We conclude that while being reasonably efficient to compute and consistent to sampling, the embedding distance does not satisfy all of the desired properties. In particular, it is not a metric, and, which is by far worse, the connection between similarity and the distance values works only in one direction.

10.3 Gromov-Hausdorff distance

Another type of similarity we met in Chapter 7 was the canonical form distance, computed as the Hausdorff distance between the minimum-distortion embeddings of two shapes into some common metric space $(\mathbb{Z}, d_{\mathbb{Z}})$. For example, we used the Euclidean space $\mathbb{R}^m$ or the m-dimensional sphere $\mathbb{S}^m$. As we mentioned, this approach suffers from an inherent inaccuracy due to the fact that usually a zero-distortion embedding of a surface into a given metric space is impossible to obtain.

Instead of having $\mathbb{Z}$ fixed for all surfaces, we can let $\mathbb{Z}$ be the best suitable space for the comparison of two given surfaces X and Y by introducing it as a variable into our optimization problem. Formally, we can write the following distance,

$$d_{\mathrm{GH}}(Y,X) = \inf_{\substack{\mathbb{Z} \\ f:X\to\mathbb{Z} \\ g:Y\to\mathbb{Z}}} d_{\mathrm{H},\mathbb{Z}}(f(X), g(Y)), \tag{10.2}$$

where the infimum is taken over all metric spaces $\mathbb{Z}$ and isometric embeddings f and g from X and Y, respectively, to $\mathbb{Z}$. d_{GH} is called the *Gromov-Hausdorff* distance and can be thought of as an extension of the Hausdorff distance. The Gromov-Hausdorff distance was introduced in 1981 by the Russian-born mathematician Mikhail Gromov [190] and first applied to the field of pattern recognition by Facundo Mémoli and Guillermo Sapiro in 2004 [269].

At this point, the reader might wonder whether the space $\mathbb{Z}$ in (10.2) exists at all. In fact, we demand that $\mathbb{Z}$ has two metric subspaces $f(X)$ and $g(Y)$ with the restriction of $d_{\mathbb{Z}}$ that are isometric to X and Y, respectively, which seems like a very strong property. However, it appears that such a space always exists; moreover, we can even reduce it to the disjoint union of X and Y. More precisely, we may let $\mathbb{Z} = X \sqcup Y$ and define a (semi-) metric $d_{\mathbb{Z}}$ such that its restrictions to X and Y coincide with d_X and d_Y (clearly, $d_{\mathbb{Z}}$ is not unique). Using this reduction, the Gromov-Hausdorff distance can be reformulated in terms of the infimum over all the metrics $d_{\mathbb{Z}}$ on $X \sqcup Y$,

$$d_{\mathrm{GH}}(Y,X) = \inf_{d_{\mathbb{Z}}} d_{\mathrm{H},(X\sqcup Y, d_{\mathbb{Z}})}(Y,X). \tag{10.3}$$

The Gromov-Hausdorff distance brings us excellent news, as it satisfies all the desired theoretical properties: it is a metric on $\mathbb{M}^*$, and it satisfies the similarity property with the constant $c = 2$ (namely, $d_{\mathrm{GH}}(X,Y) < \epsilon$ implies that X and Y are 2ϵ-isometric, and X and Y are ϵ-isometric implies that $d_{\mathrm{GH}}(X,Y) < 2\epsilon$). We leave the proof as an exercise for the reader (Problem 10.1).

At the first glance, the theoretical properties of the Gromov-Hausdorff distance make it a perfect choice for comparing non-rigid shapes. Yet, from definition (10.2), d_{GH} seems alarmingly impractical, as the minimization over all metric spaces $\mathbb{Z}$ or over all metrics $d_{\mathbb{Z}}$ on $X \sqcup Y$ is intractable. Fortunately, the Gromov-Hausdorff distance can be reformulated in terms of distances in X and Y, without resorting to the embedding space $\mathbb{Z}$:

$$d_{\mathrm{GH}}(Y,X) = \frac{1}{2} \inf_{\substack{\varphi : X \to Y \\ \psi : Y \to X}} \max\{\mathrm{dis}\,\varphi, \mathrm{dis}\,\psi, \mathrm{dis}\,(\varphi,\psi)\}. \tag{10.4}$$

Let us understand the notation first. The first two terms,

$$\begin{aligned} \mathrm{dis}\,\varphi &= \sup_{x,x' \in X} |d_X(x,x') - d_Y(\varphi(x),\varphi(x'))|; \\ \mathrm{dis}\,\psi &= \sup_{y,y' \in Y} |d_Y(y,y') - d_X(\psi(y),\psi(y'))|, \end{aligned}$$

denote the familiar distortion of the embeddings φ and ψ, respectively. On the other hand, the term $\mathrm{dis}\,(\varphi,\psi)$ is new and is used to denote

$$\mathrm{dis}\,(\varphi,\psi) = \sup_{x \in X, y \in Y} |d_X(x,\psi(y)) - d_Y(y,\varphi(x))|$$

(see Figure 10.1). This reformulation of the Gromov-Hausdorff distance can be interpreted in the following way: we try to jointly embed X into Y and Y into X such that the distortions of the embeddings φ and ψ are as low as possible. In addition, we would like φ and ψ to be as close as possible one to the inverse of the other, in the sense that the compositions $\psi \circ \varphi : X \to Y$ and $\varphi \circ \psi : Y \to X$ are as close as possible to identity mappings (Figure 10.2). We leave to the reader to prove this extraordinary transfiguration of d_{GH} (Problem 10.2).

In its alternative formulation (10.4), the Gromov-Hausdorff distance closely resembles the embedding distance. The only difference is that now we have a slightly more complicated problem involving two embeddings and three distortion terms, yet it can be solved in the same spirit. Indeed, discretizing X and Y, we obtain

$$d_{\mathrm{GH}}(Y_M,X_N) = \frac{1}{2} \min_{\substack{y'_1,\ldots,y'_N \in Y \\ x'_1,\ldots,x'_M \in X}} \max \left\{ \begin{array}{l} |d_X(x_i,x_j) - d_Y(y'_i,y'_j)|, \\ |d_Y(y_k,y_l) - d_X(x'_k,x'_l)|, \\ |d_X(x_i,x'_k) - d_Y(y_k,y'_i)| \end{array} \right\}, \tag{10.5}$$

where $i,j = 1,\ldots,N$, and $k,l = 1,\ldots,M$. This problem can be viewed as a simultaneous solution of two L_∞ GMDS problems, coupled together by the

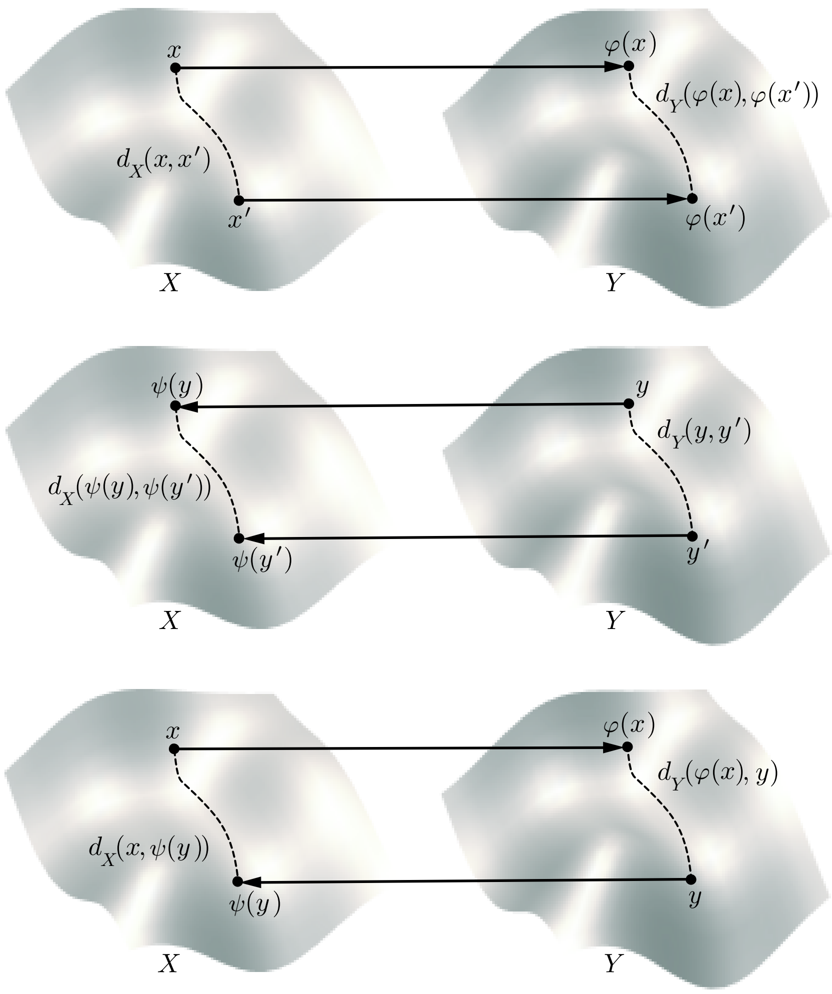

Figure 10.1. An illustration of the three distortion terms participating in the Gromov-Hausdorff distance: $\text{dis}\,\varphi$ (top row), $\text{dis}\,\psi$ (middle row), and $\text{dis}\,(\varphi, \psi)$ (bottom row).

distortion terms $|d_X(x_i, x'_k) - d_Y(y_k, y'_i)|$. As in the GMDS, the minimization is performed over the images $y'_i = \varphi(x_i)$ and $x'_k = \psi(y_k)$, instead of the mappings φ and ψ themselves. The only difference is that now we have two sets of variables: one on the surface Y and the other on the surface X.

One of the ways to solve (10.5) is by introducing an artificial variable $\epsilon \geq 0$ and casting the min-max problem to the following constrained minimization problem,

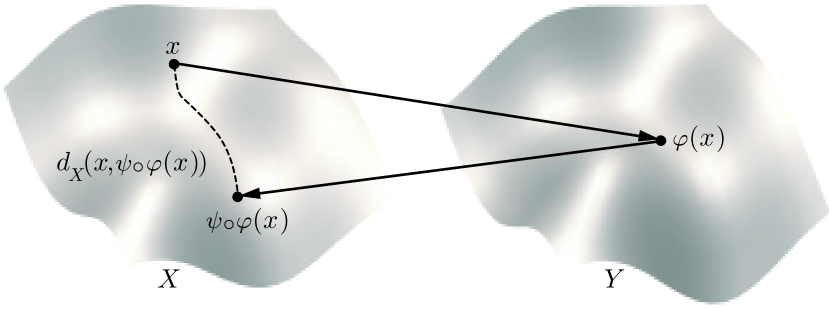

Figure 10.2. The distortion term $\mathrm{dis}\,(\varphi,\psi)$ tells us how far is φ from the inverse of ψ and vice versa, how far is ψ from the inverse of φ.

$$d_{\mathrm{GH}}(Y_M, X_N) = \min_{\substack{\epsilon\geq 0 \\ y'_1,\ldots,y'_N\in Y \\ x'_1,\ldots,x'_M\in X}} \frac{\epsilon}{2} \quad \text{s.t.} \quad \begin{cases} |d_X(x_i,x_j) - d_Y(y'_i,y'_j)| \leq \epsilon \\ |d_Y(y_k,y_l) - d_X(x'_k,x'_l)| \leq \epsilon \\ |d_X(x_i,x'_k) - d_Y(y_k,y'_i)| \leq \epsilon. \end{cases} \tag{10.6}$$

Example 10.1. As a visualization of the performance of d_{GH}, we reproduce here a numerical experiment from [67], performed on the set of non-rigid shapes. The Gromov-Hausdorff distances were computed numerically between each pair of objects. The matching results are visualized in Figure 10.3 as dissimilarities in $\mathbb{R}^2$ (i.e., the closer are the points, the smaller is the distance between the corresponding objects).

It should be noted that our analysis was done with the L_∞ distortion. The reader may wonder whether similar results can be made for some L_p formulation of the distortion. Unfortunately, it appears that the properties of the Gromov-Hausdorff distance change dramatically when L_∞ is substituted by L_p. For example, the beautiful connection to the Hausdorff distance ceases to exist (although, the intriguing question whether an L_p version of d_{GH} is connected to some L_p version of the Hausdorff distance is still open). Having said that, an L_p formulation of the Gromov-Hausdorff distance is clearly useful for practical non-rigid surface matching and may be even more useful than its L_∞ counterpart due to its lower sensitivity to noise.

10.4 Intrinsic symmetry

This chapter dedicated to shape similarity would be incomplete without mentioning an interesting and beautiful particular case of shape *self-similarity*