1

Basics

Simply stated, the problem we study in this book is: how to approximate a shape from the coordinates of a given set of points from the shape. The set of points is called a point sample, or simply a *sample* of the shape. The specific shape that we will deal with are curves in two dimensions and surfaces in three dimensions. The problem is motivated by the availability of modern scanning devices that can generate a point sample from the surface of a geometric object. For example, a range scanner can provide the depth values of the sampled points on a surface from which the three-dimensional coordinates can be extracted. Advanced hand held laser scanners can scan a machine or a body part to provide a dense sample of the surfaces. A number of applications in computer-aided design, medical imaging, geographic data processing, and drug designs, to name a few, can take advantage of the scanning technology to produce samples and then compute a digital model of a geometric shape with reconstruction algorithms. Figure 1.1 shows such an example for a sample on a surface which is approximated by a triangulated surface interpolating the input points.

The reconstruction algorithms described in this book produce a piecewise linear approximation of the sampled curves and surfaces. By approximation we mean that the output captures the topology and geometry of the sampled shape. This requires some concepts from topology which are covered in Section 1.1.

Clearly, a curve or a surface cannot be approximated from a sample unless it is dense enough to capture the features of the shape. The notions of features and dense sampling are formalized in Section 1.2.

All reconstruction algorithms described in this book use the data structures called *Voronoi diagrams* and their duals called *Delaunay triangulations*. The key properties of these data structures are described in Section 1.3.

Cambridge University Press
978-0-521-86370-4 - Curve and Surface Reconstruction: Algorithms with Mathematical Analysis
Tamal K. Dey
Excerpt
More information

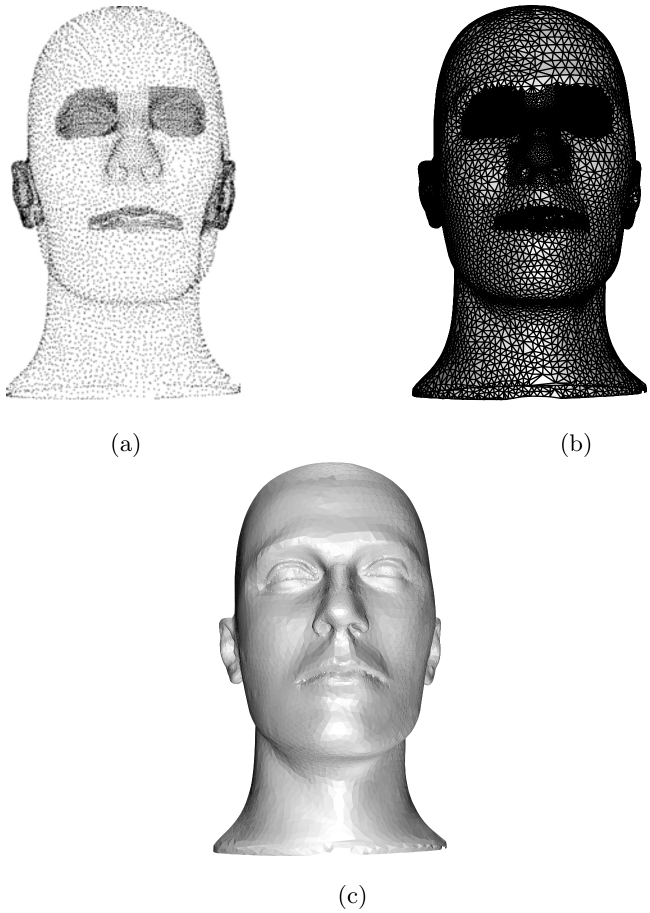

Figure 1.1. (a) A sample of MANNEQUIN, (b) a reconstruction, and (c) rendered MANNEQUIN model.

1.1 Shapes

The term *shape* can circumscribe a wide variety of meaning depending on the context. We define a shape to be a subset of an Euclidean space. Even this class is too broad for our purpose. So, we focus on a specific type of shapes called *smooth manifolds* and limit ourselves only up to three dimensions.

A global yardstick measuring similarities and differences in shapes is provided by *topology*. It deals with the connectivity of spaces. Various shapes are compared with respect to their connectivities by comparing functions over them called *maps*.

1.1.1 Spaces and Maps

In point set topology a *topological space* is defined to a be a point set $\mathbb{T}$ with a system of subsets $\mathcal{T}$ so that the following conditions hold.

1. $\emptyset, \mathbb{T} \in \mathcal{T}$,
2. $U \subseteq \mathcal{T}$ implies that the union of U is in $\mathcal{T}$,
3. $U \subseteq \mathcal{T}$ and U finite implies that the intersection of U is in $\mathcal{T}$.

The system $\mathcal{T}$ is the topology on the set $\mathbb{T}$ and its sets are *open* in $\mathbb{T}$. The *closed* sets of $\mathbb{T}$ are the subsets whose complements are open in $\mathbb{T}$. Consider the k-dimensional Euclidean space $\mathbb{R}^k$ and let us examine a topology on it. An *open ball* is the set of all points closer than a certain Euclidean distance to a given point. Define $\mathcal{T}$ as the set of open sets that are a union of a set of open balls. This system defines a topology on $\mathbb{R}^k$.

A subset $\mathbb{T}' \subseteq \mathbb{T}$ with a *subspace topology* $\mathcal{T}'$ defines a *topological subspace* where $\mathcal{T}'$ consists of all intersections between $\mathbb{T}'$ and the open sets in the topology $\mathcal{T}$ of $\mathbb{T}$. Topological spaces that we will consider are subsets of $\mathbb{R}^k$ which inherits their topology as a subspace topology. Let x denote a point in $\mathbb{R}^k$, that is, $x = \{x_1, x_2, \ldots, x_k\}$ and $\|x\| = (x_1^2 + x_2^2 + \cdots + x_k^2)^{\frac{1}{2}}$ denote its distance from the origin. Example of subspace topology are the k-ball $\mathbb{B}^k$, k-sphere $\mathbb{S}^k$, the halfspace $\mathbb{H}^k$, and the open k-ball $\mathbb{B}^k_o$ where

$$\begin{aligned}
\mathbb{B}^k &= \{x \in \mathbb{R}^k \mid \|x\| \leq 1\}\\
\mathbb{S}^k &= \{x \in \mathbb{R}^{k+1} \mid \|x\| = 1\}\\
\mathbb{H}^k &= \{x \in \mathbb{R}^k \mid x_k \geq 0\}\\
\mathbb{B}^k_o &= \mathbb{B}^k \setminus \mathbb{S}^k.
\end{aligned}$$

It is often important to distinguish topological spaces that can be covered with finitely many open balls. A *covering* of a topological space $\mathbb{T}$ is a collection of open sets whose union is $\mathbb{T}$. The topological space $\mathbb{T}$ is called *compact* if every covering of $\mathbb{T}$ can be covered with finitely many open sets included in the covering. An example of a compact topological space is the k-ball $\mathbb{B}^k$. However, the open k-ball is not compact. The *closure* of a topological space $\mathbb{X} \subseteq \mathbb{T}$ is the smallest closed set $\mathrm{Cl}\mathbb{X}$ containing $\mathbb{X}$.

Continuous functions between topological spaces play a significant role to define their similarities. A function $g : \mathbb{T}_1 \to \mathbb{T}_2$ from a topological space $\mathbb{T}_1$ to a topological space $\mathbb{T}_2$ is *continuous* if for every open set $U \subseteq \mathbb{T}_2$, the set $g^{-1}(U)$ is open in $\mathbb{T}_1$. Continuous functions are called *maps*.

Homeomorphism

Broadly speaking, two topological spaces are considered the same if one has a correspondence to the other which keeps the connectivity same. For example, the surface of a cube can be deformed into a sphere without any incision or attachment during the process. They have the same topology. A precise definition for this topological equality is given by a map called *homeomorphism*. A homeomorphism between two topological spaces is a map $h : \mathbb{T}_1 \to \mathbb{T}_2$ which is bijective and has a continuous inverse. The explicit requirement of continuous inverse can be dropped if both $\mathbb{T}_1$ and $\mathbb{T}_2$ are compact. This is because any bijective map between two compact spaces must have a continuous inverse. This fact helps us proving homeomorphisms for spaces considered in this book which are mostly compact.

Two topological spaces are *homeomorphic* if there exists a homeomorphism between them. Homeomorphism defines an equivalence relation among topological spaces. That is why two homeomorphic topological spaces are also called *topologically equivalent*. For example, the open k-ball is topologically equivalent to $\mathbb{R}^k$. Figure 1.2 shows some more topological spaces some of which are homeomorphic.

Homotopy

There is another notion of similarity among topological spaces which is weaker than homeomorphism. Intuitively, it relates spaces that can be continuously deformed to one another but may not be homeomorphic. A map $g : \mathbb{T}_1 \to \mathbb{T}_2$ is *homotopic* to another map $h : \mathbb{T}_1 \to \mathbb{T}_2$ if there is a map $H : \mathbb{T}_1 \times [0, 1] \to \mathbb{T}_2$ so that $H(x, 0) = g(x)$ and $H(x, 1) = h(x)$. The map H is called a *homotopy* between g and h.

Two spaces $\mathbb{T}_1$ and $\mathbb{T}_2$ are *homotopy equivalent* if there exist maps $g : \mathbb{T}_1 \to \mathbb{T}_2$ and $h : \mathbb{T}_2 \to \mathbb{T}_1$ so that $h \circ g$ is homotopic to the identity map $\iota_1 : \mathbb{T}_1 \to \mathbb{T}_1$ and $g \circ h$ is homotopic to the identity map $\iota_2 : \mathbb{T}_2 \to \mathbb{T}_2$. If $\mathbb{T}_2 \subset \mathbb{T}_1$, then $\mathbb{T}_2$ is a *deformation retract* of $\mathbb{T}_1$ if there is a map $r : \mathbb{T}_1 \to \mathbb{T}_2$ which is homotopic to ι_1 by a homotopy that fixes points of $\mathbb{T}_2$. In this case $\mathbb{T}_1$ and $\mathbb{T}_2$ are homotopy equivalent. Notice that homotopy relates two maps while homotopy equivalence relates two spaces. A curve and a point are not homotopy equivalent. However, one can define maps from a 1-sphere $\mathbb{S}^1$ to a curve and a point in the plane which have a homotopy.

One difference between homeomorphism and homotopy is that homeomorphic spaces have same dimension while homotopy equivalent spaces need not have same dimension. For example, the 2-ball shown in Figure 1.2(e) is homotopy equivalent to a single point though they are not homeomorphic. Any of

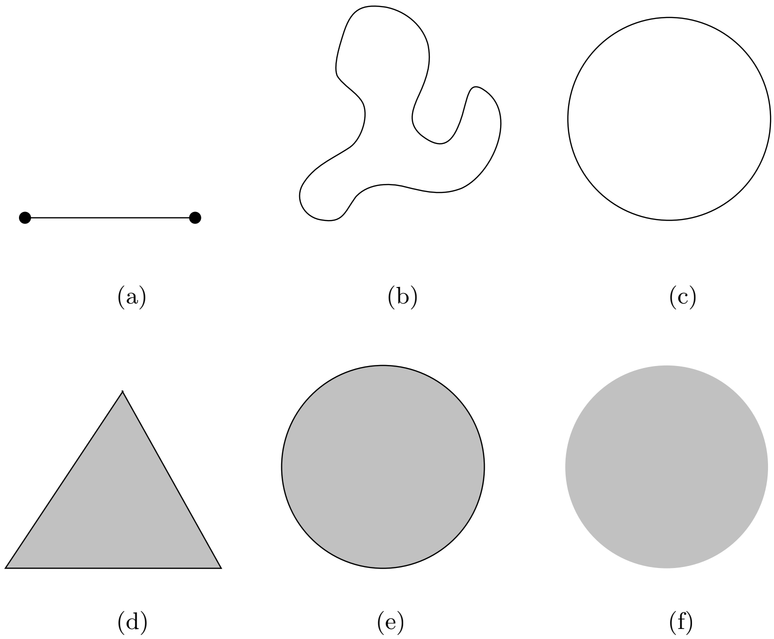

Figure 1.2. (a) 1-ball, (b) and (c) spaces homeomorphic to the 1-sphere, (d) and (e) spaces homeomorphic to the 2-ball, and (f) an open 2-ball which is not homeomorphic to the 2-ball in (e).

the end vertices of the segment in Figure 1.2(a) is a deformation retract of the segment.

Isotopy

Homeomorphism and homotopy together bring a notion of similarity in spaces which, in some sense, is stronger than each one of them individually. For example, consider a standard torus embedded in $\mathbb{R}^3$. One can knot the torus (like a knotted rope) which still embeds in $\mathbb{R}^3$. The standard torus and the knotted one are both homeomorphic. However, there is no continuous deformation of one to the other while maintaining the property of embedding. The reason is that the complement spaces of the two tori are not homotopy equivalent. This requires the notion of *isotopy*.

An *isotopy* between two spaces $\mathbb{T}_1 \subseteq \mathbb{R}^k$ and $\mathbb{T}_2 \subseteq \mathbb{R}^k$ is a map $\xi : \mathbb{T}_1 \times [0, 1] \to \mathbb{R}^k$ such that $\xi(\mathbb{T}_1, 0)$ is the identity of $\mathbb{T}_1$, $\xi(\mathbb{T}_1, 1) = \mathbb{T}_2$ and for each $t \in [0, 1]$, $\xi(\cdot, t)$ is a homeomorphism onto its image. An *ambient isotopy* between $\mathbb{T}_1$ and $\mathbb{T}_2$ is a map $\xi : \mathbb{R}^k \times [0, 1] \to \mathbb{R}^k$ such that $\xi(\cdot, 0)$ is the identity of $\mathbb{R}^k$, $\xi(\mathbb{T}_1, 1) = \mathbb{T}_2$ and for each $t \in [0, 1]$, $\xi(\cdot, t)$ is a homeomorphism of $\mathbb{R}^k$.

Observe that ambient isotopy also implies isotopy. It is also known that two spaces that have an isotopy between them also have an ambient isotopy between them. So, these two notions are equivalent. We will call $\mathbb{T}_1$ and $\mathbb{T}_2$ *isotopic* if they have an isotopy between them.

When we talk about reconstructing surfaces from sample points, we would like to claim that the reconstructed surface is not only homeomorphic to the sampled one but is also isotopic to it.

1.1.2 Manifolds

Curves and surfaces are a particular type of topological space called *manifolds*. A *neighborhood* of a point $x \in \mathbb{T}$ is an open set that contains x. A topological space is a *k-manifold* if each of its points has a neighborhood homeomorphic to the open k-ball which in turn is homeomorphic to $\mathbb{R}^k$. We will consider only k-manifolds that are subspaces of some Euclidean space.

The plane is a 2-manifold though not compact. Another example of a 2-manifold is the sphere $\mathbb{S}^2$ which is compact. Other compact 2-manifolds include *torus* with one through-hole and *double torus* with two through-holes. One can glue g tori together, called *summing* g tori, to form a 2-manifold with g through-holes. The number of through-holes in a 2-manifold is called its *genus*. A remarkable result in topology is that all compact 2-manifolds in $\mathbb{R}^3$ must be either a sphere or a sum of g tori for some $g \geq 1$.

Boundary

Surfaces in $\mathbb{R}^3$ as we know them often have boundaries. These surfaces have the property that each point has a neighborhood homeomorphic to $\mathbb{R}^2$ except the ones on the boundary. These surfaces are 2-manifolds with boundary. In general, a *k-manifold with boundary* has points with neighborhood homeomorphic to either $\mathbb{R}^k$, called the *interior points*, or the halfspace $\mathbb{H}^k$, called the *boundary points*. The boundary of a manifold M, bd M, consists of all boundary points. By this definition a manifold as defined earlier has a boundary, namely an empty one. The interior of M consists of all interior points and is denoted Int M.

It is a nice property of manifolds that if M is a k-manifold with boundary, bd M is a $(k-1)$-manifold unless it is empty. The k-ball $\mathbb{B}^k$ is an example of a k-manifold with boundary where bd $\mathbb{B}^k = \mathbb{S}^{k-1}$ is the $(k-1)$-sphere and its interior Int $\mathbb{B}^k$ is the the open k-ball. On the other hand, bd $\mathbb{S}^k = \emptyset$ and Int $\mathbb{S}^k = \mathbb{S}^k$. In Figure 1.2(a), the segment is a 1-ball where the boundary is a 0-sphere consisting of the two endpoints. In Figure 1.2(e), the 2-ball is a manifold with boundary and its boundary is the circle, a 1-sphere.

Orientability

A 2-manifold with or without boundary can be either *orientable* or *nonorientable*. We will only give an intuitive explanation of this notion. If one travels along any curve on a 2-manifold starting at a point, say p, and considers the oriented normals at each point along the curve, then one gets the same oriented normal at p when he returns to p. All 2-manifolds in $\mathbb{R}^3$ are orientable. However, 2-manifolds in $\mathbb{R}^3$ that have boundaries may not be orientable. For example, the Möbius strip, obtained by gluing the opposite edges of a rectangle with a twist, is nonorientable. The 2-manifolds embedded in four and higher dimensions may not be orientable no matter whether they have boundaries or not.

1.1.3 Complexes

Because of finite storage within a computer, a shape is often approximated with finitely many simple pieces such as vertices, edges, triangles, and tetrahedra. It is convenient and sometimes necessary to borrow the definitions and concepts from combinatorial topology for this representation.

An *affine combination* of a set of points $P = \{p_0, p_1, \ldots, p_n\} \subset \mathbb{R}^k$ is a point $p \in \mathbb{R}^k$ where $p = \Sigma_{i=0}^n \alpha_i p_i$, $\Sigma_i \alpha_i = 1$ and each α_i is a real number. In addition, if each α_i is nonnegative, the point p is a *convex combination*. The *affine hull* of P is the set of points that are an affine combination of P. The *convex hull* Conv P is the set of points that are a convex combination of P. For example, three noncollinear points in the plane have the entire $\mathbb{R}^2$ as the affine hull and the triangle with the three points as vertices as the convex hull.

A set of points is affinely independent if none of them is an affine combination of the others. A *k-polytope* is the convex hull of a set of points which has at least $k + 1$ affinely independent points. The affine hull aff μ of a polytope μ is the affine hull of its vertices.

A *k-simplex* σ is the convex hull of exactly $k + 1$ affinely independent points P. Thus, a vertex is a 0-simplex, an edge is a 1-simplex, a triangle is a 2-simplex, and a tetrahedron is a 3-simplex. A simplex $\sigma' =$ Conv T for a nonempty subset $T \subseteq P$ is called a *face* of σ. Conversely, σ is called a *coface* of σ'. A face $\sigma' \subset \sigma$ is *proper* if the vertices of σ' are a proper subset of σ. In this case σ is a *proper* coface of σ'.

A collection $\mathcal{K}$ of simplices is called a *simplicial complex* if the following conditions hold.

(i) $\sigma' \in \mathcal{K}$ if σ' is a face of any simplex $\sigma \in \mathcal{K}$.
(ii) For any two simplices $\sigma, \sigma' \in \mathcal{K}$, $\sigma \cap \sigma'$ is a face of both unless it is empty.

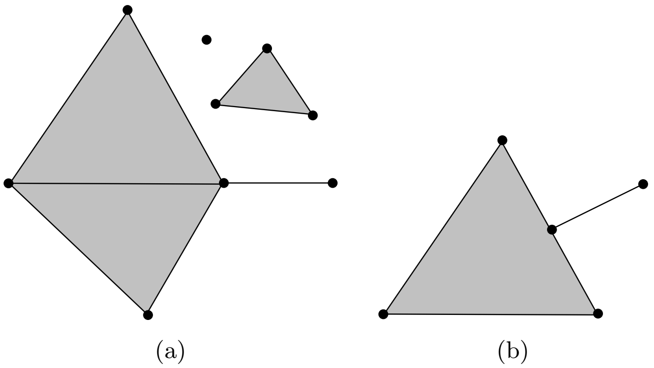

Figure 1.3. (a) A simplicial complex and (b) not a simplicial complex.

The above two conditions imply that the simplices meet nicely. The simplices in Figure 1.3(a) form a simplicial complex whereas the ones in Figure 1.3(b) do not.

Triangulation

A triangulation of a topological space $\mathbb{T}$ is a simplicial complex $\mathcal{K}$ whose underlying point set is $\mathbb{T}$. Figure 1.1(b) shows a triangulation of a 2-manifold with boundary.

Cell Complex

We also use a generalized version of simplicial complexes in this book. The definition of a cell complex is exactly same as that of the simplicial complex with simplices replaced by polytopes. A cell complex is a collection of polytopes and their faces where any two intersecting polytopes meet in a face which is also in the collection. A cell complex is a k-complex if the largest dimension of any polytope in the complex is k. We also say that two elements in a cell complex are *incident* if they intersect.

1.2 Feature Size and Sampling

We will mainly concentrate on smooth curves in $\mathbb{R}^2$ and smooth surfaces in $\mathbb{R}^3$ as the sampled spaces. The notation Σ will be used to denote this generic sampled space throughout this book. We will defer the definition of smoothness until Chapter 2 for curves and Chapter 3 for surfaces. It is sufficient to assume

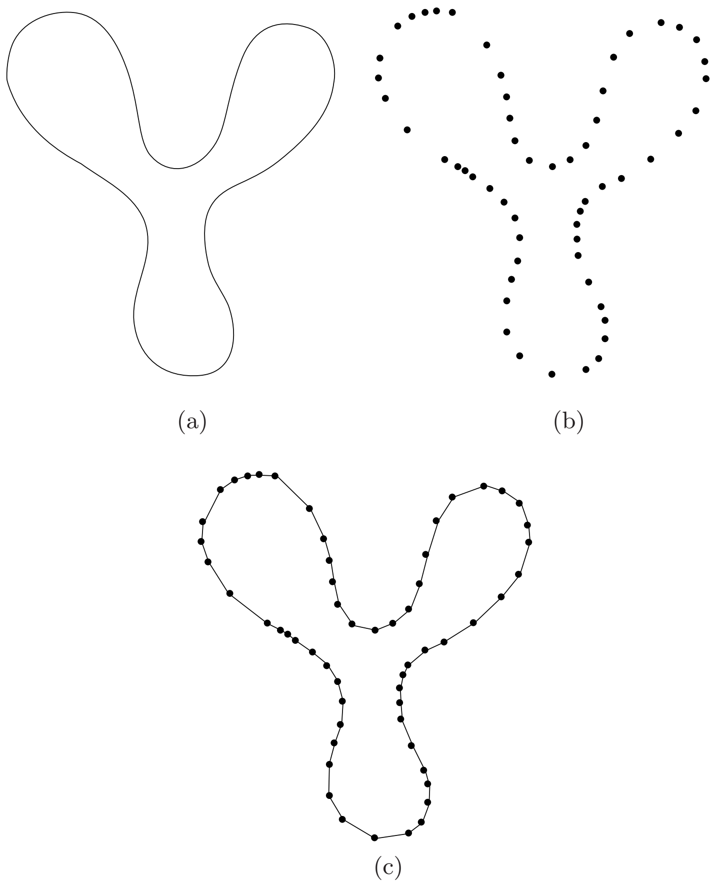

Figure 1.4. (a) A curve in the plane, (b) a sample of it, and (c) the reconstructed curve.

that Σ is a 1-manifold in $\mathbb{R}^2$ and a 2-manifold in $\mathbb{R}^3$ for the definitions and results described in this chapter.

Obviously, it is not possible to extract any meaningful information about Σ if it is not sufficiently sampled. This means features of Σ should be represented with sufficiently many sample points. Figure 1.4 shows a curve in the plane which is reconstructed from a sufficiently dense sample. But, this brings up the question of defining features. We aim for a measure that can tell us how complicated Σ is around each point $x \in \Sigma$. A geometric structure called the *medial axis* turns out to be useful to define such a measure.

Before we define the medial axis, let us fix some notations about distances and balls that will be used throughout the rest of this book. The Euclidean distance between two points $p = (p_1, p_2, \ldots, p_k)$ and $x = (x_1, x_2, \ldots, x_k)$ in $\mathbb{R}^k$ is the length $\|p - x\|$ of the vector $\overrightarrow{xp} = (p - x)$ where

$$\|p - x\| = \left\{(p_1 - x_1)^2 + (p_2 - x_2)^2 + \cdots + (p_k - x_k)^2\right\}^{\frac{1}{2}}.$$

Also, we have

$$\begin{aligned}\|p - x\| &= \{(p - x)^T (p - x)\}^{\frac{1}{2}} \\ &= \{p^T p - 2p^T x + x^T x\}^{\frac{1}{2}} \\ &= \{\|p\|^2 - 2p^T x + \|x\|^2\}^{\frac{1}{2}}.\end{aligned}$$

For a set $P \subseteq \mathbb{R}^k$ and a point $x \in \mathbb{R}^k$, let $d(x, P)$ denote the Euclidean distance of x from P; that is,

$$d(x, P) = \inf_{p \in P} \{\|p - x\|\}.$$

We will also consider distances called *Hausdorff distances* between two sets. For $X, Y \subseteq \mathbb{R}^k$ this distance is given by

$$\max\{\sup_{x \in X} d(x, Y), \sup_{y \in Y} d(y, X)\}.$$

Roughly speaking, the Hausdorff distance tells how much one set needs to be moved to be identical with the other set.

The set $B_{x,r} = \{y \mid y \in \mathbb{R}^k, \|y - x\| \leq r\}$ is a *ball* with center x and radius r. By definition $B_{x,r}$ and its boundary are homeomorphic to $\mathbb{B}^k$ and $\mathbb{S}^{k-1}$ respectively.

1.2.1 Medial Axis

The medial axis of a curve or a surface Σ is meant to capture the middle of the shape bounded by Σ. There are slightly different definitions of the medial axis in the literature. We adopt one of them and mention the differences with the others.

Assume that Σ is embedded in $\mathbb{R}^k$. A ball $B \subset \mathbb{R}^k$ is empty if the interior of B is empty of points from Σ. A ball B is maximal if every empty ball that contains B equals B. The *skeleton* Sk_Σ of Σ is the set of centers of all maximal balls. Let M_Σ^o be the set of points in $\mathbb{R}^k$ whose distance to Σ is realized by at least two points in Σ. The closure of M_Σ^o is M_Σ, that is, $M_\Sigma = \mathrm{Cl}\, M_\Sigma^o$. The following inclusions hold:

$$M_\Sigma^o \subseteq Sk_\Sigma \subseteq M_\Sigma.$$