathcal{P}_n}^{\alpha h + \beta \delta}\}$. Fix any $x \in B_S(q, \delta)$, then $|x - q| \leq d_S(x, q) \leq \delta$. Assume that the event $\{q \in \Omega_{\mathcal{P}_n}^{\alpha h + \beta \delta}\}$ holds. Then for some $p_r \in \mathcal{P}_n$, $q \in B(p_r, \alpha h + \beta \delta)$, that is $|q - p_r| \leq \alpha h + \beta \delta$. Now, note that

$$|x - p_r| \leq |x - q| + |q - p_r| \leq \alpha h + (\beta + 1)\delta$$

If we force the rightmost number to be h , we find that we must have $(1 + \beta)\delta = (1 - \alpha)h$, then $\alpha h + \beta \delta = h - \delta$. Then we have found $B_S(q, \delta) \subseteq B(p_r, h - \delta) \subset \Omega_{\mathcal{P}_n}^h$. Hence (using (8)), $\mathbb{P}(B_S(q, \delta) \subseteq \Omega_{\mathcal{P}_n}^h) \geq \mathbb{P}(q \in \Omega_{\mathcal{P}_n}^{h - \delta}) \geq 1 - \left(1 - \frac{\omega_k (h - \delta)^k + \theta_S(h - \delta)}{\mu(\mathcal{S})}\right)^{h - \delta}$. \square

We also need the next Lemma whose proof is deferred to Appendix C.

Lemma 3 (Bounding the Covering Number) Under the hypotheses of the previous Lemma and further assuming \mathcal{S} to be compact, we have that for any small enough $\delta > 0$:

$$\mathcal{N}(\mathcal{S}, \delta) \leq \frac{\mu(\mathcal{S})}{\omega_k (\delta/2)^k + \theta_S(\delta/2)} \quad (14)$$

Proposition 2 Let the set of hypotheses sustaining all of the previous lemmas hold. Let $([0, 1) \ni) x_h \triangleq \frac{\omega_k (h/2)^k + \theta_S(h/2)}{\mu(\mathcal{S})}$, where ω_k and θ_S are given as in the proof of Lemma 1. Then

$$\mathbb{P}(\mathcal{S} \not\subseteq \Omega_{\mathcal{P}_n}^h) \leq \frac{(1 - x_h)^n}{x_h} \quad (15)$$

Proof:

Consider a finite $\frac{h}{2}$ -net covering \mathcal{S} , that is $\mathcal{S} = \bigcup_{i=1}^{\mathcal{N}(\mathcal{S}, \frac{h}{2})} B_{\mathcal{S}}(q_i, \frac{h}{2})$, then

$$\begin{aligned}
\mathbb{P}(\mathcal{S} \not\subseteq \Omega_{\mathcal{P}_n}^h) &= \mathbb{P}\left(\bigcup_{x \in \mathcal{S}} \{x \notin \Omega_{\mathcal{P}_n}^h\}\right) \\
&= \mathbb{P}\left(\bigcup_{i=1}^{\mathcal{N}(\mathcal{S}, \frac{h}{2})} \bigcup_{x \in B_{\mathcal{S}}(q_i, \frac{h}{2})} \{x \notin \Omega_{\mathcal{P}_n}^h\}\right) \\
&\leq \mathcal{N}\left(\mathcal{S}, \frac{h}{2}\right) \max_{1 \leq i \leq \mathcal{N}(\mathcal{S}, \frac{h}{2})} \mathbb{P}\left(\bigcup_{x \in B_{\mathcal{S}}(q_i, \frac{h}{2})} \{x \notin \Omega_{\mathcal{P}_n}^h\}\right) \\
&= \mathcal{N}\left(\mathcal{S}, \frac{h}{2}\right) \max_{1 \leq i \leq \mathcal{N}(\mathcal{S}, \frac{h}{2})} \mathbb{P}\left(B_{\mathcal{S}}(q_i, \frac{h}{2}) \not\subseteq \Omega_{\mathcal{P}_n}^h\right) \\
&= \mathcal{N}\left(\mathcal{S}, \frac{h}{2}\right) \left(1 - \min_{1 \leq i \leq \mathcal{N}(\mathcal{S}, \frac{h}{2})} \mathbb{P}\left(B_{\mathcal{S}}(q_i, \frac{h}{2}) \subseteq \Omega_{\mathcal{P}_n}^h\right)\right)
\end{aligned}$$

Finally, using the Lemmas above we obtain the desired result. \square

It is interesting and useful to find a relation between n (number of points in the cloud), h (radii of the balls) and k (dimension of the manifold) which guarantees $\lim_{n \uparrow +\infty, h \downarrow 0} \mathbb{P}(\mathcal{S} \not\subseteq \Omega_{\mathcal{P}_n}^h) = 0$. For this purpose we will use the Proposition we just proved. Note that $h > 0$ will be small and also, if we are attempting to approximate $d_{\mathcal{S}}$, it should tend to 0.¹⁵

The proof of the following Lemma will be deferred to Appendix C.

Lemma 4 *Let $\{a_m\}_{m \in \mathbb{N}}$ be a sequence of positive numbers such that $a_m \downarrow 0$ as $m \uparrow \infty$ and $a_m < 1 \forall m \in \mathbb{N}$. Let $\{b_n\}_{n \in \mathbb{N}}$ be a sequence of positive numbers such that $b_n \uparrow \infty$ as $n \uparrow \infty$, then $\frac{(1-a_m)^{b_n}}{a_m} \rightarrow 0$ as $m, n \uparrow \infty$ iff $a_m \gtrsim \frac{\ln b_n}{b_n}$.*

Remark 2 *We see that the condition relating h , k and n should then be*

$$h^k \gtrsim \frac{\ln n}{n} \tag{16}$$

Also, under this condition we can estimate the rate at which $\frac{(1-x_h)^n}{x_h}$ approaches zero as $n \uparrow \infty$:

$$\frac{(1-x_h)^n}{x_h} = \frac{e^{n \ln(1-x_h)}}{x_h} \simeq \frac{e^{-nx_h}}{x_h} \lesssim \frac{1}{\ln n} \text{ as } n \uparrow \infty.$$

We have the following interesting corollary whose proof can be found in Appendix C:

Corollary 4 *Let \mathcal{S} be a smooth compact submanifold of \mathbb{R}^d without boundary. We have that if (16) holds, then for any $\varepsilon > 0$*

$$\lim_{h, n} \mathbb{P}\left(d_{\mathcal{H}}(\mathcal{S}, \Omega_{\mathcal{P}_n}^h) > \varepsilon\right) = 0$$

where $d_{\mathcal{H}}$ is the Hausdorff distance between sets.

¹⁵For constant $h > 0$, by definition $0 < x_h < 1$, then obviously $\frac{(1-x_h)^n}{x_h} \rightarrow 0$ as $n \uparrow \infty$.

We are now ready to state and prove the following convergence theorem.

Theorem 7 *Let \mathcal{S} be a k -dimensional smooth compact submanifold of \mathbb{R}^d . Let $\mathcal{P}_n = \{p_1, \dots, p_n\} \subseteq \mathcal{S}$ be such that $p_i \sim \mathbf{U}[\mathcal{S}]$ for $1 \leq i \leq n$. Then if $h = h_n$ is such that $h_n \downarrow 0$ and (16) holds as $n \uparrow \infty$, we have that for any $\varepsilon > 0$,*

$$\mathbb{P} \left(\max_{p, q \in \mathcal{S}} \left(d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(p, q) \right) > \varepsilon \right) \xrightarrow{n \uparrow \infty} 0$$

Proof:

From (7), we see that we need only to bound $\mathbb{P}(\mathcal{E}_\varepsilon \mid \mathcal{J}_{h,n})$ since $\mathbb{P}(\mathcal{J}_{h,n}^c)$ has already been dealt with by Proposition 2. To this end we use Chebyshev's inequality to obtain

$$\mathbb{P}(\mathcal{E}_\varepsilon \mid \mathcal{J}_{h,n}) \leq \frac{1}{\varepsilon} \mathbb{E} \left(\max_{p, q \in \mathcal{S}} \left(d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(p, q) \right) \mid \{\mathbb{1}_{[\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h]} = 1\} \right)$$

From considerations at the beginning of §4, $\max_{p, q \in \mathcal{S}} \left(d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(p, q) \right) \leq C_{\mathcal{S}} \sqrt{h_n}$ whenever $\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h$. Then, $\mathbb{P}(\mathcal{E}_\varepsilon) \leq C_{\mathcal{S}} \frac{\sqrt{h_n}}{\varepsilon} + \frac{(1-x_{h_n})^n}{x_{h_n}}$ which goes to 0 as $n \uparrow \infty$ according to Lemma 4 and posterior discussions. \square

Remark 3 *By simple considerations, one can see that under the hypotheses of the previous Theorem the rate at which $\mathbb{P} \left(\max_{p, q \in \mathcal{S}} \left(d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(p, q) \right) > \varepsilon \right) \xrightarrow{n \uparrow \infty} 0$ can be upper bounded by a constant times $\left(\left(\frac{\ln n}{n} \right)^{1/2k} + \frac{1}{\ln n} \right)$.*

Remark 4 *By imposing more stringent conditions on the relation between h and n and using better bounds we might be able to actually upgrade the convergence notion. However, as a first approach to the subject and in order to keep the techniques as simple as possible we have opted to work with this simple notion of random convergence.*

This concludes our study of distance functions on point clouds (sampled manifolds). We now turn to the even more real scenario where the points are considered to be *noisy* samples.

6 Noisy Sampling of Manifolds

We assume that we have some uncertainty on the actual position of the surface, and model this as if each point in the set of sampled points is modified by a (not yet random) perturbation of magnitude smaller than Δ . More explicitly, each p_i is given as $p_i = p + \zeta \times \vec{v}$ for some $\vec{v} \in S^{d-1}$, some p in \mathcal{S} and $\Delta \geq \zeta \geq 0$. Then we can guarantee that the point p from which p_i comes can be found inside $B(p_i, \Delta) \cap \mathcal{S}$. We are again interested in comparing $d_{\Omega_{\mathcal{P}_n}^h} : \Omega_{\mathcal{P}_n}^h \rightarrow \mathbb{R}^+ \cup \{0\}$ with $d_{\mathcal{S}} : \mathcal{S} \rightarrow \mathbb{R}^+ \cup \{0\}$, but now these functions have different domains, therefore we must be careful in defining a meaningful way of relating them. If we consider

$$\mathcal{F}_{\mathcal{S}}^\Delta \triangleq \{f \text{ s.t. } f : \Omega_{\mathcal{S}}^\Delta \rightarrow \mathcal{S}, f(p) \in B(p, \Delta) \cap \mathcal{S}\}$$

we can compare, for some $f \in \mathcal{F}_{\mathcal{S}}^\Delta$, and $1 \leq i, j \leq n$, $d_{\Omega_{\mathcal{P}_n}^h}(p_i, p_j)$ with $d_{\mathcal{S}}(f(p_i), f(p_j))$. Note that as the perturbation's magnitude goes to zero, $\mathcal{F}_{\mathcal{S}}^\Delta \ni f(p) \xrightarrow{\Delta \downarrow 0} p$, for $p \in \Omega_{\mathcal{S}}^\Delta$. The next step is to

write $\max_{1 \leq i, j \leq n} \left| d_{\Omega_{\mathcal{P}_n}^h}(p_i, p_j) - d_{\mathcal{S}}(f(p_i), f(p_j)) \right|$, the biggest error we have for our set of points. And finally, the next logical step is to look at the worst possible choice for f :

$$\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h) \triangleq \sup_{f \in \mathcal{F}_{\mathcal{S}}^{\Delta}} \max_{1 \leq i, j \leq n} \left| d_{\mathcal{S}}(f(p_i), f(p_j)) - d_{\Omega_{\mathcal{P}_n}^h}(p_i, p_j) \right| \quad (17)$$

We start by presenting deterministic bounds for the expression in (17), and only later will we be more (randomically) greedy, and in the spirit of Theorem 7, prove, for $\varepsilon > 0$, a result of the form $(\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h))$ will be a RV

$$\mathbb{P}(\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h) > \varepsilon) \xrightarrow{n \uparrow \infty} 0$$

6.1 Deterministic Setting

The idea is to prove that for some convenient function $\hat{f} \in \mathcal{F}_{\mathcal{S}}^{\Delta}$, we can write

$$\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h) \leq \max_{1 \leq i, j \leq n} \left| d_{\mathcal{S}}(\hat{f}(p_i), \hat{f}(p_j)) - d_{\Omega_{\mathcal{P}_n}^h}(p_i, p_j) \right| + \lambda(h, \Delta)$$

where $0 \leq \lambda(x, y) \xrightarrow{x, y \downarrow 0} 0$. The natural candidate for \hat{f} is the orthogonal projection onto \mathcal{S} , $\Pi_{\mathcal{S}} : \Omega_{\mathcal{S}}^H \rightarrow \mathcal{S}$, whose properties are discussed in Appendix B. Then, we see that we can reduce everything to bounding $\max_{p, q \in \mathcal{S}} \left| d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(p, q) \right|$. This is simple, since if $\mathcal{P}_n \subset \Omega_{\mathcal{S}}^{\Delta}$ then $\Omega_{\mathcal{P}_n}^h \subset \Omega_{\mathcal{S}}^{h+\Delta}$, and $d_{\mathcal{S}} \geq d_{\Omega_{\mathcal{P}_n}^h|_{\mathcal{S}}} \geq d_{\Omega_{\mathcal{S}}^{h+\Delta}|_{\mathcal{S}}}$, and finally from Theorem 5, $\|d_{\mathcal{S}} - d_{\Omega_{\mathcal{P}_n}^h}\|_{L^{\infty}(\mathcal{S})} \leq C_{\mathcal{S}}\sqrt{h+\Delta}$.

Let $\mathcal{S} \subset \Omega_{\mathcal{P}_n}^h$, $f \in \mathcal{F}_{\mathcal{S}}^{\Delta}$ and $1 \leq i, j \leq n$. Then, after using the triangle inequality a number of times we can write the bound,

$$\begin{aligned} \left| d_{\mathcal{S}}(f(p_i), f(p_j)) - d_{\Omega_{\mathcal{P}_n}^h}(p_i, p_j) \right| &\leq 2 \sup_{f \in \mathcal{F}_{\mathcal{S}}^{\Delta}} \max_{p \in \mathcal{P}_n} d_{\mathcal{S}}(f(p), \Pi_{\mathcal{S}}(p)) \\ &\quad + \max_{p, q \in \mathcal{S}} \left| d_{\mathcal{S}}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(p, q) \right| \\ &\quad + \max_{p, q \in \mathcal{P}_n} \left| d_{\Omega_{\mathcal{P}_n}^h}(p, q) - d_{\Omega_{\mathcal{P}_n}^h}(\Pi_{\mathcal{S}}(p), \Pi_{\mathcal{S}}(q)) \right| \end{aligned}$$

The last term can be bounded by 2Δ , the one in the middle has already been discussed, hence we are left with the first one. Using Corollary 2, we find that since $f(p) \in B(\Pi_{\mathcal{S}}(p), 2\Delta) \cap \mathcal{S}$ then in fact $f(p) \in B_{\mathcal{S}}(\Pi_{\mathcal{S}}(p), 2\Delta(1 + C_{\mathcal{S}}\sqrt{\Delta}))$, and $d_{\mathcal{S}}(f(p), \Pi_{\mathcal{S}}(p)) \leq 2\Delta(1 + C_{\mathcal{S}}\sqrt{2}\sqrt{\Delta})$. Summing up, under the condition $\mathcal{S} \subset \Omega_{\mathcal{P}_n}^h$, we obtain the desired result:

$$\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h) \leq C_{\mathcal{S}}\sqrt{h+\Delta} + 2\Delta \left(2 + \sqrt{2}C_{\mathcal{S}}\sqrt{\Delta} \right) \quad (18)$$

6.2 Random Setting

Assume that $\{p_1, \dots, p_n\}$ is a set of *i.i.d.* random points such that each $p_i \sim \mathbf{U}[\Omega_{\mathcal{S}}^{\Delta}]$. At this time, we want to estimate the probability of having $\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h$. It is easy to see that as a first “reality compliant” condition one should have that the noise level is not too big with respect to h . We’ll impose $h \geq \Delta$ for simplicity’s sake, as can be understood from the convergence theorem below. Since the techniques are similar to the ones used in the noise-free case, we will present its proof in Appendix C.

Theorem 8 *Let \mathcal{S} be a k -dimensional smooth compact submanifold of \mathbb{R}^d . Let $\mathcal{P}_n = \{p_1, \dots, p_n\}$ be such that $p_i \sim \mathbf{U}[\Omega_{\mathcal{S}}^{\Delta}]$ for $1 \leq i \leq n$. Then if $h = h_n$, $\Delta = \Delta_n$ are such that $\Delta_n \leq h_n$ and $h_n \downarrow 0$ and $\Delta_n^k \gtrsim \frac{\ln n}{n}$ as $n \uparrow \infty$, we have that for any $\varepsilon > 0$,*

$$\mathbb{P}(\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h) > \varepsilon) \xrightarrow{n \uparrow \infty} 0$$

We have now concluded the analysis of the most general case for noisy sampling of manifolds. Note that although the results in this and in previous sections were presented for Euclidean balls, they can easily be extended to more general covering shapes (check Corollary 1 above), e.g. following [14, 30], or using minimal spanning trees, or from the local directions of the data [46]. Similarly, the results can be extended to other sampling or noise models.

7 Implementation Details and Examples

We now present examples of distance functions and geodesics for point clouds, Figure 1; use these computations to find intrinsic Voronoi diagrams, Figure 2 (see also [34, 35]); and compare the results with those obtained with mesh-based techniques, Figure 3.¹⁶ We also present examples in high dimensions and use, following and extending [22], our results to compare manifolds given by point clouds. All these exercises are to exemplify the importance of computing distance functions and geodesics on point clouds, and are by no means exhaustive.

The theoretical results presented in previous sections show that the intrinsic distance and geodesics can be approximated by the Euclidean ones computed in the band defined by the union of balls centered at the points of the cloud. The problem is then simplified to first computing this band, and then use well known computationally optimal techniques to compute the distances and geodesics inside this band, exactly as done in [41] for implicit surfaces. The band itself can be computed in several ways, and for the examples below we have used constant radii. Locally adaptive radii can be used, based for example on diameters obtained from minimal spanning trees. Automatic and local estimation of h defining $\Omega_{\mathcal{P}_n}^h$ was not pursued in this paper and is the subject of current research.

The software implementation of the algorithm is based on using the fast Euclidean distance computation algorithms, usually referred to as *fast marching* algorithm [29, 51, 52, 55], twice. We omit the description of this algorithm since it is well known. The starting point is defining a grid over which all the computations are performed. This amounts to choosing Δ_{x_i} , the grid spacing in each direction $i = 1, \dots, d$. In the first round we compute the band $\Omega_{\mathcal{P}_n}^h = \{x \in \mathbb{R}^d : d(\mathcal{P}_n, x) \leq h\}$ by specifying a value of zero for the function $\Psi(x) = d(\mathcal{P}_n, x)$ on the points $x \in \mathcal{P}_n$. Since in general this points will not be on the grid, we use a simple multilinear interpolation procedure to specify the values on neighboring grid points. The second use of the fast distance algorithm is also simply reduced to using Ψ to define $\Omega_{\mathcal{P}_n}^h$ using the simple modification reported in [41]. The computation of geodesics was done using a simple Runge-Kutta gradient descent procedure, much in the way described in [41], with some obvious modifications.

All the code and 3D visualization was developed in C++ using both Flujos (which is written using Blitz++, see [31]) and VTK (see [56]). For matrix manipulation and visualization of other results we used matlab.

¹⁶All the figures in this paper are in color. VRML files corresponding to these examples can be found at mountains.ece.umn.edu/~guille/pc.htm.

7.1 High Dimensional Data

In this section we present a simple example for high dimensional data. We embedded a circle of radius 15 in \mathbb{R}^5 , and use a grid of size $34 \times 4 \times 4 \times 4 \times 34$ (with uniform spacing $\Delta x = 1$) such that each of the sample points is of the form $p_i = 15 (\cos(\frac{2\pi i}{N}), 0, 0, 0, \sin(\frac{2\pi i}{N})) + (17, 2, 2, 2, 17)$, for $1 \leq i \leq N$. We then used our approach to compute the (approximate) distance function d_h in a band in \mathbb{R}^5 , and then, the error $e_{ij} = |d_S(p_i, p_j) - d_h(p_i, p_j)|$ for $i, j \in \{1, \dots, N\}$. In our experiments we used $h = 2.5 > \Delta x \sqrt{5}$.¹⁷ We randomly sampled 500 points from the $N = 1000$ points used to construct the union of balls to build the 500×500 error matrix $((e_{ij}))$. We found $\max_{ij} \{e_{ij}\} = 2.0275$, that is a 4.3% L_∞ -error. In Figure 4 we show the histogram of all the (500^2) entries of $((e_{ij}))$. We should also note that when following the dimensionality reduction approach in [53], with the geodesic distance computation here proposed, the correct dimensionality of the circle was obtained.

In high dimensions, when the grid is too large, our current numerical implementation becomes unusable. The problem stems from the fact that we require too much memory space, most of which is not really used, since the computations are conducted only in a band around $\mathcal{P} \subset \mathbb{R}^d$. To be more precise, the memory requirements of our current direct implementation, which uses a d -dimensional array to make the computations, are $\simeq (\max_i l_i)^d$, whereas, we really need a storage capacity of order $\mu_k(\mathcal{S})h^{d-k}$, where l_i is the size of \mathcal{P} s bounding box along the i 'th direction, $1 \leq i \leq d$, and $\mu_k(\mathcal{S})$ is the measure of the k -dimensional manifold \mathcal{S} (embedded in \mathbb{R}^d). This memory problem is to be addressed by a computation that is not based on discretizing the whole band (note of course that the theoretical foundations presented in this paper are independent of the particular implementation). We are currently working on addressing this specific issue.

7.2 Object Recognition

The goal of this application is to use our framework to compare manifolds given by point clouds. The comparison is done in an intrinsic way, that is, isometrically (bending) invariant. This application is motivated by [22], where they use geodesic distances (computed using a graph based approach) to compare 3D triangulated surfaces. In contrast with [22], we compare point clouds using our framework (which, is not only based in the original raw data, but it is also, as shown in Appendix A, more robust to noise than mesh approaches as those of [22] and valid in any dimensions), and use a different procedure/similarity metric between the manifolds. The authors in [22] basically project into low dimensional manifolds and use eigenvalues and eigenvectors of a centralized matrix related to the *distance matrices* (matrices which in each entry (i, j) have the value of the intrinsic distance between (projected) points p_i and p_j of the cloud)), which are clearly not sufficient to distinguish non-isometric objects (non-isometric objects can have distance matrices with the same eigenvalues). A different study, based on direct comparisons of distance matrices, is used here and detailed in Appendix E.

Our task then is to compare two manifolds in an intrinsic way, i.e., we want to check whether they are isometric or not. We want to check this condition by using point clouds representing each one of the manifolds. Let \mathcal{S}_1 and \mathcal{S}_2 be two submanifolds of \mathbb{R}^d and sampled on each of them the two point clouds $\mathcal{P}_n^{(1)} \subset \mathcal{S}_1$ and $\mathcal{P}_n^{(2)} \subset \mathcal{S}_2$. Then, following our theory we compute the corresponding distances in the offset bands for these two sets of points, $d_{\Omega_{\mathcal{P}_n^{(1)}}^{h_1}}$ and $d_{\Omega_{\mathcal{P}_n^{(2)}}^{h_2}}$, and for points subsets $\{q_1^{(1)}, \dots, q_m^{(1)}\} = \mathcal{Q}_m^{(1)} \subseteq \mathcal{P}_n^{(1)}$, $\{q_1^{(2)}, \dots, q_m^{(2)}\} = \mathcal{Q}_m^{(2)} \subseteq \mathcal{P}_n^{(2)}$, we compute the corresponding $m \times m$

¹⁷For a discussion on how to make a preliminary estimation of the value of h see [41].

Dataset	Number of Points in the Cloud (n)	Grid size used
Bunny	15862	80 × 80 × 70
MAN2	26186	120 × 90 × 200
MAN3	26186	120 × 90 × 165
MAN5	26186	120 × 85 × 160
WOMAN2	29624	120 × 105 × 175
WOMAN3	29624	120 × 100 × 180

Table 1: Information about the models used in our recognition experiments.

pair-wise distance matrices (as defined above)

$$D_1 = \left(\left(d_{\Omega_{\mathcal{P}_n^{(1)}}^{h_1}}(q_i^{(1)}, q_j^{(1)}) \right) \right) \text{ and } D_2 = \left(\left(d_{\Omega_{\mathcal{P}_n^{(2)}}^{h_2}}(q_i^{(2)}, q_j^{(2)}) \right) \right)$$

Let \mathcal{PM}_m be the set of $m \times m$ permutation matrices and $\|\cdot\|$ a unitary transformation invariant norm¹⁸ (let's fix Frobenius norm: $\|A\| = \sqrt{\sum_i \sum_j a_{ij}^2}$). Then we define the \mathcal{J} -distance between (distance) matrices D_1 and D_2 as

$$d_{\mathcal{J}}(D_1, D_2) \triangleq \min_{P \in \mathcal{PM}_m} \|D_1 - PD_2P^T\|$$

Clearly, if $d_{\mathcal{J}}(D_1, D_2) = 0$ then we have an isometry between the discrete metric sets $(\mathcal{Q}_n^{(1)}, d_{\Omega_{\mathcal{P}_n^{(1)}}^{h_1}})$ and $(\mathcal{Q}_n^{(2)}, d_{\Omega_{\mathcal{P}_n^{(2)}}^{h_2}})$. This should allow us to establish a *rough isometry* (see [13], §4.4) between \mathcal{S}_1 and \mathcal{S}_2 with interesting constants.

The exact details on how this metric is approximated, and how the sub-sets of points \mathcal{Q} are selected is presented in Appendix E. For the experiments regarding recognition of shapes we used the datasets listed in Table 1.

In Figure 5 we present the histogram of the error $e(100)/100$ for 20 different 100×100 distance matrices corresponding to the full Bunny model, with the 100 points chosen as in the “packing procedure” described in Appendix E, where the exact definition of $e(\cdot)$ is also given (Equation (27)). We computed the mean of $e(100)/100$ over the $19 \times 18 \times \dots \times 1 = 190$ comparison-experiments to be 0.4774 with standard deviation 0.0189. This can be interpreted as when one considers a large enough set of points, the information contained in the packing set is representative of the metric information of the manifold, independently of the particular choice of the packing set. This claim needs some further theoretical justification which could come if a result of the following fashion were proved.

Conjecture 1 *Let \mathcal{S} be a smooth compact k -dimensional submanifold of \mathbb{R}^k such that its Ricci curvature is bounded below by $\kappa(n-1)$ with $\kappa \leq 0$. Let $\mathcal{Q}_m^{(r)} \subset \mathcal{S}$, $r = 1, 2$ be such that $d_{\mathcal{S}}(q_i^{(r)}, q_j^{(r)}) \geq \varepsilon$ and $B_{\mathcal{S}}(\mathcal{Q}_m^{(r)}, R)$ covers \mathcal{S} for some $R > \varepsilon > 0$. Then with D_1 and D_2 defined as before*

$$d_{\mathcal{J}}(D_1, D_2) \leq 2mC_{\mathcal{S}}\sqrt{h} + C(R, \varepsilon, m)$$

¹⁸ $\|AU\| = \|A\|$ for any matrix A and any unitary matrix U .

MODEL	Man2	Man3	Man5	Woman2	Woman3
Man2	*	0.0514	0.0570	0.4690	0.4853
Man3	*	*	0.0206	0.4701	0.4859
Man5	*	*	*	0.4702	0.4862
Woman2	*	*	*	*	0.2639
Woman3	*	*	*	*	*

Table 2: Cross comparisons for the Human models using the error measure $e(300)/300$ normalized by the maximum of the errors.

where the exact form of $C(R, \varepsilon, m)$ is to be determined, leading to an optimal choice of m (the size of the sub-set).

Using the same procedure, described in Appendix E, to choose the sets $\mathcal{Q}_m^{(i)}$ we computed the errors (according to $e(D_1, D_2)$) for 5 human artificial models, 3 of them are bendings of a man and 2 bendings of a woman, see figures 6 and 7. Details on these models are also given in Table 1. The results of this cross comparisons are presented in Table 2 below.

These examples show how our geodesic distance computation technique, when complemented with the matrix metric in Appendix E, can be used to compare manifolds given by point clouds, in a bending invariant fashion and without explicit manifold reconstruction. More exhaustive experimentation and additional theoretical justification will be reported elsewhere.

Before concluding, we should comment that as frequently done in the literature, we could normalize the geodesic distances, if scale invariance is also required. Moreover, we could also consider in the distance matrix only non-zero entries for local neighborhoods. In addition, the use of techniques for computing eigenvalues and eigenvectors such as those in Coifman *et al.*, [16], work on high-dimensional geometric multiscale analysis should be explored.

8 Concluding Remarks

In this paper, we have shown how to compute distance functions and geodesics intrinsic to a generic manifold defined by a point cloud, without the intermediate step of manifold reconstruction. The basic idea is to use well developed computational algorithms for computing Euclidean distances in an offset band surrounding the manifold, and use these to approximate the intrinsic distance. The underlying theoretical results were complemented by experimental illustrations.

As mentioned in the introduction, an alternative technique to compute geodesic distances was introduced in [8, 53] (see also [25]). In contrast with our work, the effects of noise were not addressed in [8, 25]. Moreover, as one can see from considerations in Appendix A, our framework is more robust to noise. We should note that the memory requirements of the current way of implementing our framework are large, and this needs to be addressed for very high dimensions. In particular, we are interested in direct ways of computing distances inside regions defined by union of balls, without the need to use the Hamilton-Jacobi approach.

We are currently working on the use of this framework to create multiresolution representations of point clouds (in collaboration with N. Dyn, see also [10, 18, 20, 47]), to further perform object recognition for larger libraries, and to compute basic geometric characteristics of the underlying manifold, all this of course without reconstructing the manifold (see [45] for recent results on normal computations for 2D and 3D noisy point clouds). Further applications of our framework for high

dimensional data are also currently being addressed, beyond the preliminary (toy) results reported in §7. Of particular interest in this direction is the combination of this work with the one developed by Coifman and colleagues and the recent one in [25].

Acknowledgments

We acknowledge useful conversations on the topic of this paper with L. Aspirot, P. Bermolén, R. Coifman, D. Donoho, N. Dyn, J. Giesen, O. Gil, R. Gulliver, R. Kimmel, A. Pardo and O. Zeitouni. We thank M. Levoy and the Digital Michelangelo Project for data provided for this project. This work was supported by a grant from the Office of Naval Research ONR-N00014-97-1-0509, the Presidential Early Career Award for Scientists and Engineers (PECASE), and a National Science Foundation CAREER Award. The research of F.M. is also supported by CSIC-Uruguay.

A Comparison with Mesh-based Strategies for Distance Calculation in Presence of Noise

We now make some very basic comparisons between our approach to geodesic distance computations and those based on graph approximations to the manifold, such as the one in Isomap [53, 25].¹⁹ The goal is to show that such graph-based techniques are more sensitive to noise in the point cloud sample. This is expected, since the geodesic in such techniques goes through the noisy samples, while in our approach, they just go through the union of balls. We only make our argument for the 1D case, while the high dimensional cases can be similarly studied.

A.1 1D Theoretical Case

Let's consider a rectilinear segment of length L and $n + 1$ equi-spaced points p_1, \dots, p_{n+1} in that segment. Consider the *noisy* points $q_i = p_i + \zeta_i \vec{n}$ where \vec{n} is the normal to the segment and ζ_i $1 \leq i \leq n$ are independent RV uniformly distributed in $[-\Delta, \Delta]$. Let $l = L/n$ denote the distance between adjacent p_i s. Let d_g^Δ denote the length of the polygonal path $\overline{q_1 q_2 \dots q_{n+1}}$ and $d_0 = L$. Then obviously $d_g^\Delta \geq d_0$ for any realization of the RVs ζ_i . Let $d_i = \|p_i - p_{i+1}\|$, then by Pythagoras theorem $d_i = \sqrt{l^2 + z_i^2}$, where $z_i = \zeta_i - \zeta_{i+1}$ are RVs with triangular density in $[-2\Delta, 2\Delta]$.

Then we compute $\mathbb{E}(d_i) = \frac{1}{2\Delta} \int_{-2\Delta}^{2\Delta} \sqrt{l^2 + z^2} (1 - \frac{|z|}{2\Delta}) dz$. The result is

$$\mathbb{E}(d_i) = \sqrt{l^2 + 4\Delta^2} + \frac{l^2}{2\Delta} \ln \left(\frac{2\Delta + \sqrt{l^2 + 4\Delta^2}}{l} \right) - \frac{1}{6\Delta^2} \left((l^2 + 4\Delta^2)^{3/2} - l^3 \right)$$

Now assuming $\frac{\Delta}{l} \ll 1$, we find that up to first order $\mathbb{E}(d_i) \simeq l + \Delta$, and

$$\mathbb{E}(d_g^\Delta - d_0) \simeq n\Delta$$

From this we also get²⁰

$$p_g \triangleq \mathbb{P}(d_g^\Delta - d_0 > \varepsilon) \lesssim \frac{n\Delta}{\varepsilon}$$

¹⁹Isomap builds a mesh by locally connecting the (noisy) samples.

²⁰Also, with similar arguments we can prove that $\max_{\zeta_1, \dots, \zeta_{n+1}} (d_g^\Delta - d_0) \simeq \frac{2n^2\Delta^2}{L}$.

On the other hand, for our approximation, d_h^Δ , if the segment is contained in the union of the balls centered at the sampling points, $d_h^\Delta = d_0$. The probability of covering the segment by the band can be made arbitrarily close to 1 by increasing n . More precisely, one can prove that if p stands for the value of the probability of *not covering* the segment, then $p \leq k \frac{L}{\Delta} (1 - k' \frac{\Delta}{L})^n$, for some positive constants k and k' . Then, we can write

$$p_h \stackrel{\Delta}{=} \mathbb{P} (d_h^\Delta - d_0 > \varepsilon) \leq \frac{k'' L}{\varepsilon \Delta} (1 - k' \frac{\Delta}{L})^{n+1}$$

The comparison is now easy. We see that in order to have p_g vanish as $n \uparrow \infty$, Δ must go to zero *faster* than $\frac{1}{n}$. Whereas, we know that by requiring $\Delta \simeq \frac{\ln n}{n} \gtrsim \frac{1}{n}$ we have $p_h \downarrow 0$ as $n \uparrow \infty$. This means that the graph approximation of the distance is more sensitive to noise than ours.²¹ This gives some evidence on why our approach is more robust than popular mesh-based ones. Next we present results of some simulations carried in order to further verify our claim.

A.2 Simulations

In the table below we present results of simulations carried out for the SwissRoll dataset [53], see Figure 3 . We used 10000 points to define the manifold. We then generated 10000 noise vectors, being each component uniform with power one and zero mean. Then, we generated noisy datasets from the noiseless SwissRoll dataset by adding the noise vector times a constant n_k to each vector of the noiseless initial dataset. We then chose 1000 corresponding points in each dataset and computed the intrinsic pairwise distance approximation obtaining the matrices $\{(D_{ij}^{g,n_k})\}$ and $\{(D_{ij}^{h,n_k})\}$ for the graph-based and our approach respectively, where $k = 1, 2, \dots, 5$, $i, j \in [1, 1000]$, and n_k denotes the noise level. We then computed the values of $\max_{ij} |D_{ij}^{g,n_k} - D_{ij}^{g,0}|$ and $\max_{ij} |D_{ij}^{h,n_k} - D_{ij}^{h,0}|$ for each k , where $D_{ij}^{g,0}$ and $D_{ij}^{h,0}$ stand for noiseless intrinsic distance approximations. In the table, h indicates the radii and k the size of the neighborhood for Isomap. The graph approximation shows less robustness to noise than our method, as was argued above. This is also true for the sensitivity,²² where our approach outperforms the graph-based one by at least one order of magnitude. Note that the sensitivity for our approach can be formally studied from Theorem 3.

Noise Power (n_k^2)	$\max_{ij} D_{ij}^{g,n_k} - D_{ij}^{g,0} $	k	$\max_{ij} D_{ij}^{h,n_k} - D_{ij}^{h,0} $	h
0.0001	2.5222	7	0.5266	1.8
0.01	4.6409	7	0.9430	1.8
0.04	5.1737	7	1.2489	1.8
0.09	5.3292	7	1.4682	1.8
0.16	5.4651	7	1.7965	1.8

B Properties of Euclidean Distance Functions

The references for this section are [2], pp.12-16, and [23].

Theorem 9 ([2]) *Let $\Gamma \subset \mathbb{R}^d$ be a compact, smooth manifold without boundary. Then $\eta(x) \stackrel{\Delta}{=} \frac{1}{2}d^2(\Gamma, x)$ is smooth in a tubular neighborhood U of Γ . Also, in U it satisfies $\|D\eta\|^2 = 2\eta$.*

²¹Another way of seeing this is noting that for a fixed noise level Δ , by increasing n we actually worsen the graph approximation, whereas we are making our approximation better.

²²Sensitivity is defined as $\left|1 - \frac{\text{distance for noisy points}}{\text{distance for clean points}}\right|$.

Corollary 5 *The projection operator $\Pi : U \rightarrow \Gamma$, for a given $x \in U$, can be written as $\Pi(x) = x - D\eta(x)$. Moreover, this operator is smooth.*

Remark 5 *Differentiation of the relation $\langle D\eta, D\eta \rangle = 2\eta$ gives us $D^2\eta D\eta = D\eta$. Differentiating once more we also find $D^3\eta D\eta = D^2\eta$.*

Theorem 10 ([2]) *Let Γ and U be as in the previous Theorem and let $y \in U$ and $x = y - D\eta(x) \in \Gamma$, $k = \dim(\Gamma)$. Then, denoting by $\lambda_1, \dots, \lambda_n$ the eigenvalues of $D^2\eta(y)$,*

$$\lambda_i(y) = \begin{cases} \frac{d(\Gamma, y)\kappa_i(x)}{1+d(\Gamma, y)\kappa_i(x)} & \text{if } 1 \leq i \leq k \\ 1 & \text{if } k < i \leq n \end{cases}$$

where $\kappa_i(x)$ are the principal curvatures of Γ at x along $Dd(\Gamma, y) \in N_x\Gamma$, where $N_x\Gamma$ is the normal space to Γ at x .

C Deferred Proofs

Proof: (Corollary 3)

We present only a sketch of the proof. Let \mathcal{M} be an extension of \mathcal{S} such that \mathcal{S} is still strongly convex in \mathcal{M} . Then we choose h sufficiently small so as to have a well defined projection operator from $\Omega_{\mathcal{S}}^h$ to \mathcal{M} . Next, for any $x, y \in \mathcal{S}$ consider γ_h the $\Omega_{\mathcal{S}}^h$ minimizing geodesic, $\mathbf{L}(\gamma_h) = d_h(x, y)$. As in Theorem 5 and using the convexity of \mathcal{S} in \mathcal{M} we have $\mathbf{L}(\gamma_h) = d_h(x, y) \leq d_{\mathcal{S}}(x, y) = d_{\mathcal{M}}(x, y) \leq \mathbf{L}(\Pi(\gamma_h))$, since $\Pi(\gamma_h) \subset \mathcal{M}$ but may not be a minimizing path. Then $0 \leq d_{\mathcal{S}}(x, y) - d_h(x, y) \leq |\mathbf{L}(\Pi(\gamma_h)) - \mathbf{L}(\gamma_h)|$ and from here on, if γ_h were $C^{1,1}([0, 1], \mathbb{R}^d)$ we could proceed as in Theorem 5. Since there exists a unique $\gamma_0 \subset \mathcal{S}$ joining x, y , by Theorem 4, γ_h uniformly converges to γ_0 . But since $\Pi(\gamma_h) \rightrightarrows \gamma_h$ too, then $\Pi(\gamma_h) \rightrightarrows \gamma_0$. Now, if we show that the interior of $\Pi(\gamma_h)$ is contained in the interior of \mathcal{S} we would have the degree of smoothness of γ_h that we need (γ_h would only be touching those parts of $\partial\Omega_{\mathcal{S}}^h$ which are smooth). But from the uniform convergence of $\Pi(\gamma_h)$ towards γ_0 and the fact that the interior of γ_0 is contained in the interior of \mathcal{S} the conclusion follows. \square

Proof: (Lemma 3)

We now estimate the covering number $\mathcal{N}(\mathcal{S}, \delta)$. The idea is constructive, very simple and of course standard. We consider the following procedure (adopted from [8]): Let q_1 be any point in \mathcal{S} , then choose $q_2 \in \mathcal{S} \setminus B_{\mathcal{S}}(q_1, \delta)$. Then choose $q_3 \in \mathcal{S} \setminus \{B_{\mathcal{S}}(q_1, \delta) \cup B_{\mathcal{S}}(q_2, \delta)\}$. Iterate this procedure until it is not longer possible to choose any point $q \in \mathcal{S} \setminus \{\cup_{k=1}^{\mathcal{N}(\mathcal{S}, \delta)} B_{\mathcal{S}}(q_k, \delta)\}$, in such a case $\mathcal{S} = \cup_{k=1}^{\mathcal{N}(\mathcal{S}, \delta)} B_{\mathcal{S}}(q_k, \delta)$. Note that $B_{\mathcal{S}}(q_k, \frac{\delta}{2}) \cap B_{\mathcal{S}}(q_l, \frac{\delta}{2}) = \emptyset$ if $k \neq l$, therefore we can bound $\mathcal{N}(\mathcal{S}, \delta) \leq \frac{\mu(\mathcal{S})}{\min_{x \in \mathcal{S}} \mu(B_{\mathcal{S}}(x, \delta/2))}$ and therefore, using again Bishop-Günther inequalities in the same manner as in Lemma 1, we find (14). \square

Proof: (Lemma 4)

Taking logarithms we see that the condition we want to satisfy is equivalent to requiring that $b_n \ln(1 - a_m) - \ln b_n \rightarrow -\infty$. Since $a_m \downarrow 0$, $\ln(1 - a_m) \simeq -a_m$ for large enough m , we should then look at $b_n a_m - \ln(a_m) \rightarrow \infty \Leftrightarrow b_n a_m + \ln(b_n a_m) - \ln n \rightarrow \infty$. Now, the last term pushes towards $-\infty$, hence the product $b_n a_m$ has no choice but to tend to ∞ . Then $\ln(b_n a_m)$ becomes negligible when compared to $b_n a_m$ for large enough n and m . Therefore the condition is now $b_n a_m - \ln b_n \rightarrow \infty$, that is $b_n(a_m - \frac{\ln b_n}{b_n}) \rightarrow \infty$ what concludes the argument. \square

Proof: (Corollary 4)

Note first that the random variable $d_{\mathcal{H}}(\mathcal{S}, \Omega_{\mathcal{P}_n}^h)$ is bounded by $\max\{\text{diam}(\mathcal{S}) + h, h\}$. By definition of Hausdorff distance, $d_{\mathcal{H}}(\mathcal{S}, \Omega_{\mathcal{P}_n}^h) = \max\left(\sup_{x \in \mathcal{S}} d(x, \Omega_{\mathcal{P}_n}^h), \sup_{y \in \Omega_{\mathcal{P}_n}^h} d(y, \mathcal{S})\right)$. Then, $\sup_{x \in \mathcal{S}} d(x, \Omega_{\mathcal{P}_n}^h) \leq \text{diam}(\mathcal{S}) + h$ by the triangle inequality, and $\sup_{y \in \Omega_{\mathcal{P}_n}^h} d(y, \mathcal{S}) \leq h$, trivially.

Now, we can write $\mathbb{E}(d_{\mathcal{H}}(\mathcal{S}, \Omega_{\mathcal{P}_n}^h)) = \mathbb{E}\left(\mathbb{E}\left(d_{\mathcal{H}}(\mathcal{S}, \Omega_{\mathcal{P}}^h) \Big| \mathbb{1}_{[\mathcal{S} \subseteq \Omega_{\mathcal{P}}^h]}\right)\right)$, but the inner expected value can be bounded by h , when $\mathbb{1}_{[\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h]} = 1$, and by $\max\{\text{diam}(\mathcal{S}) + h, h\}$, when $\mathbb{1}_{[\mathcal{S} \subseteq \Omega_{\mathcal{P}}^h]} = 0$. Using Chebyshev's inequality we find

$$\begin{aligned} \mathbb{P}(d_{\mathcal{H}}(\mathcal{S}, \Omega_{\mathcal{P}_n}^h) > \delta) &\leq \frac{h}{\delta} \mathbb{P}(\{\mathcal{S} \subseteq \Omega_{\mathcal{P}}^h\}) + \frac{\max\{\text{diam}(\mathcal{S}) + h, h\}}{\delta} (1 - \mathbb{P}(\{\mathcal{S} \subseteq \Omega_{\mathcal{P}}^h\})) \\ &\leq \frac{h}{\delta} + \frac{\text{diam}(\mathcal{S}) + h}{\delta} (1 - \mathbb{P}(\{\mathcal{S} \subseteq \Omega_{\mathcal{P}}^h\})) \end{aligned}$$

a quantity that goes to zero for any fixed $\delta > 0$ as $h \downarrow 0$ and $n \uparrow \infty$ provided (16) holds. \square

Proof: (Theorem 8)

Since the proof is almost identical to that of Theorem 7, many steps will be skipped. Note that since \mathcal{S} is compact, there exists an upper bound K for all its sectional curvatures. This will allow us to use the volume comparison theorems as before.

We can start from the adequate version of equation (7). We must bound both $\mathbb{P}(\{\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h\}^c)$ and $\mathbb{P}(\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h) > \varepsilon \mid \{\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h\})$. The second term can be bounded in an identical way as its $\Delta = 0$ counterpart was, obtaining

$$\mathbb{P}(\mathcal{L}_{\mathcal{S}}(\mathcal{P}_n; \Delta, h) > \varepsilon \mid \{\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h\}) \leq \frac{C_{\mathcal{S}} \sqrt{h + \Delta} + 2\Delta (2 + \sqrt{2} C_{\mathcal{S}} \sqrt{\Delta})}{\varepsilon} \quad (19)$$

which vanishes as $n \uparrow \infty$.

Now we upper bound $\mathbb{P}(\{\mathcal{S} \subseteq \Omega_{\mathcal{P}_n}^h\}^c)$. Everything carries over in the same fashion as in the proof of Lemma 1, except that now we must take into consideration that the p_i s are not necessarily on \mathcal{S} but inside $\Omega_{\mathcal{S}}^{\Delta}$. Following the described steps we obtain

$$\mathbb{P}(\{x \notin \Omega_{\mathcal{P}_n}^h \cap \mathcal{S}\}) \leq \left(1 - \frac{\mu(B(x, h) \cap \Omega_{\mathcal{S}}^{\Delta})}{\mu(\Omega_{\mathcal{S}}^{\Delta})}\right)^n \quad (20)$$

Notice that since we are working with $h \geq \Delta$, we have $B(x, \Delta) \subset B(x, h) \cap \Omega_{\mathcal{S}}^{\Delta}$ (see Figure 8) and we can rewrite the bound in (20) as

$$\mathbb{P}(\{x \notin \Omega_{\mathcal{P}_n}^h \cap \mathcal{S}\}) \leq \left(1 - \frac{\mu(B(x, \Delta))}{\mu(\Omega_{\mathcal{S}}^{\Delta})}\right)^n \quad (21)$$

$$= \left(1 - \frac{\mu(B(\cdot, \Delta))}{\mu(\Omega_{\mathcal{S}}^{\Delta})}\right)^n \quad (22)$$

We can bound this quantity using formulas akin to Weyl's Tube Theorem. More precisely, as explained in Appendix D, we can write

$$\mu(\Omega_{\mathcal{S}}^{\Delta}) = \mu(\mathcal{S}) v(d - k, \Delta) + \varphi_{\mathcal{S}}(\Delta)$$

where $\frac{\varphi_{\mathcal{S}}(\Delta)}{\Delta^{d-k+1}} \rightarrow 0$ as $\Delta \rightarrow 0$ and $v(D, R)$ is the volume of the ball of radius R in D -dimensional Euclidean space.

Now, for $x \in \mathcal{S}$ we must find a bound for $\mathbb{P}(B_{\mathcal{S}}(x, \delta) \not\subseteq \Omega_{\mathcal{P}_n}^h)$, but like in the proof of Lemma 2, $\mathbb{P}(B_{\mathcal{S}}(x, \delta) \not\subseteq \Omega_{\mathcal{P}_n}^h) \leq \mathbb{P}(x \notin \Omega_{\mathcal{P}_n}^{h-\delta} \cap \mathcal{S})$, which can be bounded by (21). Also the bound (14) for the covering number still works in this case, then we can write $\mathbb{P}(\mathcal{S} \not\subseteq \Omega_{\mathcal{P}_n}^h \cap \mathcal{S}) \leq \frac{(1 - y_{\Delta})^n}{x_h}$, where $y_{\Delta} \triangleq \frac{\mu(B(\cdot, \Delta))}{\mu(\Omega_{\mathcal{S}}^{\Delta})}$. Also since $h \geq \Delta$, $x_h \geq x_{\Delta}$, then

$$\mathbb{P}(\mathcal{S} \not\subseteq \Omega_{\mathcal{P}_n}^h) \leq \frac{(1 - y_{\Delta})^n}{x_{\Delta}}$$

But with Δ small enough, $y_\Delta \simeq \alpha \Delta^k$ and $x_\Delta \simeq \beta \Delta^k$, then Lemma 4 and the hypotheses guarantee that $\mathbb{P}(\{\mathcal{S} \subset \Omega_{\mathbb{P}_n}^h\}^c) \rightarrow 0$ as $n \uparrow \infty$. \square

D Basic Differential Geometry Facts

In this section we collect some facts that were used throughout the article, following [26].

D.1 Measure of a d -dimensional Ball

Recall the definition of the Γ function:

$$\Gamma(\alpha) = \int_0^{+\infty} e^{-t} t^{\alpha-1} dt$$

Theorem 11 *The volume of d -dimensional ball of radius r is given by*

$$v(d, r) \triangleq \mu(B(\cdot, r)) = \omega_d r^d$$

where $\omega_d = \frac{2\pi^{d/2}}{d\Gamma(d/2)}$

D.2 Bishop-Günther Inequalities for the Measure of a Geodesic Ball

Theorem 12 *Let \mathcal{S} be a complete k -dimensional Riemannian manifold, assume r to be smaller than the distance between $m \in \mathcal{S}$ and $\text{Cut}(m, \mathcal{S})$ (cut locus of the points m in \mathcal{S}). Let $K^{\mathcal{S}}$ be the sectional curvatures of \mathcal{S} and γ a constant. Then if $\widehat{V}_\gamma(r) \triangleq \frac{2\pi^{k/2}}{\Gamma(k/2)} \int_0^r \left(\frac{\sin(t\sqrt{\gamma})}{\sqrt{\gamma}}\right)^{k-1} dt$*

$$K^{\mathcal{S}} \geq \gamma \text{ implies } \mu(B_{\mathcal{S}}(m, r)) \leq \widehat{V}_\gamma(r) \quad (23)$$

$$K^{\mathcal{S}} \leq \gamma \text{ implies } \mu(B_{\mathcal{S}}(m, r)) \geq \widehat{V}_\gamma(r) \quad (24)$$

Proposition 3 *We have the following Taylor expansion for $\widehat{V}_\gamma(r)$, the volume of a geodesic ball in a space of constant sectional curvature γ :*

$$\widehat{V}_\gamma(r) = \omega_k r^k \left(1 - r^2 \frac{\gamma}{6} \frac{k(k-1)}{k+2}\right) + \phi(r)$$

where $\frac{\phi(r)}{r^{k+2}} \rightarrow 0$ as $r \downarrow 0$.

D.3 Weyl's Tube Theorem

Theorem 13 *Let \mathcal{S} be a k -dimensional manifold topologically embedded in \mathbb{R}^d . Assume that \mathcal{S} is compact closure, and that every point in the tube $T(\mathcal{S}, r) = \{x \in \mathbb{R}^d \text{ s.t. } d(\mathcal{S}, x) \leq r\}$ has a unique shortest geodesic connecting it with \mathcal{S} , then the volume $\mu(T(\mathcal{S}, r))$ of the tube is given by*

$$\mu(T(\mathcal{S}, r)) = r \frac{\sqrt{\pi}}{\Gamma(3/2)} \sum_{i=0}^{\lfloor \frac{d-1}{2} \rfloor} \frac{k_{2i}(\mathcal{S}) r^{2i}}{I(i)} \quad (25)$$

where $I(i) = 1 \cdot 3 \cdot 5 \cdot \dots \cdot (2i+1)$ and the numbers k_{2i} depend on the curvature structure of \mathcal{S} . For our purposes we only need know that $k_0 = \mu(\mathcal{S})$.

Corollary 6 *The volume of the tube $T(\mathcal{S}, r)$ can be expanded as*

$$\mu(T(\mathcal{S}, r)) = \mu(\mathcal{S}) v(d - k; r) + \phi_{\mathcal{S}}(r)$$

where $\frac{\phi_{\mathcal{S}}(r)}{r^{d-k}} \xrightarrow{r \downarrow 0} 0$

E Details on Object Recognition

The ideal objective is to actually compute the \mathcal{J} -distance between D_1 and D_2 as described in §7.2, however, this is a very hard problem since there are $m!$ $m \times m$ permutation matrices. The choice of m is subject to compromise, on one hand we want it to be big enough so as to capture the metric structure of \mathcal{S}_i with the information given by $(\Omega_n^{(i)}, d_{\Omega_n^{(i)}})$, on the other hand we want to be able to actually make the computations involved without too much processing cost. Therefore we should attempt to circumvent this $m!$ search space by exploiting some other information we might have.

One possibility to bypass this difficulty is to try to upper bound the \mathcal{J} -distance by some difference between eigenvalues of the matrices. However, it turns out that one can easily find two distance matrices which have positive \mathcal{J} -distance (they are not *cogredient*) but have the same spectra. Then, an upper bound should take into account also other term that measures our inability to really differentiate distance functions by only looking at their eigenvalues. Of course this information must then be contained in the eigenvectors.²³

A way of dealing with this particular issue is working with the spectral factorization of each of the matrices. Let $D_1 = QDQ^T$ and $D_2 = \hat{Q}\hat{D}\hat{Q}^T$, where Q and \hat{Q} are unitary matrices and D and \hat{D} are diagonal matrices whose entries are the eigenvalues of D_1 and D_2 respectively. Note that we are not saying anything about the order in which those eigenvalues are presented, for convenience let $D_{11} = |D_{11}| \geq |D_{22}| \geq \dots \geq |D_{mm}|$ and $\hat{D}_{11} = |\hat{D}_{11}| \geq |\hat{D}_{22}| \geq \dots \geq |\hat{D}_{mm}|$.²⁴ Then we can write with little effort

$$\min_{P \in \mathcal{PM}_m} \|D_1 - PD_2P^T\| \leq \|(Q - P\hat{Q})D\| + \|(Q - P\hat{Q})\hat{D}\| + \|D - \hat{D}\| \quad (26)$$

Note that if W is any matrix and T is diagonal, then $\|WT\|^2 = \sum_k \|W_{(:,k)}t_{kk}\|^2 = \sum_k \|W_{(:,k)}\|^2 t_{kk}^2$ where $W_{(:,k)}$ is the k -th column vector of W . Using this observation, we note that the first two terms in (26) can be bounded as follows (let $Q = (q_1 | \dots | q_m)$, and $\hat{Q} = (\hat{q}_1 | \dots | \hat{q}_m)$):

$$\|(Q - P\hat{Q})D\| + \|(Q - P\hat{Q})\hat{D}\| = \sqrt{\sum_k D_{kk}^2 \|q_k - P\hat{q}_k\|^2} + \sqrt{\sum_k \hat{D}_{kk}^2 \|q_k - P\hat{q}_k\|^2}$$

Now, using the trivial inequality $\frac{\sqrt{a} + \sqrt{b}}{2} \leq \sqrt{\frac{a+b}{2}} \forall a, b \geq 0$, we finally arrive at the expression

$$\min_{P \in \mathcal{PM}_m} \|D_1 - PD_2P^T\| \leq \sqrt{\sum_{k=1}^m (D_{kk} - \hat{D}_{kk})^2} + \sqrt{2} \sqrt{\sum_{k=1}^m (D_{kk}^2 + \hat{D}_{kk}^2) \|q_k - P\hat{q}_k\|^2} \quad (27)$$

²³Another idea for example is the following: We know that the sought for isometry (if it exists) must be a Lipschitz continuous map, therefore it makes no sense to consider the huge set of transformations spanned by \mathcal{PM}_m . We leave the exploitation of this idea for future work.

²⁴We've used Frobenius Theorem [42], which asserts that nonnegative matrices have a positive largest absolute value eigenvalue.

This inequality holds for any $P \in \mathcal{PM}_m$. It is important to note that in case D_1 and D_2 are cogredient, then all their eigenvectors will also be related through that same permutation, therefore this inequality is sharp.²⁵

Note that in the second term of (27), the values of $\|q_i - P\hat{q}_i\|$ are weighted by $(D_{ii}^2 + \hat{D}_{ii}^2)$, so one can think that since $\|q_i - P\hat{q}_i\| \leq 2$, the most important terms of the sum will be those for which $(D_{ii}^2 + \hat{D}_{ii}^2)$ is large. This is not a rigorous consideration, but gives some guidelines on how to compute an approximate bound when the sizes of the distance matrices are prohibitively large.

In some situations, the choice of the sub-sampled set size m that guarantees a good metric approximation in the sense discussed above, might be too large, making the computation of the full bound (27) onerous. But still a measure of similarity must be provided which does not require the computation of all of the eigenvalues and eigenvectors of each distance matrix. Therefore, in order to estimate $d_J(D_1, D_2)$, we use the following idea: Instead of computing all the eigenvalues and eigenvectors of the matrices D_1 and D_2 , compute the $N \ll m$ more important ones, where important means, in the light of the expression for the bound, those with the largest moduli, at least for the part of the bound involving eigenvectors. Then, for a (computationally) reasonable N we define the approximate error bound (we still let P be any convenient choice of a permutation matrix)

$$e(N) \triangleq \sqrt{\sum_{k=1}^N (D_{kk} - \hat{D}_{kk})^2} + \sqrt{2} \sqrt{\sum_{k=1}^N (D_{kk}^2 + \hat{D}_{kk}^2) \|q_k - P\hat{q}_k\|^2} \quad (28)$$

Now, we fix the permutation P as follows: Let S be the permutation matrix such that Sq_1 is a column vector whose components are sorted from largest to smallest. Do the same with \hat{q}_1 to obtain \hat{S} , then we would compare Sq_1 with $\hat{S}\hat{q}_1$, what amounts to comparing q_1 with $S^T\hat{S}$, hence we let $P = S^T\hat{S}$. We could again use a more sophisticated way of choosing P but this one suffices for our demonstration purposes and of course, achieves equality in (27) when both matrices are cogredient.

Another possibility is to directly compare the distance matrices according to the expression $\|D_1 - PD_2P^T\|$ using a certain sensible choice for P . We first put both matrices in a ‘‘canonical’’ order. Let (i_1, j_1) be one position on the matrix D_1 with the maximum value. We then order the rest of the points in the set according to their distances to either $q_{i_1}^{(1)}$ or $q_{i_2}^{(1)}$ from smallest to largest.²⁶ This induces an ordering for the matrix D_1 , being P_1 be the underlying permutation matrix. We do the same with D_2 and obtain P_2 . Finally we let

$$e_G(D_1, D_2) \triangleq \|D_1 - P_1^T P_2 D_2 P_2^T P_1\|$$

and note that obviously $d_J(D_1, D_2) \leq e_G(D_1, D_2)$ and that the inequality is sharp.

E.1 Choice of the Point Cloud Subset $Q^{(i)}$

In general the number of points in the cloud is too big. This means that the actual computation of the distance matrices, if done using all the points in the cloud, and subsequent eigenvalues and

²⁵Note that from (26) one can obtain, $d_J(D_1, D_2) \leq \|D - \hat{D}\| + (\|D\| + \|\hat{D}\|) \|Q\hat{Q}^T - P\|$, then one further idea to be explored is how to best approximate a given unitary matrix by a permutation matrix. This would allow us not only to obtain an explicit bound for the \mathcal{J} -distance, but it would also provide us with a low metric distortion way of mapping S_1 ($\mathcal{P}_n^{(1)}$) into S_2 ($\mathcal{P}_n^{(2)}$), with applications like texture mapping, brain warping, etc.

²⁶In our current implementation we don't worry about repeated distance values, since this can be easily handled. We will present more details and further refinements elsewhere.

eigenvectors computations (if needed) become onerous. Therefore we need a procedure which allows us to select a small cardinality subset \mathcal{Q}_m of \mathcal{P}_n for which we will actually compute the approximate distance matrix, but still using \mathcal{P}_n to define the offset $\Omega_{\mathcal{P}_n}^h$ inside which the computations are performed. This subset $\mathcal{C}_r \subset \mathcal{P}_n$ must be “representative” of the geometry of the underlying manifold. One way of selecting those points is by not allowing them to cluster inside any region of the manifold. This can be accomplished in practise using the “packing idea” in [22]: Given $m < n$, choose the first point $c_1 \in \mathcal{C}_m$ randomly, then proceed by always choosing a point as far as possible from the set of points that have already been chosen. End the process when m points have been chosen. This is the procedure used in the experiments.

References

- [1] R. Alexander & S. Alexander, “Geodesics in Riemannian manifolds with boundary,” *Indiana University Mathematics Journal* **30:4**, 1981.
- [2] L. Ambrosio and H.M. Soner, “Level set approach to mean curvature flow in arbitrary codimension,” *J. Diff. Geom.* **43**, pp. 693-737, 1996.
- [3] N. Amenta, S. Choi, and R. Kolluri, “The power crust, unions of balls, and the medial axis transform,” *Computational Geometry: Theory and Applications* **19**, pp. 127-153, 2001.
- [4] N. Amenta, S. Choi, and R. Kolluri, “The power crust,” *Proceedings of 6th ACM Symposium on Solid Modeling*, pp. 249-260, 2001.
- [5] N. Amenta and R. Kolluri, “Accurate and efficient unions of balls,” *ACM Symposium on Computational Geometry*, pp. 119-128, 2000.
- [6] T. Apostol, *Mathematical Analysis*, Addison-Wesley Series in Mathematics, 1974.
- [7] A. Bartesaghi and G. Sapiro “A System for the Generation of Curves on 3D Brain Images”. *Human Brain Mapping*, Volume 14, Issue 1, 2001.
- [8] M. Bernstein, V. de Silva, J. Langford, and J. Tenenbaum, “Graph approximations to geodesics on embedded manifolds,” <http://isomap.stanford.edu/BdSLT.ps>
- [9] **Blitz++** website: www.oonumerics.org/blitz
- [10] J-D. Boissonnat and F. Cazals, “Coarse-to-fine surface simplification with geometric guarantees,” in A. Chalmers and T-M. Rhyne, Editors, *EUROGRAPHICS '01*, Manchester, 2001.
- [11] M. Botsch, A. Wiratanaya, and L. Kobbelt, “Efficient high quality rendering of point sampled geometry,” *EUROGRAPHICS Workshop on Rendering*, 2002.
- [12] E. Calabi, P. J. Olver, and A. Tannenbaum, “Affine geometry, curve flows, and invariant numerical approximations,” *Adv. in Math.* **124**, pp. 154-196, 1996.
- [13] I. Chavel, *Riemannian Geometry: A Modern Introduction*, Cambridge University Press, 1993.
- [14] R. Coifman, personal communication, December 2002.
- [15] R. Coifman, personal communication, January 2003 (talk presented at IPAM-UCLA).

- [16] R. Coifman, personal communication, April 2003 (talk presented at University of Minnesota).
- [17] T. K. Dey, J. Giesen, S. Goswami, and W. Zhao, "Shape dimension and approximation from samples," *Discrete and Comput. Geom.* **29**, pp. 419-434, 2003.
- [18] T. K. Dey, J. Giesen, and J. Hudson, "Decimating samples for mesh simplification," *Proc. 13th Canadian Conference on Computational Geometry*, pp. 85-88, 2001.
- [19] D. L. Donoho and C. Grimes, "When does ISOMAP recover the natural parametrization of families of articulated images?," *Technical Report 2002-27: Department of Statistics, Stanford University*, 2002.
- [20] N. Dyn, M. S. Floater, and A. Iske, "Adaptive thinning for bivariate scattered data," *Journal of Computational and Applied Mathematics* **145(2)**, pp. 505-517, 2002.
- [21] M. doCarmo, *Riemannian Geometry*. Birkhäuser Boston, 1992.
- [22] A. Elad (Elbaz) and R. Kimmel, "Bending invariant representations for surfaces," *In Proc. of CVPR'01*, Hawaii, Dec. 2001.
- [23] H. Federer, "Curvature measures," *Trans. Amer. Math. Soc.* **93**, pp. 418-491, 1959.
- [24] M. S. Floater and A. Iske, "Thinning algorithms for scattered data interpolation," *BIT Numerical Mathematics* **38**, pp. 705-720, 1998.
- [25] J. Giesen and U. Wagner, "Shape dimension and intrinsic metric from samples of manifolds with high co-dimension," *ACM Symposium on Computational Geometry*, San Diego, June 2003.
- [26] A. Gray, *Tubes*, Addison Wesley, 1990.
- [27] M. Gross et al, *Point Based Computer Graphics*, EUROGRAPHICS Lecture Notes, 2002 (graphics.stanford.edu/~niloy/research/papers/ETH/PointBasedComputerGraphics_TutorialNotes.pdf).
- [28] R. Gulliver. Personal communication, March 2003.
- [29] J. Helmsen, E. G. Puckett, P. Collela, and M. Dorr, "Two new methods for simulating photolithography development in 3D," *Proc. SPIE Microlithography IX*, pp. 253, 1996.
- [30] P. W. Jones, "Rectifiable sets and the traveling salesman problem," *Invent. Math.* **102**, pp. 1-15, 1990.
- [31] Flujos website. <http://iie.fing.edu.uy/investigacion/grupos/gti/flujos/flujos.html>
- [32] R. Kimmel and J. A. Sethian, "Computing geodesic paths on manifolds," *Proc. National Academy of Sciences* **95:15**, pp. 8431-8435, 1998.
- [33] A. N. Kolmogorov and S. V. Fomin, *Elements of the Theory of Functions and Functional Analysis*, Dover, 1999.
- [34] R. Kunze, F. E. Wolter, and T. Rausch, "Geodesic Voronoi diagrams on parametric surfaces," *IEEE*, 1997.
- [35] G. Leibon and D. Letscher, "Delaunay triangulations and Voronoi diagrams for Riemannian manifolds," *Computational Geometry 2000*, Hong Kong, 2000.

- [36] G. Lerman, “How to partition a low-dimensional data set into disjoint clusters of different geometric structure,” preprint, 2000.
- [37] L. Linsen, “Point cloud representation,” *CS Technical Report*, University of Karlsruhe, 2001.
- [38] L. Linsen and H. Prautzsch, “Local versus global triangulations,” *EUROGRAPHICS '01*, 2001.
- [39] C. Mantegazza and A.C. Menucci, *Hamilton-Jacobi Equations and Distance Functions on Riemannian Manifolds*. Obtained from: <http://cvgmt.sns.it/papers/manmen99/>
- [40] A. Marino and D. Scolozzi, “Geodetiche con ostacolo,” *Boll. Un. Mat. Ital.* **6:2-B**, pp. 1-31, 1983.
- [41] F. Mémoli and G. Sapiro, “Fast computation of weighted distance functions and geodesics on implicit hyper-surfaces,” *Journal of Computational Physics* **173:2**, pp. 730-764, 2001.
- [42] H. Minc, *Nonnegative Matrices*, Wiley Interscience Series in Discrete Mathematics and Optimization, 1988.
- [43] J. S. B. Mitchell, “An algorithmic approach to some problems in terrain navigation,” *Artificial Intelligence* **37**, pp. 171-201, 1988.
- [44] J. S. B. Mitchell, D. Payton, and D. Keirse, “Planning and reasoning for autonomous vehicle control,” *International Journal of Intelligent Systems* **2**, pp. 129-198, 1987.
- [45] N. J. Mitra and A. Nguyen, “Estimating surface normals in noisy point cloud data,” *ACM Symposium on Computational Geometry*, San Diego, June 2003.
- [46] M. Pauly and M. Gross, “Spectral processing of point-sampled geometry,” *ACM SIGGRAPH*, pp. 379-386, 2001
- [47] M. Pauly, M. Gross, and L. Kobbelt, “Efficient simplification of point-sampled surfaces,” *IEEE Visualization 2002*.
- [48] S. Rusinkiewicz and M. Levoy, “QSplat: A multiresolution point rendering system for large meshes,” *Computer Graphics (SIGGRAPH 2000 Proceedings)*, 2000.
- [49] T. Sakai, *Riemannian Geometry*, AMS Translations of Mathematical Monographs **149**, 1996.
- [50] E. Schwartz, A. Shaw, and E. Wolfson, “A numerical solution to the generalized mapmaker’s problem: Flattening nonconvex polyhedral surfaces,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11:9**, pp. 1005 -1008, 1989.
- [51] J. Sethian, “Fast marching level set methods for three-dimensional photolithography development,” *Proc. SPIE International Symposium on Microlithography*, Santa Clara, California, March, 1996.
- [52] J. A. Sethian, “A fast marching level-set method for monotonically advancing fronts,” *Proc. Nat. Acad. Sci.* **93:4**, pp. 1591-1595, 1996.
- [53] J. B. Tenenbaum, V. de Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction,” *Science*, pp. 2319-2323, December 2000.

- [54] R. Tsai, L.-T. Cheng, S. Osher, and H.-K. Zhao, “Fast sweeping algorithms for a class of Hamilton-Jacobi equations,” *UCLA-CAM Report*, October 2001 (<http://www.math.ucla.edu/applied/cam/index.html>).
- [55] J. N. Tsitsiklis, “Efficient algorithms for globally optimal trajectories,” *IEEE Transactions on Automatic Control* **40** pp. 1528-1538, 1995.
- [56] VTK website.www.vtk.org
- [57] F. E. Wolter, “Cut loci in bordered and unbordered Riemannian manifolds,” *Doctoral Dissertation, Technische Universität Berlin*, 1985.
- [58] M. Zwicker, M. Pauly, O. Knoll, and M. Gross, “PointShop 3D: An interactive system for point-based surface editing,” *SIGGRAPH*, 2002.







Figure 1: *Top: Intrinsic distance function for a point cloud. A point is selected in the head of the David, and the intrinsic distance is computed following the framework here introduced. The point cloud is colored according to their intrinsic distance to the selected point, going from bright red (close) to dark blue (far). The offset band, given by the union of balls, is shown next to the distance figure. Bottom: Same as before, with a geodesic curve between two selected points. (This is a color figure.)*



Figure 2: *Voronoi diagram for point clouds. Four points (left) and two points (right) are selected on the cloud, and the point cloud is divided (colored) according to their geodesic distance to these four points. Note that this is a surface Voronoi, based on geodesics computed with our proposed framework, not an Euclidean one. (This is a color figure.)*




Figure 3: *Examples of geodesic computations. This data is used to study the algorithm robustness to noise, see Appendix A. (This is a color figure.)*




Figure 4: *Histogram for the error in the case of a circle embedded in \mathbb{R}^5 .*




Figure 5: *Histogram showing the errors for different selections of point clouds on the Bunny model.*

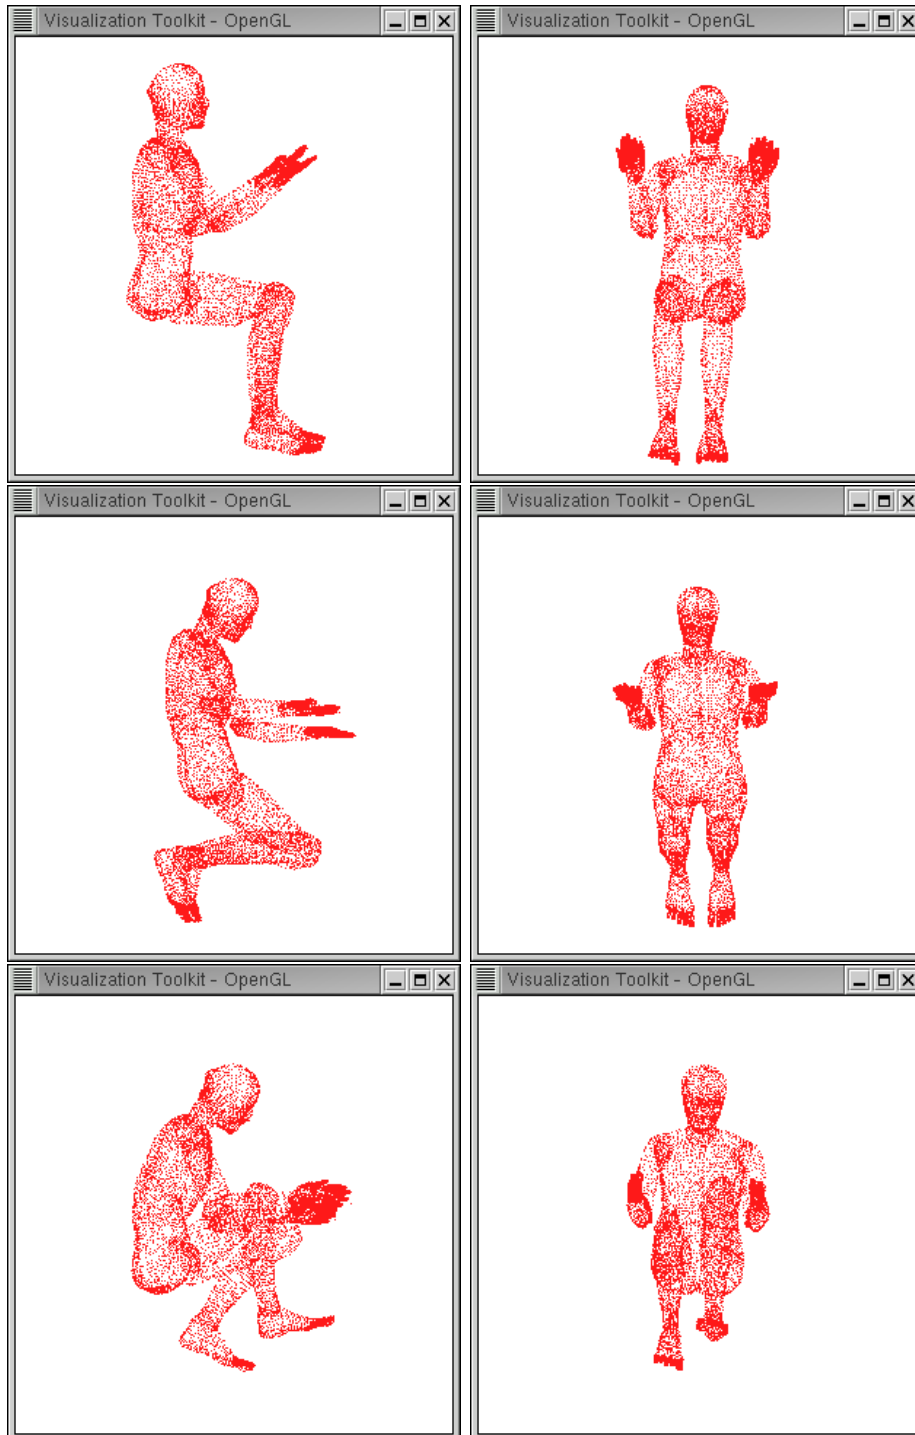


Figure 6: *MAN* models. From top to bottom (two views of each model): *MAN2*, *MAN3*, and *MAN5*.




Figure 7: *WOMAN* models. From top to bottom (two views of each model): *WOMAN2* and *WOMAN3*.




Figure 8: $B(x, \Delta) \subset B(x, h) \cap \Omega_s^\Delta$