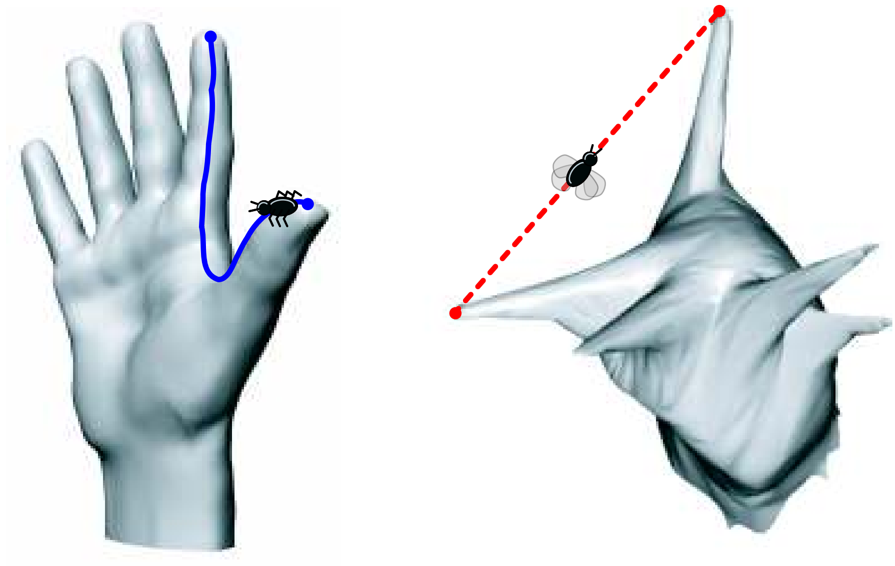

Surface Canonical form

Figure 7.1. Illustration of the isometric embedding problem. Left: original shape; right: canonical form.

7.1 Isometric embedding problem

Let us come back for a moment to the example we have already used in Chapter 2 to illustrate the concept of intrinsic geometry. Assume that our shape is inhabited by an insect, which always chooses the shortest path to crawl between any two points. Now, imagine that there is another shape embedded into $\mathbb{R}^m$, whose points correspond with those of the original shape, while the Euclidean distances between the points are equal to the original geodesic ones. On this new shape, there lives another winged insect, which flies along straight lines between the points.

The lengths of the paths are the isometry-invariant description of our object. Because the distances traveled by both insects are equal, the descriptions produced by them are the same. However, for our application, the second insect's point of view is preferable, as his world is Euclidean. The advantage stems from the smaller number of degrees of freedom that influence our description: whereas the first insect would not feel any isometric deformation of the shape, the only way we can fool the second one is by applying rigid transformations, which are limited to rotations, translations, or reflections.

Formalizing the above intuition, given a shape (X, d_X), we would like to find a map $f : (X, d_X) \to (\mathbb{R}^m, d_{\mathbb{R}^m})$, such that

$$d_X(x, x') = d_{\mathbb{R}^m}(f(x), f(x')),$$

for all $x, x' \in X$. Such an f is an isometric embedding, and the space $\mathbb{R}^m$ is referred to as the *embedding space* in this context. The image $f(X)$, which we call the *canonical form* of X, can be used as an extrinsic representation of the intrinsic geometry of X (see example in Figure 7.1). Note that we regard $f(X)$ as a metric space with the restricted Euclidean metric $d_{\mathbb{R}^m}|_{f(X)}$. Up to isometries in $\mathbb{R}^m$, it defines an equivalence class of all the shapes that are indistinguishable from the point of view of intrinsic geometry. In simple words, two isometric shapes have identical canonical forms, possibly differing by an isometry of $\mathbb{R}^m$.

In a sense, this embedding allows us to undo the non-uniqueness of the way the metric structure of X is realized in $\mathbb{R}^m$ (all the possible bendings), thus reducing its vast number of degrees of freedom. Consequently, considering the canonical forms instead of the shapes themselves, we translate the non-rigid shape similarity problem into a much simpler problem of rigid similarity, with which we already know how to deal. This simple idea, proposed by Asi Elad and R. K. [147, 149], allows us to define the similarity between two shapes as an extrinsic distance between their canonical forms, measured by means of ICP or the moments method as shown in Algorithm 7.1. We call this distance the *canonical form distance* and denote it by d_{CF}.

Though originally formulated with the particular choice of $\mathbb{R}^3$ (i.e., $m = 3$) as the embedding space, the canonical forms approach can be generalized to any embedding space, as we will see in Chapter 9. Here, we stick to the Euclidean embedding, but assume m to be arbitrary. Thus far, the canonical forms method seems an ideal recipe for our problem of non-rigid shape comparison. However, there is still a question whether a shape X is isometrically embeddable into $\mathbb{R}^m$.

Unfortunately, the answer is usually negative. As the simplest case, consider the problem of embedding a sphere into $\mathbb{R}^2$. This problem arose in cartography centuries ago. One of the fundamental problems in map-making is creating a planar map of the Earth, which reproduces, in the best way, the distances between geographic objects. That is, equipped with a simple ruler, we can measure distances on the map, which represent geodesic distances on the Earth (Figure 7.2). Every cartographer knows that it is impossible to create a map of the Earth that preserves all the geodesic distances.[1] This, as a matter of fact, is a consequence of the *theorema egregium*: because the Gaussian curvature of the sphere is positive, whereas the plane has zero curvature, these two surfaces cannot be isometric.

input : shapes (X, d_X) and (Y, d_Y).
output: canonical forms distance $d_{\mathrm{CF}}(X, Y)$.

1 Find the isometric embedding f and g of X and Y into $\mathbb{R}^m$.
2 Compute $d_{\mathrm{CF}}(X, Y)$ as $d_{\mathrm{MOM}}(f(X), g(Y))$ or $d_{\mathrm{ICP}}(f(X), g(Y))$.

Algorithm 7.1. Idealized canonical forms distance computation.

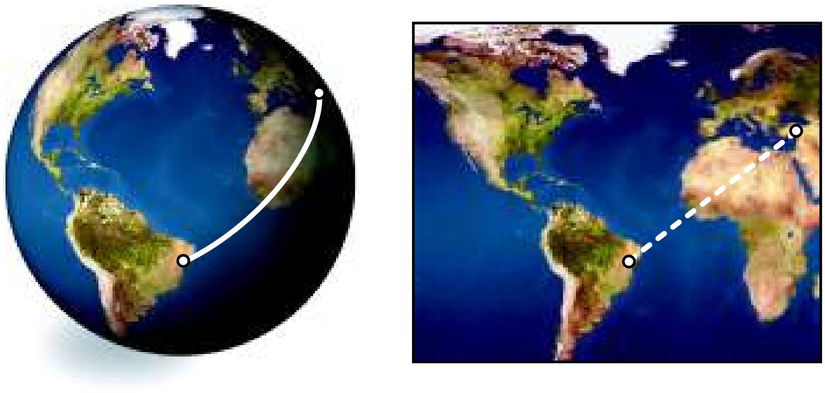

Figure 7.2. The problem of isometric embedding arising in cartography: the spherical surface of the Earth (shown is the upper hemisphere, left) is to be mapped into the plane so that it preserves the geodesic distances (right). A consequence of *theorema egregium* is that such a map does not exist, and a distortion of the distance is inevitable.

Yet, maybe by increasing the embedding space dimension, i.e., trying to embed the sphere into $\mathbb{R}^3, \mathbb{R}^4, \mathbb{R}^5$, and so on, we could succeed in finding an isometric embedding? Even this appears to be impossible. The following example, shown by Nathan Linial [253], demonstrates that even a very simple discrete metric space consisting of only four points cannot be isometrically embedded into a Euclidean space of any finite dimension.

Example 7.1 (Linial's example). Consider four points $x_1, \dots, x_4$, sampled on the sphere of radius $R = \frac{2}{\pi}$ as shown in Figure 7.3 (one point at the north pole and three points along the equator). The distances between the points are given by the following matrix

$$D_X = \begin{pmatrix} 0 & 1 & 2 & 1 \\ 1 & 0 & 1 & 1 \\ 2 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}.$$

We denote the embedded points by $z_1, \dots, z_4$ and assume that the embedding is distortionless, that is, $(D_X)_{ij} = d_{\mathbb{R}^m}(z_i, z_j)$ for $i, j = 1, \dots, 4$. Let us consider first the triangle with vertices z_1, z_2, z_3, with edges of lengths $d_{\mathbb{R}^m}(z_1, z_2) = d_{\mathbb{R}^m}(z_2, z_3) = 1$ and $d_{\mathbb{R}^m}(z_1, z_3) = 2$. Because $d_{\mathbb{R}^m}(z_1, z_3) = d_{\mathbb{R}^m}(z_1, z_2) + d_{\mathbb{R}^m}(z_2, z_3)$, the triangle is flat, i.e., the points z_1, z_2, z_3 are collinear. Applying the same reasoning, we conclude that the points z_1, z_4, z_3 are collinear, which implies that $z_2 = z_4$ and consequently, $d_{\mathbb{R}^m}(z_2, z_4) = 0$, contradicting the assumption that $d_{\mathbb{R}^m}(z_2, z_4) = d_X(x_2, x_4) = 1$. Because we

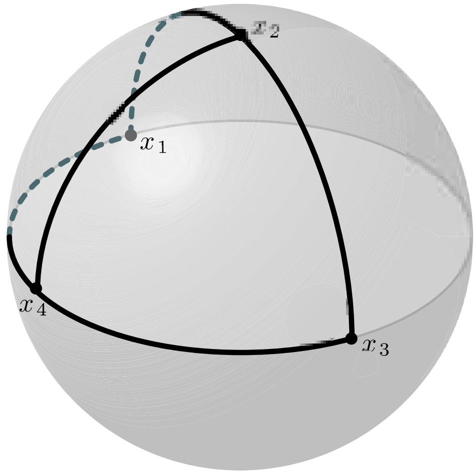

Figure 7.3. Linial's example of a metric space obtained by sampling the sphere at four points, which cannot be embedded into a Euclidean space of any finite dimension.

have not assumed any particular m, we conclude that the given structure cannot be isometrically embedded into a Euclidean space of any finite dimension. Moreover, if this is the case for such a simple object as a sphere, our conclusion is that a general shape cannot be isometrically embedded into $\mathbb{R}^m$.

It is important to emphasize that this result by no means contradicts the Nash embedding theorem. Nash guarantees that any Riemannian structure can be realized as a length metric induced by a Euclidean metric, whereas we are trying to realize it using the restricted Euclidean metric.

Although the embedding error makes it impossible to find a truly isometric embedding, we could try constructing an approximate representation of the shape X, looking for a *minimum-distortion embedding*, i.e., such f that distorts d_X the least, in the sense of some criterion. In Chapter 2, we defined the distortion, which reads in our problem as

$$\operatorname{dis} f = \sup_{x,x' \in X} |d_X(x,x') - d_{\mathbb{R}^m}(f(x), f(x'))|.$$

Adopting this criterion, we can measure how the distances on the original shape differ from those in the embedding space in the sense of the L_∞-norm. In practice, it is useful to replace the L_∞ criterion by an L_p analog,

$$\sigma_p = \int_{X \times X} |d_{\mathbb{R}^m}(f(x), f(x')) - d_X(x,x')|^p \, da \times da, \tag{7.1}$$

where da denotes the area element on X. The L_∞ criterion can be obtained as the limit of $(\sigma_p)^{1/p}$ when $p \to \infty$.

In the discrete setting, when the shape X is sampled at N points $\{x_1, \ldots, x_N\}$, the L_∞ criterion becomes

$$\sigma_\infty = \max_{i,j=1,\ldots,N} |d_{\mathbb{R}^m}(f(x_i), f(x_j)) - d_X(x_i, x_j)|,$$

and the discrete version of the L_p criterion can be expressed as

$$\sigma_p = \sum_{i>j} a_i a_j |d_{\mathbb{R}^m}(f(x_i), f(x_j)) - d_X(x_i, x_j)|^p, \quad (7.2)$$

where a_i and a_j are discrete area elements corresponding with the points x_i, x_j. If the shape is sampled uniformly, we can simplify σ_p by setting $a_i = 1/N$.

The canonical form obtained by means of a minimum-distortion embedding is only an approximate representation of the shape's intrinsic geometry. Nevertheless, we can still measure the similarity of shapes and the distance between their canonical forms, of course, having in mind that the distortion introduced by the embedding would influence the accuracy of such a similarity.

7.2 Multidimensional scaling

An important question is how to find the minimum-distortion embedding in practice. Assume that the shape X is uniformly sampled at points $\{x_1, \ldots, x_N\}$ and σ_2 is used as the distortion criterion. We are looking for the minimum-distortion embedding into $\mathbb{R}^m$,

$$f = \underset{f:X\to\mathbb{R}^m}{\operatorname{argmin}} \sum_{i>j} |d_{\mathbb{R}^m}(f(x_i), f(x_j)) - d_X(x_i, x_j)|^2.$$

Denoting by $z_i = f(x_i)$ the m-dimensional Euclidean coordinates of the image of the shape sample x_i under f and arranging them into an $N \times m$ matrix $Z = (z_i^j)$, we can rewrite our distortion criterion as

$$\sigma_2(Z; D_X) = \sum_{i>j} |d_{ij}(Z) - d_X(x_i, x_j)|^2.$$

Here $D_X = (d_X(x_i, x_j))$ is an $N \times N$ matrix of geodesic distances and $d_{ij}(Z)$ is a shorthand notation for the Euclidean distance between the i-th and the j-th points on the canonical form, $d_{\mathbb{R}^m}(z_i, z_j) = \|z_i - z_j\|_2$. In this formulation, we find the coordinates of the discrete canonical form directly as the solution of a nonlinear least-squares problem,

$$Z^* = \underset{Z\in\mathbb{R}^{N\times m}}{\operatorname{argmin}} \sigma_2(Z). \quad (7.3)$$

Problem (7.3) is a non-convex optimization problem [380] in Nm variables, in which the objective function $\sigma_2 : \mathbb{R}^{N\times m} \to \mathbb{R}$ is defined over the space of