

Shape Dimension and Approximation from Samples*

Tamal K. Dey, Joachim Giesen, Samrat Goswami, and Wulue Zhao

Department of CIS, Ohio State University,
Columbus, OH 43210, USA
{tamaldey,giesen,goswami,zhaow}@cis.ohio-state.edu

Abstract. There are many scientific and engineering applications where an automatic detection of shape dimension from sample data is necessary. Topological dimensions of shapes constitute an important global feature of them. We present a Voronoi-based dimension detection algorithm that assigns a dimension to a sample point which is the topological dimension of the manifold it belongs to. Based on this dimension detection, we also present an algorithm to approximate shapes of arbitrary dimension from their samples. Our empirical results with data sets in three dimensions support our theory.

1. Introduction

Interpretation of shapes from its samples is needed in many scientific and engineering applications. As a result, defining and detecting features that contain useful information about the shape are subjects of active research in shape modeling [4], [14], [22], [23]. A global feature of a shape is its dimension which has been defined in various ways to accommodate intricacies and varieties in shapes. Topological and Hausdorff dimensions are two such definitions that capture the global feature, the first one stresses on the space connectivity and the second one on the space filling property [15], [17], [20]. In this work we focus on the topological dimension of the shapes that are smooth manifolds embedded in a Euclidean space. Data collected for scientific analysis through natural phenomena or simulations lie on such manifolds and can reveal important information if the underlying dimension is detected automatically.

Automatic *dimension detection* is a major challenge in the fields of learning theory, pattern recognition [5] and artificial intelligence in general [21], [25]. In these applications point samples can be generated from an otherwise unknown manifold. The

* This work is partially supported by NSF under Grants CCR-9988216 and DMS-0138456 (subcontract: Stanford University).

dimension of the manifold when detected from the samples, provides useful feedback to the learning process and sometimes can be used for various matching purposes in pattern recognition. A related problem studied in the literature is *dimensionality reduction*, where a low-dimensional representation of a point set is sought possibly with some small distortions to the input [19], [24]. Principal component analysis (PCA) and the multidimensional scaling (MDS) are two prevalent methods used for this problem [8]. These methods approximate the dimension after allowing small distortions to the input. We stress that this paper studies the problem of exact dimension detection rather than approximating it.

In two and three dimensions, the topological dimension of a shape becomes obvious with a visual inspection of the data points. However, often a shape needs to be computed from the samples for further processing. For example, in the study of dynamical systems, the quasiperiodic and chaotic orbits can be automatically recognized from sample points in the phase portraits if automatic shape recognition can be performed [25]. To this end one may try reconstruction algorithms for curves [2], [11], [18] and for surfaces [1], [3], [6], [13]. Unfortunately, these reconstruction algorithms are of no use if the dimension of the shape is not determined a priori. For example, a surface reconstruction algorithm cannot produce a curve out of a sample that has been derived from a curve. The situation becomes worse when data have samples from shapes of various dimensions such as the data from a phase portrait of a sampled dynamical system. In this case, one first needs to separate the samples according to their dimensions and then apply respective reconstruction steps. The α -shapes pioneered by Edelsbrunner et al. [12], [13] can reconstruct shapes of arbitrary dimension from their point samples. However, this algorithm requires that the input sample density be globally uniform whereas our method requires only locally uniform density. Also, the theoretical guarantees of α -shapes in detecting dimensions have not been studied.

In this paper we present an algorithm that determines the dimension of the shape from which samples are derived. Specifically, it assigns a dimension k to a sample point if it belongs to a manifold of dimension k . We also present a shape approximation algorithm which filters simplices from a Delaunay complex approximating the sampled shape subsequent to the dimension detection step. The assumption we make on the sampling is that it is feature dependent, i.e., it is dense wherever small features exist and is relatively sparse wherever intricate features are absent. This is similar to the ε -sampling condition introduced in [1] and [2]. However, as argued later, we need a stronger assumption on sampling for automatic dimension detection and this assumption cannot be avoided if one wants to guarantee correctness.

Our approach is based on the Voronoi diagrams of the samples. The dimension of a small neighborhood of a point in a manifold depends on the dimension of the manifold itself. For example, a point on a curve has a neighborhood homeomorphic to a 1-disk and a point on a surface has a neighborhood homeomorphic to a 2-disk. Under an appropriate sampling condition, the structure of the Voronoi cells of the sample points contain information about the dimension and the shape of these neighborhoods. We analyze these structures carefully and extract this information. We test our algorithms on a number of data sets in three dimensions. Empirical results support the theory based on which we design the dimension detection algorithm DIMENSION and the shape approximation algorithm COCONESHAPE.

2. Sampling and Voronoi Geometry

Let $\mathcal{M} = \{M_1, M_2, \dots, M_\ell\}$ be a collection of smooth compact manifolds embedded in \mathbb{R}^d and let $M = |\mathcal{M}|$ be its underlying space. In general, \mathcal{M} can have manifolds of different dimensions, i.e., M_i , $i \in \{1, 2, \dots, \ell\}$, may be a k -manifold for any k , $1 \leq k \leq d$. In order to determine the dimensions of the manifolds from their samples, we must assume some density condition on the samples, otherwise they may be too sparse to provide any meaningful information about M .

2.1. Sample Density

We adapt the density assumption of Amenta et al. [2]. It is based on the *local feature size* of a shape which measures the distance to the *medial axis*. The medial axis of M is the closure of the set of points in \mathbb{R}^d that have two or more closest points on M . The local feature size is a function $f: M \rightarrow \mathbb{R}$ where $f(x)$ is the Euclidean distance to the medial axis of a point $x \in M$. Intuitively, $f(x)$ is a measure of how complicated M is in the neighborhood of x . A nice property of $f(\cdot)$ is that it is 1-Lipschitz, i.e., $f(x) \leq f(y) + \|x - y\|$ for any two points x and y in M .

A sample $P \subset M$ is called an ε -sample of M if each point $x \in M$ has a sample point $p \in P$ within an $\varepsilon f(x)$ distance. This means that each point in M has a sample within an ε factor of its local feature size. We argue that this ε -sampling assumption is not suitable for dimension detection even for small values of ε since it cannot prevent ambiguity.

Consider the sample shown in Fig. 1. The sample is taken from the surface of a pear and satisfies the ε -sampling condition for $\varepsilon < 0.4$. However, if it is not known that the sample is from a surface, the points in circular arrangements near the top may as well be taken as a sample of the respective circles. Any dimension detection algorithm would be confused in this case as to concluding if the samples are generated from a surface or a set of curves. In general, an ε -sample of a manifold M can be arranged so that a subset of the points samples a lower-dimensional manifold $M' \subset M$. This prompts us to adopt a stricter sampling condition to disambiguate the dimension of the underlying shapes.




Fig. 1. An ambiguous sample.

Definition 2.1. A sample P of a collection of smooth manifolds \mathcal{M} is called an (ε, δ) -sample of M for $\varepsilon/2 \leq \delta < \varepsilon < 1$ if each point $x \in M$ has a sample point within an $\varepsilon f(x)$ distance and each sample point p has all other sample points at least $\delta f(p)$ away from it.

Remark. The above definition requires that the sample be dense with respect to the local feature size, but at the same time sample points cannot be arbitrarily close to form an arbitrary pattern such as curves on surfaces. Notice that this requirement is much less strict than the condition where the sampling has to be uniform everywhere. We chose the lower bound on δ as $\varepsilon/2$ for making further calculations precise though any other constant greater than 1 will be equally valid for our theoretical analysis. Similar sampling conditions have been studied before in the context of surface reconstruction in [9] and [16].

2.2. Tangent and Normal Spaces

Since we are dealing with a collection \mathcal{M} of smooth manifolds in \mathbb{R}^d , results from differential geometry ensure that a tangent space at each point $p \in M$ is well defined [7]. The dimension of the tangent space at p coincides with the topological dimension of the manifold containing p . We also refer to this dimension as the *dimension of p* . Our goal is to determine the dimension of the sample points from their Voronoi cells.

Let T_p denote the *tangent space* at p , i.e., the affine subspace of \mathbb{R}^d spanned by tangent vectors at p . The affine subspace spanned by normals to T_p constitute the *normal space* N_p at p . The dimensions of N_p and T_p add to d .

The main tool we use for detecting dimension is the geometric structure of the Voronoi cells. Let V_P and D_P denote the Voronoi diagram and its dual, the Delaunay triangulation, for a sample P . A Voronoi cell V_p for a sample point p is the set of all points that are no farther away from p than any other sample point in P . Let \tilde{T}_p denote the set of all points in T_p that are no farther away from p than from any other sample point, i.e., $\tilde{T}_p = T_p \cap V_p$. Similarly, let $\tilde{N}_p = N_p \cap V_p$ denote the set of points in N_p that are no farther away from p than from any other sample point. We call \tilde{T}_p and \tilde{N}_p the *tangent* and *normal polytopes* of p , respectively. The main observation, based on which the dimension detection proceeds, is that the Voronoi cell V_p of $p \in P$ approximates the Minkowski sum of \tilde{T}_p and \tilde{N}_p . Figure 2 illustrates the above fact. This figure shows the Voronoi cells of points of various dimensions in an (ε, δ) -sample of M in three dimensions. In the left picture the tangent polytope is the segment going through p and the normal polytope is shaded. This is the case for a sample point on a curve. In the middle picture the tangent polytope is shaded and the normal polytope is the line segment going through p . This is the case for a sample point on a surface. In the right picture the tangent polytope is the entire Voronoi cell and the normal polytope is the single point p . This is the case for an interior point in a solid.

2.3. Voronoi Subpolytopes

The Voronoi cell V_p contains information about the dimension of \tilde{T}_p and hence T_p . Our task is to separate \tilde{T}_p from \tilde{N}_p from the Voronoi cell V_p . We achieve this by examin-




Fig. 2. Tangent and normal polytopes of a sample on a curve (left), surface (middle) and solid (right).

ing the structure of a sequence of *Voronoi subpolytopes* as defined recursively below. The definition of Voronoi subpolytopes is based on the concept of *poles* for Voronoi cells which was introduced by Amenta and Bern [1]. We generalize it here for Voronoi subpolytopes. For convenience we use the notation $\angle(\mathbf{v}, \mathbf{w})$ to denote the *acute* angle between the lines supporting two vectors \mathbf{v} and \mathbf{w} .

Definition 2.2. The Voronoi subpolytopes for a sample point $p \in M$ are special subsets $V_p^i \subseteq V_p$, $i = 1, \dots, d$, of the Voronoi cell V_p . Let $V_p^d = V_p$. Assume that V_p^i is already defined. Let v_p^i be the farthest point in V_p^i from p . We call v_p^i the *pole* of V_p^i and the vector $\mathbf{v}_p^i = v_p^i - p$ its *pole vector*. If V_p^i is unbounded, v_p^i is taken at infinity, and the direction of \mathbf{v}_p^i is taken as the average of all directions given by unbounded edges. The Voronoi subpolytope V_p^{i-1} is the minimal polytope containing all points $\{x: \angle((x - p), \mathbf{v}_p^i) = \pi/2\}$.

Clearly, $V_p^{i-1} \subset V_p^i$ is a polytope orthogonal to the pole vector of V_p^i . In Fig. 3 we show the Voronoi subpolytopes for a sample point on a curve and on a surface, respectively in three dimensions.

We introduce the definition of *height* to measure the structure of Voronoi subpolytopes.

Definition 2.3. The height H_p^i of a Voronoi subpolytope V_p^i is the length $\|\mathbf{v}_p^i\| = \|v_p^i - p\|$.

The height of V_p^i measures its elongation, but we also need a measure for the “fatness” of V_p^i which is measured by the height of V_p^{i-1} .




Fig. 3. Voronoi subpolytopes for a point on a curve (left) and a surface (right). V_p^3 are the entire Voronoi cells, V_p^2 are the shaded polygons and V_p^1 are the line segments going through p .

3. Geometric Guarantees

Let $p \in P$ be an interior point in a manifold of dimension k , $1 \leq k \leq d$, where P is an (ε, δ) -sample of M , the underlying space of a collection of manifolds \mathcal{M} in \mathbb{R}^d . We assume that p satisfies the following condition.

Assumption 3.1. V_p does not contain any boundary point of M .

One way to satisfy the above assumption for all sample points is to consider only manifolds without boundary. Of course, this excludes d -manifolds since they necessarily have boundaries in \mathbb{R}^d .

We claim that V_p^k approximates the tangent polytope \tilde{T}_p and thus cannot have large height. All other higher-dimensional Voronoi subpolytopes are long and thin along some normal direction and they have large heights. The claims are based on the following two observations:

- (i) All Voronoi subpolytopes V_p^i , $k < i \leq d$, contain a boundary point of \tilde{N}_p .
- (ii) The pole vectors of Voronoi subpolytopes V_p^i , $k < i \leq d$, approximate some vector in the normal space N_p .

Fact (i) implies that the height H_p^i is at least as large as $f(p)$, the local feature size at p , for $k < i \leq d$. Fact (ii) implies that the affine space of V_p^k is normal to k orthogonal vectors each of which approximates a vector in the normal space N_p . This in turn means that the affine space of V_p^k lies close to T_p and V_p^k approximates the tangent polytope \tilde{T}_p . We extend a result of [3] to show that the diameter of the tangent polytope is small, specifically $O(\varepsilon)f(p)$. It follows that the height H_p^i is $O(\varepsilon)f(p)$ for $1 \leq i \leq k$.

Facts (i) and (ii) hold vacuously for $k = d$ since $k < d$ is not satisfied. In this case $V_p^k = V_p$ is the tangent polytope and its height is small due to the ε -sampling condition.

Lemma 3.1. *Let $v \in V_p^i$ be any point in the Voronoi subpolytope V_p^i where $k < i \leq d$ and $\|v - p\| > \mu f(p)$. Then there is a normal vector $\mathbf{n}_p \in N_p$ so that*

$$\angle((v - p), \mathbf{n}_p) \leq \arcsin \frac{\varepsilon}{\mu(1 - \varepsilon)} + \arcsin \frac{\varepsilon}{1 - \varepsilon}.$$

Proof. If $v \in \tilde{N}_p$, then the condition is trivially satisfied. So, assume $v \notin \tilde{N}_p$. First we establish that there is a point $z \in M$ so that the line going through v and z intersects $\tilde{N}_p \cap V_p$, the boundary of the normal polytope \tilde{N}_p . Consider the cone $X_v = \{x: x = tv + (1 - t)y\}$ for all $t \in [0, 1]$ and $y \in \tilde{N}_p$. Since \tilde{N}_p has dimension $d - k$, the dimension of X_v is $d - k + 1$. The manifold M intersects X_v at p transversally and $L = M \cap X_v$ must be a manifold of dimension $(d - k + 1) + k - d = 1$. We argue that L must intersect the boundary of X_v at a point $z \notin \tilde{N}_p$. If $z \in \tilde{N}_p$, then there is a medial axis point violating the sampling condition. To see this consider a ball tangential to M at p growing towards z , i.e., its center moving along pz from p towards z . It must touch M at another point before or when it reaches z implying that there is a medial axis point within the segment connecting z and p . This means z has the nearest sample point p at a distance more than $f(z)$, a contradiction to our sampling condition. Thus, z can lie only on $bd(X_v) \setminus \tilde{N}_p$ establishing our claim. See Fig. 4 for an illustration.

Let the line going through v and z meet the boundary of \tilde{N}_p at m . Clearly, $m \in bd(V_p)$. Consider the triangle pvm . Take \mathbf{n}_p as the vector $r - p$ where r is the other endpoint of the line segment in \tilde{N}_p going through p and m . We are interested in the angle $\angle rpv$. We have

$$\angle rpv = \angle pmv + \angle pvm.$$

Since $z \in M$ has p as the nearest sample point we have $\|z - p\| \leq \varepsilon f(z)$. Applying the Lipschitz condition between $f(z)$ and $f(p)$ we get $f(z) \leq (\varepsilon/(1 - \varepsilon))f(p)$. Therefore,

$$\|z - p\| \leq \frac{\varepsilon}{1 - \varepsilon} f(p).$$




Fig. 4. The polytope X_v in Lemma 3.1 is three-dimensional for a curve point (left) and is two-dimensional for a surface point p (right) in three dimensions. The edges of this polytope are dashed. \tilde{N}_p is the lightly shaded polygon for the curve point, and it is the segment between m and r for the surface point.

Also, $\|m - p\| \geq f(p)$ since there is an empty ball touching M at p with its center on the medial axis and on the segment pm . To see this, again consider growing a ball touching M at p and moving its center from p towards m until it touches another point on M . The center of such a ball must be in V_p since the ball is empty. It follows that

$$\begin{aligned} \angle pmv &\leq \arcsin \frac{\|z - p\|}{\|m - p\|} \\ &\leq \arcsin \frac{\varepsilon}{1 - \varepsilon}. \end{aligned}$$

Considering the triangle pvm we have

$$\begin{aligned} \angle pvm &\leq \arcsin \frac{\|z - p\|}{\|v - p\|} \\ &\leq \arcsin \frac{\varepsilon}{\mu(1 - \varepsilon)}. \end{aligned}$$

This establishes the claim of the lemma. \square

Lemma 3.2. $H_p^i \geq f(p)$ for $k < i \leq d$.

Proof. The Voronoi subpolytope V_p^i intersects the normal polytope \tilde{N}_p in a polytope of dimension $i + (d - k) - d = i - k$ which is greater than 0 since $i > k$. It means that V_p^i has a point, say x , where $x \in bd(V_p^i) \cap \tilde{N}_p \subset bd(V_p) \cap \tilde{N}_p$. Consider growing a ball touching M at p always keeping the center on px until it touches M at another point. The center of this grown ball is a medial axis point that must lie in V_p since the ball is empty. Thus, $\|x - p\| \geq f(p)$. In particular, the length of the pole vector $H_p^i = \|\mathbf{v}_p^i - p\|$ must be at least $f(p)$. \square

In what follows we use $\alpha = 2 \arcsin \varepsilon / (1 - \varepsilon)$.

Corollary 3.1. *There exists a normal vector $\mathbf{n}_p \in N_p$ for each pole vector \mathbf{v}_p^i , $k < i \leq d$, so that $\angle(\mathbf{v}_p^i, \mathbf{n}_p) \leq \alpha$.*

Proof. From Lemma 3.2 we get $\|\mathbf{v}_p^i\| \geq f(p)$ for $k < i \leq d$. So, plugging $\mu = 1$ into Lemma 3.1 we obtain that $\angle(\mathbf{v}_p^i, \mathbf{n}_p) \leq 2 \arcsin \varepsilon / (1 - \varepsilon) = \alpha$ for some vector $\mathbf{n}_p \in N_p$. \square

Lemma 3.3. *Let $x \in V_p^i$ be any point on the boundary of the Voronoi subpolytope V_p^i , where $i \leq k$.*

$$\frac{\delta}{2} f(p) \leq \|x - p\| \leq \frac{\varepsilon}{1 - \varepsilon} \sec\left(\frac{\alpha}{2}(1 + 4\sqrt{d - k})\right) f(p).$$

Proof. The lower bound on $\|x - p\|$ follows from the δ -condition in the sampling density. Since any sample point must be at least $\delta f(p)$ away from p , any point on the boundary of V_p is at least $(\delta/2)f(p)$ distance away from p . The point x lies on the boundary of V_p^i and hence on the boundary of V_p .

For the upper bound we first exclude the case $k = d$. In this case $V_p^k = V_p$ is the tangent polytope with height $H_p^d \leq \varepsilon f(p)$ due to the sampling condition. Thus, the upper bound holds trivially in this case. For the other cases when $k < d$ we first show that the vector $\mathbf{x}_p = x - p$ makes a large angle with any vector in the normal space N_p . Let $\mathbf{n}_p \in N_p$ be a vector so that $\angle(\mathbf{x}_p, \mathbf{n}_p)$ is the smallest.

From Lemma 3.1 it follows that for each pole vector \mathbf{v}_p^i , $k < i \leq d$, there exists a vector \mathbf{n}_p^i in the normal space so that $\angle(\mathbf{v}_p^i, \mathbf{n}_p^i) \leq \alpha$, where $\alpha = 2 \arcsin \varepsilon (1 - \varepsilon)$ is small. This result and the fact that the vectors \mathbf{v}_p^i are orthogonal to each other imply that the vectors \mathbf{n}_p^i are linearly independent and thus span N_p . Consider the vector $\mathbf{v}_p = \sum_{i=k+1}^d a_i \mathbf{v}_p^i$ if $\mathbf{n}_p = \sum_{i=k+1}^d a_i \mathbf{n}_p^i$. Without loss of generality assume that the vectors $\mathbf{v}_p^i, \mathbf{n}_p^i$ for $i = k+1, \dots, d$ are unit. We can also choose $\sum_{i=k+1}^d a_i^2 = 1$ which only scales \mathbf{n}_p . With these choices \mathbf{v}_p becomes a unit vector. Let $\angle(\mathbf{v}_p, \mathbf{n}_p) = \beta$. We show that β is small.

We can write $\mathbf{n}_p^i = \mathbf{v}_p^i + \mathbf{w}^i$ where \mathbf{w}^i is a small vector with $\|\mathbf{w}^i\| \leq \alpha$ since \mathbf{v}_p^i is a unit vector and $\angle(\mathbf{v}_p^i, \mathbf{n}_p^i) \leq \alpha$. Therefore,

$$\begin{aligned} \mathbf{n}_p &= \sum_{i=k+1}^d a_i \mathbf{n}_p^i \\ &= \sum_{i=k+1}^d a_i \mathbf{v}_p^i + \sum_{i=k+1}^d a_i \mathbf{w}^i \\ &= \mathbf{v}_p + \sum_{i=k+1}^d a_i \mathbf{w}^i. \end{aligned}$$

It follows that

$$\begin{aligned} \sin \beta &\leq \left\| \sum_{i=k+1}^d a_i \mathbf{w}^i \right\| \\ &\leq \sum_{i=k+1}^d |a_i| \cdot \|\mathbf{w}^i\| \\ &\leq \alpha \sum_{i=k+1}^d |a_i|. \end{aligned}$$

The term $\sum_{i=k+1}^d |a_i|$ is at most $\sqrt{d-k}$ by Cauchy-Schwartz's inequality with the condition $\sum_{i=k+1}^d a_i^2 = 1$. Therefore, for small β ,

$$\frac{\beta}{2} \leq \sin \beta \leq \alpha \sqrt{d-k}.$$

Now consider the angle between \mathbf{x}_p and \mathbf{n}_p . By the triangular inequality

$$\begin{aligned}\angle(\mathbf{x}_p, \mathbf{n}_p) &\geq \angle(\mathbf{x}_p, \mathbf{v}_p) - \angle(\mathbf{v}_p, \mathbf{n}_p) \\ &= \frac{\pi}{2} - \beta.\end{aligned}$$

Using the contrapositive of Lemma 3.1 we obtain that $\|x - p\| = \mu f(p)$ where

$$\begin{aligned}\mu &\leq \left(\frac{\varepsilon}{1-\varepsilon}\right) \frac{1}{\sin(\pi/2 - \beta - \arcsin \varepsilon/(1-\varepsilon))} \\ &\leq \left(\frac{\varepsilon}{1-\varepsilon}\right) \frac{1}{\cos(\beta + \alpha/2)} \\ &\leq \left(\frac{\varepsilon}{1-\varepsilon}\right) \sec\left(\frac{\alpha}{2}(1 + 4\sqrt{d-k})\right),\end{aligned}$$

completing the proof. \square

Combining Lemmas 3.2 and 3.3 we obtain the following theorem.

Theorem 3.1. *Let $p \in P$ be any point derived from a manifold of dimension k embedded in \mathbb{R}^d where P is an (ε, δ) -sample and V_p does not contain any boundary point. Then the following conditions hold:*

1. $H_p^i \geq f(p)$ for $k < i \leq d$.
2. $(\delta/2)f(p) \leq H_p^i \leq \varepsilon/(1-\varepsilon) \sec((\alpha/2)(1 + 4\sqrt{d-k}))f(p)$ for $1 \leq i \leq k$.

4. Dimension Detection

Theorem 3.1 is the basis of our algorithm for dimension detection. From Theorem 3.1 we have $H_p^1 = O(\varepsilon)f(p)$ if we assume that the dimension of p is at least one. Certainly, H_p^1 is small if ε is sufficiently small. On the other hand, $H_p^i \geq f(p)$ are large for $k < i \leq d$. Then one possible algorithm to determine k may proceed as follows. Compare the heights H_p^i with H_p^1 in decreasing sequence of i starting with $i = d$. Continue the sequence of comparisons as long as the ratio H_p^1/H_p^i is less than a user chosen parameter ρ . The comparison stops when the ratio goes above ρ and we note the value of the corresponding i . This strategy has the potential problem that the heights H_p^i for $1 \leq i \leq k$ are small compared with $f(p)$, but they may not be small with respect to H_p^1 . The remedy is obtained by the δ -condition on the (ε, δ) -sampling. We note that,

$$\begin{aligned}\frac{H_p^1}{H_p^i} &\geq \frac{\delta(1-\varepsilon)}{2\varepsilon \sec((\alpha/2)(1 + 4\sqrt{d-k}))} \\ &> \frac{(1-\varepsilon)}{4 \sec((\alpha/2)(1 + 4\sqrt{d-k}))} \\ &= \Omega(1) \quad \text{for } 1 \leq i \leq k,\end{aligned}$$

and for $k < i \leq d$,

$$\begin{aligned} \frac{H_p^1}{H_p^i} &\leq \left(\frac{\varepsilon}{1-\varepsilon} \right) \sec \left(\frac{\alpha}{2} (1 + 4\sqrt{d-k}) \right) \\ &= O(\varepsilon). \end{aligned}$$

Thus, a choice of ρ in the range

$$\left[\left(\frac{\varepsilon}{1-\varepsilon} \right) \sec \left(\frac{\alpha}{2} (1 + 4\sqrt{d-k}) \right), \frac{(1-\varepsilon)}{4 \sec((\alpha/2)(1 + 4\sqrt{d-k}))} \right]$$

suffices for our algorithm DIMENSION. So, for example, in dimension five, with ε less than 0.03, $\rho \in [0.04, 0.2]$ will suffice.

DIMENSION (P, ρ, d)

```

1  Compute  $V_p$ 
2  for all  $p \in P$ 
3    compute  $H_p^1, \dots, H_p^d$ 
4     $i := d$ 
5    while  $H_p^1/H_p^i < \rho$ 
6       $i := i - 1$ 
7    endwhile
8     $\dim(p) := i$ 
9  endfor

```

5. Shape Approximation

We use DIMENSION to design a shape approximation algorithm that can approximate M with a piecewise linear complex interpolating the sample P . This algorithm is dimension independent in that it does not need to know a priori the dimension of the shape from which the sample is derived. The algorithm can be seen as a generalization of our COCONE algorithm that reconstructs surfaces in three dimensions [3], [10]. We need the following definition of *cocone* for a sample point p .

Definition 5.1. Let $p \in P$ be a sample point from a manifold of dimension k . The cocone for p , C_p , is defined as the set of all points $x \in V_p$ so that the segment connecting x and p makes an acute angle less than $\pi/8$ with V_p^k . See Fig. 5 for examples in two and three dimensions.

The approximation algorithm COCONESHAPE filters a subset of k simplices incident to p from the Delaunay triangulation of P . In the COCONE algorithm for surface reconstruction we selected a set of Delaunay triangles incident to p that are dual to the Voronoi edges intersected by C_p . Generalizing this idea, we compute the set of k -simplices incident to p that are dual to the $(d-k)$ -dimensional Voronoi faces intersected by C_p .




Fig. 5. Cocone for a curve point in two dimensions (left) and a surface point in three dimensions (right).

All such simplices computed over all sample points constitute what we call the set of *candidate simplices*. We claim that all candidate simplices lie close to the sampled manifold M .

COCONESHAPE (P, ρ)

- 1 Compute V_p
- 2 DIMENSION(P, ρ, d)
- 3 $T := \emptyset$
- 4 **for** all $p \in P$
- 5 Let $k := \dim(p)$
- 6 Compute V_p^k and C_p
- 7 Compute F , the $(d - k)$ -dimensional Voronoi faces intersected by C_p
- 8 $T := T \cup \text{dual}(F)$
- 9 **endfor**
- 10 Output T

Lemma 5.1. *Let t be a candidate simplex computed by COCONESHAPE. There is a ball circumscribing t with radius $O(\varepsilon)f(p)$ where p is any vertex of t .*

Proof. It follows from the definition of cocone C_p and the proof of Lemma 3.3 that any point $x \in C_p$ satisfies the condition $\angle((x - p), \mathbf{n}_p) \geq 3\pi/8 - O(\varepsilon)$ where \mathbf{n}_p is an anglewise closest vector to $x - p$ in N_p . Using the contrapositive of Lemma 3.1, this implies that $\|x - p\| = O(\varepsilon)f(p)$. Since t is dual to a Voronoi face F that intersects C_p , there exists a ball centered at a point $c = C_p \cap F$ with radius $O(\varepsilon)f(p)$ which circumscribes t . \square

The above lemma implies that each point on an output simplex has a point p on the manifold M within a small distance of $O(\varepsilon)f(p)$. Also, each point in M has a vertex p on

the output complex within an $\varepsilon f(\rho)$ distance. These two facts imply that the Hausdorff distance between M and T is small relative to the feature size.

6. Experimental Results

We experimented with DIMENSION on three-dimensional data sets. DIMENSION is implemented in C++ using the computational geometry algorithms library CGAL [27]. We found that robust Delaunay triangulation/Voronoi diagram computations in the presence of degeneracies and finite precision arithmetic are absolutely necessary for valid output. To this end we used the filtered floating point arithmetic of CGAL. This simulates exact arithmetic only on a demand basis. Thus, it provides the advantage of exact arithmetic with a nominal increase in running time which is observed to be a factor of two in our case. The reported running times are due to experiments performed on a PC with a 933 MHz Pentium III processor and 512 MB main memory. We tested our implementation of DIMENSION on various datasets. See Figs. 6 and 7 and Table 1. It turned out that a value of $\rho = 0.3$ gives good results in practice. The results we report here are all obtained using this value.

We developed the algorithm DIMENSION for manifolds without boundaries, but tested it also on manifolds with boundaries. In three dimensions we can have one and two manifolds without boundaries. However, 3-manifolds embedded in three dimensions necessarily have boundaries. The boundary of a k -manifold is a $(k - 1)$ -manifold, i.e., the boundary of a volume is a surface. The samples on this surface have Voronoi cells elongated along the normal to the surface. Thus, DIMENSION detects these points as lying on a 2-manifold. Similarly, we observe that points on the boundary of a surface are detected as lying on a 1-manifold. See the datasets FOOT, ENGINE and BALL in Fig. 6.

The shapes are approximated with our COCONESHape algorithm once the dimensions are detected by DIMENSION. In three dimensions we can extract a manifold out of the candidate simplices which can be proved to be homeomorphic to the original curve or surface. Our COCONE software works on this principle to reconstruct surfaces in three dimensions [26]. With the manifold extraction step COCONESHape is exactly what COCONE does for surface reconstruction. Thus, COCONESHape acts like COCONE for samples that are assigned dimension two. It includes samples from the surfaces of the three manifolds. In Fig. 6 BALL is an example of a sample from a three manifold in




Fig. 6. Output of the algorithm DIMENSION on the datasets CURVES, FOOT, ENGINE and BALL. Points classified one-dimensional by the algorithm are colored red (leftmost), points classified two-dimensional are colored green (middle two) and points classified three-dimensional are colored blue (rightmost), respectively.




Fig. 7. Output of the algorithm COCONESHAPE on the dataset SCENE (top) which has objects of different dimensions and the reconstruction of the same dataset (bottom). This picture needs to be seen in color.

Table 1. Experimental data.

Object	Number of points	DIMENSION time (sec.)	COCONESHape time (sec.)
CURVES	259	2	2
BALL	805	4	1
CACTUS	3,280	46	47
ENGINE	11,360	237	138
FOOT	20,021	110	106
SCENE	29,285	227	189

three dimensions and the reconstruction of its boundary with COCONESHape after the manifold extraction step.

The sample points on the boundaries of the surfaces have dimension one. While computing candidate triangles COCONESHape chooses triangles incident to these boundary samples. They are selected by their neighbors whose dimension is correctly assigned as two. Figure 6 shows that the boundary samples above the ankle of the foot have been connected to the rest of the foot correctly. We conclude our experiments with a data extracted from manifolds of different dimensions. The result is shown in Fig. 7.

We summarize our experimental data on dimension detection and reconstruction in Table 1.

7. Future Research

We need to test the algorithm DIMENSION with data sets in higher dimensions. Unfortunately, real data in higher dimensions are not as easily available as in the three dimensions. Also, we need to compute Voronoi diagrams in higher dimensions that are robust against numerical errors. Currently we are in the process of these developments.

The theory developed in this paper applies to manifolds without boundaries. Although in our experiments in three dimensions we used heuristics to detect the boundaries, an algorithm based on sound theory needs to be developed. Feature recognitions, particularly in higher dimensions, can benefit from such complete information. We believe that the structure of the Voronoi cells can again guide us to decipher this information. Currently research is in progress along this direction.

Acknowledgments

The first author thanks Marshall Bern for suggesting the problem of dimension detection from sample points and for his insightful comments. We also thank Rephael Wenger for his comments on the first draft.

References

1. N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, **22** (1999), 481–504.
2. N. Amenta, M. Bern and D. Eppstein. The crust and the β -skeleton: combinatorial curve reconstruction. *Graph. Models Image Process.*, **60** (1998), 125–135.
3. N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Appl.*, **12** (2002), 125–141.
4. C. L. Bajaj, V. Pascucci and D. R. Schikore. Visualization of scalar topology for structural enhancement. *Proc. IEEE Visualization 98*, 1998, pp. 18–23.
5. C. Bregler and S. M. Omohundro. Nonlinear manifold learning for visual speech recognition. *Proc. 5th Internat. Conf. Comput. Vision*, 1995, pp. 494–499.
6. J. D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbor interpolation of distance functions. *Proc. 16th Ann. Sympos. Comput. Geom.*, 2000, pp. 223–232.
7. M. P. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, NJ, 1976.
8. T. F. Cox and M. A. A. Cox. *Multidimensional Scaling*. Chapman and Hall, London, 1994.
9. T. K. Dey, S. Funke and E. A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. *Proc. European Workshop on Comput. Geom.*, Berlin, March 2001.
10. T. K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. *Proc. 17th Ann. Sympos. Comput. Geom.*, 2001, pp. 257–263.
11. T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. *Proc. 12th Ann. ACM–SIAM Sympos. Discrete Algorithms*, 1999, pp. 893–894.
12. H. Edelsbrunner, D. G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory*, **29** (1983), 551–559.
13. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics*, **13** (1994), 43–72.
14. H. Edelsbrunner, M. A. Facello and J. Liang. On the definition and the construction of pockets in macromolecules. *Discrete Appl. Math.*, **88** (1998), 83–102.
15. G. A. Edgar. *Measure, Topology and Fractal Geometry*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1990.
16. J. Erickson. Nice point sets can have nasty Delaunay triangulations. *Proc. 17th Ann. Sympos. Comput. Geom.*, 2001, pp. 96–105.
17. K. J. Falconer. *The Geometry of Fractal Sets*. Cambridge University Press, Cambridge, 1985.
18. J. Giesen. Curve reconstruction, the TSP, and Menger’s theorem on length. *Proc. 15th Ann. Sympos. Comput. Geom.*, 1999, pp. 207–216.
19. P. Indyk. Algorithmic applications of low-distortion geometric embeddings. *Proc. 42nd Ann. IEEE Sympos. Foundations Comput. Sci.*, 2001, pp. 10–33.
20. B. B. Mandelbrot. *The Fractal Geometry of Nature*. Freeman, New York, 1983.
21. T. Martinetz and K. Schulten. Topology preserving networks. *Neural Networks*, **7** (1994), 507–522.
22. D. Silver and X. Wang. Tracking scalar features in unstructured datasets. *Proc. IEEE Visualization 98*, 1998, pp. 79–86.
23. C-K. Tang and G. Medioni. Extremal feature extraction from 3-D vector and noisy scalar fields. *Proc. IEEE Visualization 98*, 1998, pp. 95–102.
24. J. B. Tenenbaum, V. de Silva and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, **290** (2000), 5500.
25. K. M.-K. Yip. *KAM: A System for Intelligently Guiding Numerical Experimentation by Computer*. The MIT Press, Cambridge, MA, 1991.
26. <http://www.cis.ohio-state.edu/~tamaldehy/cocone.html>.
27. <http://www.cgal.org>.

Received December 18, 2001, and in revised form July 25, 2002. Online publication January 30, 2003.