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 space. 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 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 sufficently 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 *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 $p = (p_1, p_2, ..., p_k)$ and $x = (x_1, x_2, .., x_k)$ in $\mathbb{R}^k$ is

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

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,

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

The set $B_{x,r} = \{y \mid y \in \mathbb{R}^k, \|y - x\| \leq r\}$ is a *ball* with radius r and center x. 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* M of Σ is the set of centers of all maximal balls whose interiors are empty of the points from Σ. In other words, each point of M is the center of a ball that meets Σ only tangentially. We call each ball $B_{x,r}$, $x \in M$, a *medial ball* where $r = d(x, \Sigma)$.

Figure 1.5 (a) shows a subset of the medial axis of a curve. Notice that the medial axis may have a branching point such as v and boundary points such as u and w. Also, the medial axis need not be connected. For example, the part of the medial axis in the region bounded by the curve may be disjoint from the rest, see Figure 1.5 (a). In fact, if Σ is smooth, the two parts of

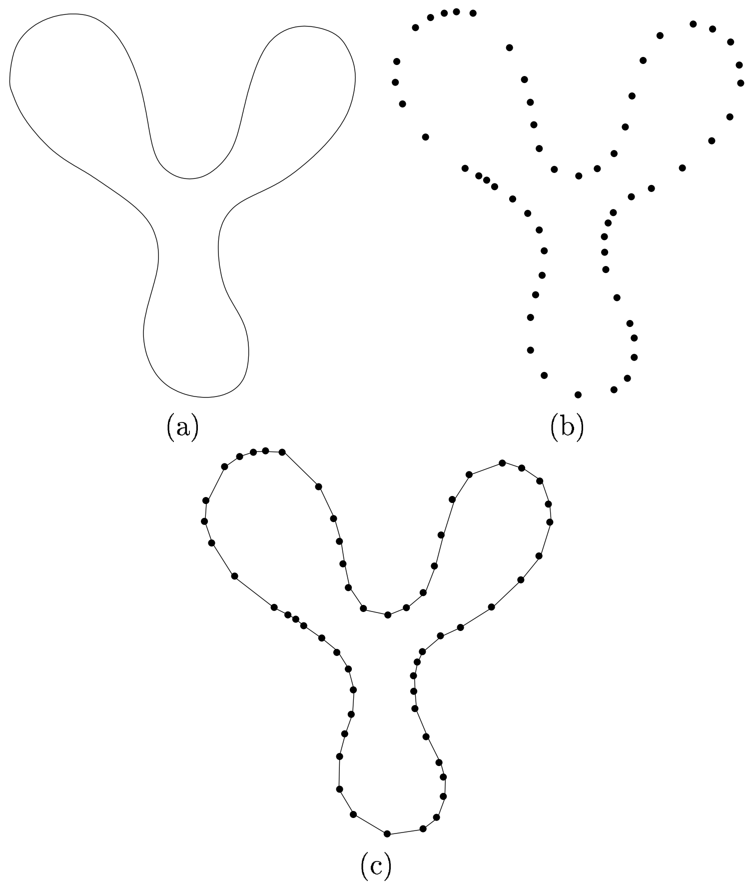

Figure 1.4: (a) A curve in the plane, (b) a sample of it, (c) reconstructed curve.

the medial axis are indeed disjoint. The subset of the medial axis residing in the unbounded component of $\mathbb{R}^2 \setminus \Sigma$ is called the *outer* medial axis. The rest is called the *inner* medial axis.

It follows from the definition that, if one grows a ball around a point on the medial axis, it will either meet Σ for the first time tangentially in two or more points or remains tangent to Σ at a single point, see Figure 1.5 (b). Conversely, for a point $x \in \Sigma$ one can start growing a ball keeping it tangent to Σ at x until it hits another point $y \in \Sigma$ or becomes maximally empty. In the latter case we can think $x = y$. At this moment the ball is medial and the segments joining the center m to x and y are normal to Σ at x and y

respectively, see Figure 1.5.

If we move along the medial axis and consider the medial balls as we move, the radius of the medial balls increases or decreases accordingly to maintain the tangency with Σ, and at the boundaries coincides with the radius of the *curvature ball* where all tangent points merge into a single one. See Figure 1.5 (b).

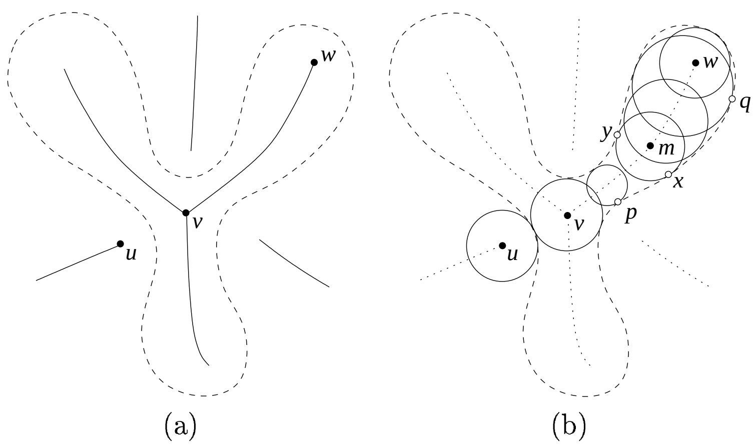

Figure 1.5: (a) A subset of the medial axis of the curve in Figure 1.4, (b) medial ball centered at v touches the curve in three points, whereas the ones with centers u and w touch it in only one point and coincide with the curvature ball.

It will be useful for our proofs later to know the following property of balls intersecting the sampled space Σ. The proof of the lemma assumes that Σ is either a smooth curve or a smooth surface whose definitions are given in later chapters. Also, the proof uses some concepts from differential topology (critical point theory) which are exposed in chapter 8. The readers may skip the proof at this point if they are not familiar with these concepts.

We say that a topological space X is a k-ball or a k-sphere if X is homeomorphic to $\mathbb{B}^k$ or $\mathbb{S}^k$ respectively.

Lemma 1.1 (Feature Ball.) *If a d-ball $B = B_{c,r}$ intersects a k-manifold $\Sigma \subset \mathbb{R}^d$ at more than one point with either (a) $B \cap \Sigma$ is not a k-ball, or (b) $\mathrm{bd}\,(B \cap \Sigma)$ is not a $(k-1)$-sphere, then a medial axis point is in B.*

PROOF. First we show that if B intersects Σ at more than one point and B is tangent to Σ at a point, B contains a medial axis point. Let x be the point of this tangency. Shrink B further keeping it tangent to Σ at x. This means the center of B moves towards x along a normal direction at x. We stop when B meets Σ only tangentially. Observe that, since $B \cap \Sigma \neq x$ to start with, this happens eventually when B is maximally empty. At this moment B becomes a medial ball and its center is a medial axis point which must lie in the original ball B, refer to Figure 1.6.

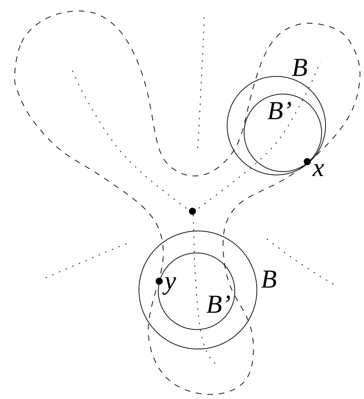

Figure 1.6: (a) The ball B intersecting the upper right lobe is shrunk till it becomes tangent to another point other than x. The new ball B' intersects the medial axis. (b) The ball B intersecting the lower lobe is shrunk radially to the ball B' that is tangent to the curve at y and also intersects the curve in other points. B' can further be shrunk till it meets the curve only tangentially.

Now consider that the condition (b) holds. Define a function $h\colon B \cap \Sigma \to \mathbb{R}$ where $h(x)$ is the distance of x from the center c of B. Let m be a point in Σ so that $h(m)$ is a global minimum. If there is more than one such global minimum, the ball B becomes tangent to Σ at more than one point when radially shrunk to a radius of $h(m)$. Then, we can apply the previous argument to claim that B contains a medial axis point. So, assume that there is only global minimum m of h.

We claim that the function h has a critical point p in $\mathrm{Int}\,(B \cap \Sigma)$ other than m where B becomes tangent to Σ. If not, as we shrink B centrally the level set $\mathrm{bd}\,(B \cap \Sigma)$ does not change topology until it reaches the minimum m when it vanishes. Since m is a minimum, there is a small enough $\delta > 0$ so that $B_{c,h(m)+\delta} \cap \Sigma$ is a k-ball. The boundary of this k-ball given by $(\mathrm{bd}\, B_{c,h(m)+\delta}) \cap \Sigma$ should be a $(k-1)$-sphere. This contradicts the fact that $\mathrm{bd}\,(B \cap \Sigma)$ is not a $(k-1)$-sphere and remains that way till the end.

Therefore, there is a critical point, say $p \neq m$ of h in $\text{Int}\,(B \cap \Sigma)$. At this point p, the ball $B_{c,||p-c||}$ becomes tangent to Σ, see also Figure 1.6. Now we can apply our previous argument to claim that B contains a medial axis point.

Next, consider that the condition (a) holds. If the condition (b) also holds, we have the previous argument. So, assume that $\text{bd}\,(B \cap \Sigma)$ is a $(k-1)$-sphere and $B \cap \Sigma$ is not a k-ball. Again, we claim that the function h as defined earlier has a critical point other than m. If not, consider the subset of Σ swept by B while shrinking it till it meets Σ only at m. This subset is homeomorphic to the quotient space $(\mathbb{S}^{k-1} \times I)/\mathbb{S}^{k-1}$ where I is the closed unit interval in $\mathbb{R}$. This quotient space is a k-ball which contradicts the fact that $B \cap \Sigma$ is not a k-ball to begin with. Therefore, as B is continually shrunk, it becomes tangent to Σ at a point $p \neq m$. Apply the previous argument to claim that B has a medial axis point.

□

Figure 1.7 illustrates the different cases of Feature Ball Lemma.

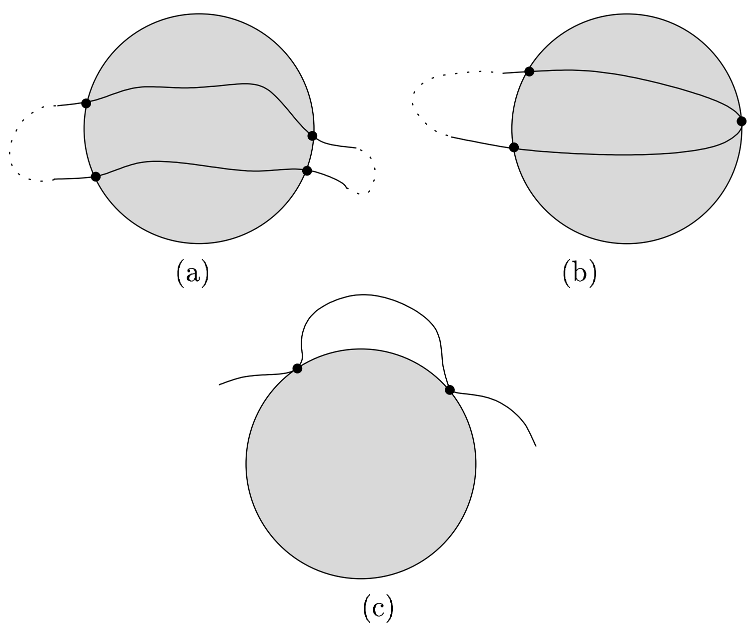

Figure 1.7: (a) $B \cap \Sigma$ is not a 1-ball, (b) $B \cap \Sigma$ is a 1-ball, but $\text{bd}\, B \cap \Sigma$ is not a 0-sphere, (c) $\text{bd}\, B \cap \Sigma$ is a 0-sphere, but $B \cap \Sigma$ is not a 1-ball.

1.2.2 Local feature size

The medial axis M with the distance r_m to Σ at each point $m \in M$ captures the shape of Σ. In fact, Σ is the boundary of the union of all medial balls centering points of the inner (or outer) medial axis. So, as a first attempt to capture local feature sizes we define the following two functions on Σ.

$\rho_i, \rho_o : \Sigma \rightarrow \mathbb{R}$ where $\rho_i(x)$, $\rho_o(x)$ are the radii of the inner and outer medial balls respectively both of which are tangent to Σ at x.

The functions ρ_i and ρ_o are continuous. However, we need a stronger form of continuity on the local feature size function to carry out the proofs. This property, called the *Lipschitz property*, stipulates that the difference in the function values at two points is bounded by a constant times the distance between the points. Keeping this in mind we define the *local feature size* as the function values of

$$f : \Sigma \rightarrow \mathbb{R} \text{ where } f(x) = d(x, M),$$

that is, $f(x)$ is the distance of $x \in \Sigma$ to the medial axis M. As can be observed from Figure 1.5 (b), the local feature size, $f(p)$ at p, is small compared to the local feature size, $f(q)$ at q, in accordance with our intuitive notion of features. It follows from the definitions that $f(x) \leq \min\{\rho_i(x), \rho_o(x)\}$. Lipschitz property of the local feature size function f follows easily from the definition.

Lemma 1.2 (Lipschitz Continuity.) *$f(x) \leq f(y) + \|x - y\|$ for any two points x and y in Σ.*

PROOF. Let m be a point on the medial axis so that $f(y) = \|y - m\|$. By triangular inequality,

$$\begin{aligned} \|x - m\| &\leq \|y - m\| + \|x - y\|, \text{ and} \\ f(x) \leq \|x - m\| &\leq f(y) + \|x - y\|. \end{aligned}$$

□

1.2.3 Sampling

A *sample* P of Σ is a set of points from Σ. Once we have quantized the feature size, we would require the sample respect the features, i.e., we require more sample points where the local feature size is small compared to the regions where it is not.

Definition 1.1 *A sample P of Σ is an ε-sample if each point $x \in \Sigma$ has a sample point $p \in P$ so that $\|x - p\| \leq \varepsilon f(x)$.*

The value of ε has to be smaller than 1 to have a dense sample. In fact, practical experiments suggest that $\varepsilon < 0.4$ constitutes a dense sample useful for reconstructing Σ from P. An ε-sample is also an ε'-sample for any $\varepsilon' > \varepsilon$. The definition of ε-sample allows a sample arbitrarily dense anywhere on Σ. It only puts a lower bound on the density. Figure 1.8 illustrates a sample of a circle which is a 0.2-sample. By definition, it is also a 0.3-sample of the same.

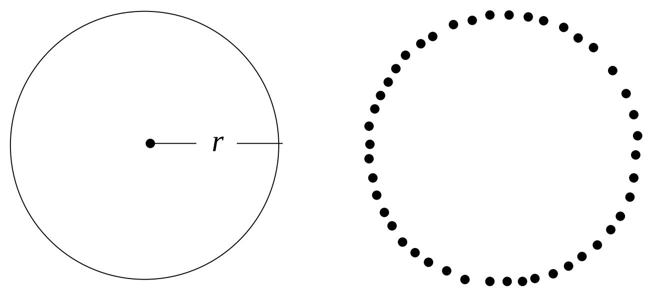

Figure 1.8: Local feature size at any point on the circle is equal to the radius r. Each point on the circle has a sample point within $0.2r$ distance.

An useful application of the Lipschitz Continuity Lemma (1.2) is that the distance between two points expressed in terms of the local feature size of one can be expressed in terms of that of the other.

Lemma 1.3 (Feature Translation.) *For any two points x, y in Σ with $\|x - y\| \leq \varepsilon f(x)$ we have*

(i) $f(x) \leq \frac{1}{1-\varepsilon} f(y)$, *and*

(ii) $\|x - y\| \leq \frac{\varepsilon}{1-\varepsilon} f(y)$.

PROOF. We have

$$\begin{aligned} f(x) &\leq f(y) + \|x - y\| \\ f(x) &\leq f(y) + \varepsilon f(x) \\ f(x) &\leq \frac{1}{1-\varepsilon} f(y) \text{ proving (i).} \end{aligned}$$

Plug the above inequality in $\|x - y\| \leq \varepsilon f(x)$ to obtain (ii). □

Uniform sampling. The definition of ε-sampling allows non-uniform sampling over Σ. A *globally uniform* sampling is more restrictive. It means that sampling is equally dense everywhere. Local feature size does not play a role in such samplings. There could be various definitions of globally uniform samples. We will say a sample $P \subset \Sigma$ is globally δ-uniform if any point $x \in \Sigma$ has a point in P within $\delta > 0$ distance. In between globally uniform and non-uniform samplings, there is another one called the *locally uniform sampling*. This sampling respects feature sizes and is uniform only locally. We say $P \subset \Sigma$ is *locally* (ε, δ)-*uniform* for $\delta > \varepsilon > 0$ if each point $x \in \Sigma$ has a point in P within $\varepsilon f(x)$ distance, and no point $p \in P$ has another point $q \in P$ within $\frac{\varepsilon}{\delta} f(p)$ distance.