Int J Comput Vis
DOI 10.1007/s11263-007-0078-4

Analysis of Two-Dimensional Non-Rigid Shapes

Alexander M. Bronstein · Michael M. Bronstein · Alfred M. Bruckstein · Ron Kimmel

Received: 3 November 2006 / Accepted: 23 July 2007

Abstract Analysis of deformable two-dimensional shapes is an important problem, encountered in numerous pattern recognition, computer vision and computer graphics applications. In this paper, we address three major problems in the analysis of non-rigid shapes: similarity, partial similarity, and correspondence. We present an axiomatic construction of similarity criteria for deformation-invariant shape comparison, based on intrinsic geometric properties of the shapes, and show that such criteria are related to the Gromov-Hausdorff distance. Next, we extend the problem of similarity computation to shapes which have similar parts but are dissimilar when considered as a whole, and present a construction of set-valued distances, based on the notion of Pareto optimality. Finally, we show that the correspondence between non-rigid shapes can be obtained as a byproduct of the non-rigid similarity problem. As a numerical framework, we use the generalized multidimensional scaling (GMDS) method, which is the numerical core of the three problems addressed in this paper.

Keywords Non-rigid shapes · Partial similarity · Pareto optimum · Multidimensional scaling · GMDS · Gromov-Hausdorff distance · Intrinsic geometry

A.M. Bronstein · M.M. Bronstein (✉) · A.M. Bruckstein · R. Kimmel
Department of Computer Science, Technion—Israel Institute of Technology, 32000 Haifa, Israel
e-mail: bronstein@ieee.org

A.M. Bronstein
e-mail: alexbron@ieee.org

A.M. Bruckstein
e-mail: freddy@cs.technion.ac.il

R. Kimmel
e-mail: ron@cs.technion.ac.il

1 Introduction

Many of the objects surrounding us in the world are non-rigid and, due to their physical properties, can undergo deformations. Such objects are encountered at various resolution levels—from amoebae on microscopic scales to bodies of humans and animals on macroscopic ones. Modeling and understanding the behavior of such objects is an important problem in pattern recognition, computer vision and computer graphics, and has recently attracted significant attention in different applications (Bronstein et al. 2007b).

We outline three major problems in the analysis of non-rigid shapes, which are explored in this paper:

- *Deformation-invariant comparison* (Fig. 1a): finding a similarity criterion between shapes insensitive to the deformations they undergo;
- *Partial comparison* (Fig. 1b): finding similarity of deformable shapes which have only partial similarity, i.e., have similar as well as dissimilar parts.
- *Correspondence* (Fig. 1c): finding correspondence between points on deformable shapes.

The problems of similarity and correspondence are intimately related, and in most cases, solving one problem allows to solve the other. Broadly speaking, similarity and correspondence can be thought of as two archetype problems in the analysis of non-rigid shapes; the similarity problem is often encountered in computer vision and pattern recognition applications, whereas that of correspondence arises in computer graphics and geometry processing. Partial similarity is a more general setting of the shape similarity problem, in which the shapes have similar parts but are dissimilar when considered as a whole.

The main difficulty in analyzing non-rigid shapes stems from the fact that their geometry varies and it is not clear

Fig. 1 (Color online) Three main problems in the analysis of non-rigid shapes

a. Comparison of non-rigid shapes

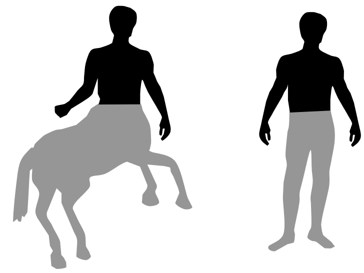

b. Partial comparison of non-rigid shapes

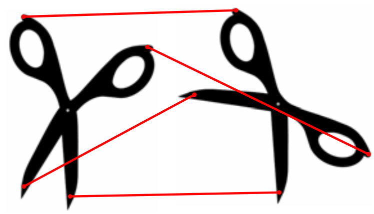

c. Correspondence between non-rigid shapes

what quantities characterize the shape and which can be attributed to the deformation. Recently, several methods for deformation-invariant description of shapes have been proposed, targeting mainly three-dimensional objects. Elad and Kimmel (2001) used geodesic distances as invariant descriptors of three-dimensional non-rigid shapes under the class of isometric deformations. Their approach created a representation of the intrinsic geometry of a shape (referred to as the *canonical form*) by finding a minimum-distortion embedding into a Euclidean space. The embedding was performed by applying *multidimensional scaling* (MDS) to the geodesic distances. This approach showed good results in the problem of expression-invariant three-dimensional face recognition, where the deformations of the facial surface due to expressions were modeled as near-isometries (Bronstein et al. 2003, 2005a). However, the canonical forms approach allows only for an approximate representation of the intrinsic geometry, since usually a shape cannot be isometrically embedded into a Euclidean space.

In a follow-up work, Mémoli and Sapiro (2005) proposed using the Gromov-Hausdorff distance, introduced in Gromov (1981), in order to compare the intrinsic geometries of three-dimensional shapes. Their paper was the first use of this distance in pattern recognition. The Gromov-Hausdorff distance has appealing theoretical properties, and in particular, lacks the inherent inaccuracy of the canonical forms, but its computation is NP-hard. Mémoli and Sapiro suggested an algorithm that approximates the Gromov-Hausdorff distance in polynomial time by computing a different distance related to it by a probabilistic bound. More recently, Bronstein et al. developed an approach, according to which the computation of the Gromov-Hausdorff distance is formulated as a continuous MDS-like problem and solved efficiently using a local minimization algorithm (Bronstein et al. 2006a, 2006b). This numerical framework was given the name of *generalized MDS* (GMDS). GMDS appeared superior to the canonical forms approach in face recognition (Bronstein et al. 2006c) and face animation applications (Bronstein et al. 2006d).

In this paper, we consider a two-dimensional setting of non-rigid shape analysis, where shapes are planar and can be thought of as "silhouettes" of deformable objects. Analysis of such shapes is often encountered in the computer vision literature (Fischler and Elschlager 1973; Ullman 1989; Grenander et al. 1991; Mumford 1991; Stark and Bowyer 1991; Lades et al. 1993; Rivlin et al. 1992; Geiger et al. 1998, 2003; Gdalyahu and Weinshall 1999; Latecki and Lakamper 2000; Cheng et al. 2001; Belongie et al. 2002; Felzenszwalb and Huttenlocher 2005), typically as a subset of the more generic problem of image analysis (Platel et al. 2005; Berg et al. 2005; Jacobs and Ling 2005).

One of the mainstream approaches is representing shape contours as planar curves and posing the shape similarity as a problem of deformable curve comparison, generally referred to as *elastic matching*. The latter problem is usually solved by deforming one curve into another and defining the similarity of curves as the "difficulty" to perform such a deformation. Different criteria of such "difficulty" were proposed (Burr 1981; Tappert 1982; Hildreth 1983; Kass et al. 1988; Yuille et al. 1989; Cohen et al. 1992; Jain et al. 1996; Felzenszwalb 2005), in most cases, inspired by physical considerations. Elastic matching can be performed in a hierarchical manner (Felzenszwalb and Schwartz 2007). A more recent viewpoint, pioneered by Michor and Mumford (2003) and later extended in Yezzi and Mennucci (2004), Mio et al. (2007), Charpiat et al. (2007), considered the space of curves as an infinite-dimensional Riemannian manifold and endowed it with a distance structure, which was used to measure the similarity of two curves.

Another mainstream approach suggest computing the similarity of non-rigid shapes by dividing them into parts and comparing the parts as separate objects (Binford 1987; Brooks 1981; Hoffman and Richards 1984; Biderman 1985; Bajcsy and Solina 1987; Pentland 1987; Connell and Brady 1987; Hel-Or and Werman 1994; Kupeev and Wolfson 1994; Siddiqi and Kimia 1996), which allows, at least in theory, to address the problem of partial comparison as well. However, there are several difficulties in such approaches. The first one is the problem of division of the shape into meaningful parts. There is no obvious definition of what is a part of a shape, and thus, results may vary depending on what method is used to divide the shape. The second difficulty is the question of how to "integrate" similarities between different parts into a global similarity measure of the entire shape (Geiger et al. 1998).

A simplified approach to non-rigid shape analysis is based on the *articulated shape* model, which assumes that non-rigid shapes are composed of rigid parts, each of which has a certain freedom to move (Zhang et al. 2004). In the recent work of Ling and Jacobs (2005), this model was used implicitly in order to claim that the *intrinsic geometry* of such shapes is nearly invariant. The geodesic distances measured in the shapes are used as deformation-insensitive descriptors, in the spirit of Elad and Kimmel (2001).

In this paper, we approach the problem of non-rigid shape analysis from the intrinsic geometric point of view, following Ling and Jacobs (2005). We start with formulating a set of desired properties that a good similarity or correspondence criterion should satisfy. We show that the Gromov-Hausdorff distance satisfies these properties, while the canonical forms distance does not. We apply the axiomatic construction to the correspondence problem, and extend it in order to cope with the partial similarity. Numerically, all the problems are formulated as instances of GMDS, which allows for a computationally-efficient solution.

The paper consists of eight sections and is organized as follows. In Sect. 2, we present a model of deformable shapes. In Sect. 3, we formulate a set of axioms that an ideal deformation-invariant similarity criterion should satisfy, and compare how different distances fit into this axiomatic construction. In Sect. 4, we introduce set-valued distances based on the notion of Pareto optimality to address the problem of partial similarity. Section 5 addresses the problem of correspondence between non-rigid shapes. Section 6 deals with numerical computation of the distance and correspondence between non-rigid shapes using GMDS. In Sect. 7, we present experimental validations of our approach. Section 8 concludes the paper. The proofs of the main results are given in the Appendix.

2 Isometric Model for Deformable Shapes

2.1 Definitions

A two-dimensional *shape* $\mathcal{S}$ is modeled as a compact two-dimensional manifold with boundary, embedded in $\mathbb{R}^2$. The space of all shapes, in which $\mathcal{S}$ corresponds to a point, is denoted by $\mathbb{M}$. A *minimal geodesic* is the shortest path between points s_1, s_2 in $\mathcal{S}$. It consists of linear segments and parts of the boundary (Ling and Jacobs 2005). The *geodesic distance* $d_{\mathcal{S}}(s_1, s_2)$ is the length of the minimal geodesic between s_1 and s_2. It is important to stress the difference between the *induced* and the *restricted* metric. The latter, denoted by $d_{\mathbb{R}^2}|_{\mathcal{S}}$, measures the distances in $\mathcal{S}$ using the metric of $\mathbb{R}^2$, i.e., $d_{\mathbb{R}^2}|_{\mathcal{S}}(s_1, s_2) = d_{\mathbb{R}^2}(s_1, s_2)$ for all $s_1, s_2 \in \mathcal{S}$. The induced metric $d_{\mathcal{S}}$, on the other hand, measures the length of the geodesics in $\mathcal{S}$. The pair $(\mathcal{S}, d_{\mathcal{S}})$ is a *metric space*; quantities expressible in terms of $d_{\mathcal{S}}$ are referred to as *intrinsic*. The intrinsic geometry of a two-dimensional shape is completely defined by its boundary. This is a fundamental difference between two-dimensional shapes (flat manifolds) and three-dimensional shapes, which may have non-trivial curvature. Note that although $(\mathcal{S}, d_{\mathcal{S}})$ is part of a larger metric space $(\mathbb{R}^2, d_{\mathbb{R}^2})$, from the intrinsic point of view, there exists nothing "outside" $\mathcal{S}$. We further assume that the *measure* $\mu_{\mathcal{S}}$, induced by the Riemannian structure, is defined on $\mathcal{S}$. Informally speaking, for a subset $\mathcal{S}' \subseteq \mathcal{S}$, we can think of $\mu_{\mathcal{S}}(\mathcal{S}')$ as of the area of $\mathcal{S}'$ and express it in units of squared distance.

In practical applications, shapes are usually represented as discrete binary images sampled at a finite number of points (pixels). A set $\mathcal{S}^r \subset \mathcal{S}$ is said to be an *r-covering* of $\mathcal{S}$, if $\bigcup_{i=1}^{N} B_{\mathcal{S}}(s_i, r) = \mathcal{S}$, where $B_{\mathcal{S}}(s_i, r)$ denotes a ball of radius r with respect to the metric $d_{\mathcal{S}}$, centered at s_i. Since the shapes are assumed to be compact, every shape has a finite r-covering $\mathcal{S}^r_N = \{s_1, \ldots, s_N\}$ for every $r > 0$. The measure $\mu_{\mathcal{S}}$ is discretized by constructing a discrete measure $\mu_{\mathcal{S}_N} = \{\mu_1, \ldots, \mu_N\}$, assigning to each $s_i \in \mathcal{S}^r_N$ the area of the corresponding *Voronoi cell*,

$$\mu_i = \mu(\{s \in \mathcal{S} : d_{\mathcal{S}}(s, s_i) < d_{\mathcal{S}}(s, s_j)\ \forall j \neq i\}). \qquad (1)$$

For brevity, we denote the discrete metric measure space $(\mathcal{S}^r_N, d_{\mathcal{S}}|_{\mathcal{S}^r_N}, \mu_{\mathcal{S}_N})$ by $\mathcal{S}^r_N$ and refer to it as an *r-sampling* of $\mathcal{S}$.

2.2 Isometric Shapes

Let $\mathcal{S}, \mathcal{Q}$ be two shapes in $\mathbb{M}$. A map $f : \mathcal{S} \to \mathcal{Q}$ is said to have *distortion* ϵ if

$$\operatorname{dis} f \equiv \sup_{s_1, s_2 \in \mathcal{S}} |d_{\mathcal{S}}(s_1, s_2) - d_{\mathcal{Q}}(f(s_1), f(s_2))| = \epsilon. \tag{2}$$

We call such an f an ϵ-*isometric embedding* of $\mathcal{S}$ into $\mathcal{Q}$. If in addition f is ϵ-*surjective*, i.e. $\bigcup_{q \in f(\mathcal{S})} B_{\mathcal{Q}}(q, \epsilon) = \mathcal{Q}$, we call f an ϵ-*isometry* and say that the shapes $\mathcal{S}$ and $\mathcal{Q}$ are ϵ-*isometric*. In the particular case of $\epsilon = 0$, the shapes are said to be *isometric* and f is called an *isometry*. True isometries are cardinally different from ϵ-isometries. Particularly, an isometry is always bi-Lipschitz continuous (Burago et al. 2001), which is not necessarily true for an ϵ-isometry.

Isometries from $\mathcal{S}$ to itself are called *self-isometries*; with the function composition operator, self-isometries form the *isometry group* which we denote by $\operatorname{Iso}(\mathcal{S})$. The most obvious self-isometry is the identity map $id : \mathcal{S} \to \mathcal{S}$, which copies every point on $\mathcal{S}$ into itself. Normally, $\operatorname{Iso}(\mathcal{S})$ would contain only id, a case in which it is said to be *trivial*. However, if the shape has symmetries, the isometry group is not trivial and may contain other self-isometries different from the identity. For example, if $\mathcal{S}$ is a planar triangle with two equal sides unequal to the third, the isometry group is the cyclic group $\mathbb{Z}/2\mathbb{Z}$. This group contains only two elements: the identity transformation and the reflection transformation, which flips the triangle about its symmetry axis.

2.3 Articulated Shapes

A shape $\mathcal{S}$ consisting of K disjoint *parts* $\mathcal{S}_1, \ldots, \mathcal{S}_K$ and L *joints* $\mathcal{J}_1, \ldots, \mathcal{J}_L$, such that

$$\mathcal{S} = (\mathcal{S}_1 \cup \cdots \cup \mathcal{S}_K) \cup (\mathcal{J}_1 \cup \cdots \cup \mathcal{J}_L), \tag{3}$$

is called an *articulated shape*. An example of an articulated shape is shown in Fig. 2. We call an articulated shape with $\sum_{i=1}^{L} \operatorname{diam}(\mathcal{J}_i) \leq \epsilon$ an ϵ-*articulated shape* (here, $\operatorname{diam}(\mathcal{J}_i) = \sup_{s,s' \in \mathcal{J}_i} d_{\mathcal{S}}(s, s')$ denotes the *diameter* of $\mathcal{J}_i$).

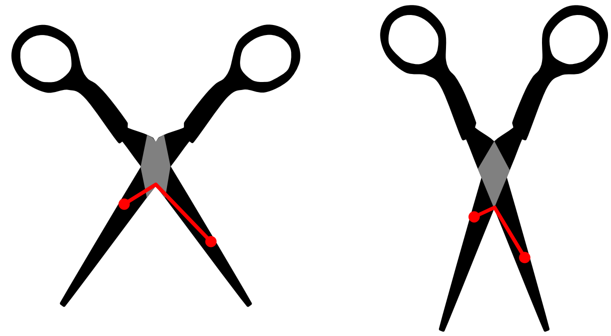

Fig. 2 (Color online) Example of an ϵ-articulated shape, consisting of four parts (*black*) and one joint (*gray*). The geodesic distance between two points is shown in *red*. Note that the geodesic distances change is bounded by the diameter of the joint

We denote by $\mathbb{M}_\epsilon$ the space of all ϵ-articulated shapes; $\mathbb{M}$ coincides with $\mathbb{M}_\infty$.

Given $\mathcal{S} \in \mathbb{M}_\epsilon$, an *articulation* is a topology-preserving map $f : \mathcal{S} \to \mathcal{S}'$, which isometrically maps parts $\mathcal{S}_1, \ldots, \mathcal{S}_K$ into parts $\mathcal{S}'_1, \ldots, \mathcal{S}'_K$, and maps joints $\mathcal{J}_1, \ldots, \mathcal{J}_L$ into joints $\mathcal{J}'_1, \ldots, \mathcal{J}'_L$ such that $\sum_{i=1}^{L} \operatorname{diam}(\mathcal{J}'_i) \leq \epsilon$, or in other words, $\mathcal{S}'$ is also an ϵ-articulated shape.

Proposition 1 *Articulations of an ϵ-articulated shape are ϵ-isometries.*

The proof is technical and can be found in Ling and Jacobs (2005). The converse of Proposition 1 is not true: an ϵ-isometry is not necessarily an articulation. Figure 3 illustrates this difference showing the skeleton of a human palm, which is an ϵ-articulated shape (left). The skeleton is articulated by moving the bones while keeping them connected (middle); the two postures of the skeleton are ϵ-isometric. On the other hand, Fig. 3 (right) shows another ϵ-isometry of the skeleton, which is not an articulation. Another difference between articulations of ϵ-articulated shapes and ϵ-isometries is the *closure property*. A composition of two articulations leaves the shape ϵ-articulated; on the other hand, a composition of two ϵ-isometries is generally a 2ϵ-isometry.

Fig. 3 The difference between an articulation of an ϵ-articulated shape (*middle*) and an ϵ-isometry (*right*)

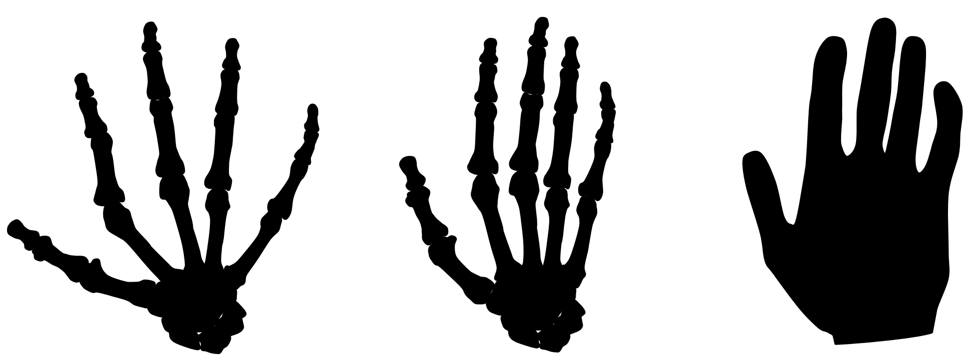

An *ideal* or 0-*articulated shape* has point joints; its articulations are true isometries. Such shapes rarely occur in practice, yet, the joints can be often assumed significantly smaller compared to the parts, i.e., $\min_{i=1,\ldots,K} \mathrm{diam}(\mathcal{S}_i) \gg \epsilon$ (Ling and Jacobs 2005).

3 Axiomatic Approach to Shape Comparison

Our starting point in the analysis of shapes is the problem of shape comparison. We will refer to this problem as that of *full comparison*, to distinguish it from *partial comparison* discussed later. When we say that two shapes are similar or dissimilar, we can quantitatively express this degree of dissimilarity as a distance $d_{\mathrm{F}} : \mathbb{M} \times \mathbb{M} \to \mathbb{R}$. The definition of similarity is rather a semantic question and cannot be addressed in a univalent manner. Ling and Jacobs (2005) showed that the comparison of intrinsic geometries is a good similarity criterion for articulated shapes. Here, we present an axiomatic construction of a d_{F} by first listing a set of properties such a distance should satisfy, and then introducing a distance that satisfies these properties.

3.1 Isometry-Invariant and Articulation-Invariant Distances

Since we use the intrinsic geometry to compare shapes, the most fundamental property that must hold is isometry-invariance, which implies that two shapes that are related by an isometry are similar. An ideal distance must satisfy for all $\mathcal{S}, \mathcal{Q}, \mathcal{R} \in \mathbb{M}$ the following list of axioms; we denote such a distance by d_{F}:

(F1) *Non-negativity*: $d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}) \geq 0$.

(F2) *Symmetry*: $d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}) = d_{\mathrm{F}}(\mathcal{S}, \mathcal{Q})$.

(F3) *Triangle inequality*: $d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}) \leq d_{\mathrm{F}}(\mathcal{Q}, \mathcal{R}) + d_{\mathrm{F}}(\mathcal{R}, \mathcal{S})$.

(F4) *Isometry-invariant similarity*: (i) If $d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}) \leq \epsilon$, then $\mathcal{S}$ and $\mathcal{Q}$ are $c\epsilon$-isometric; (ii) if $\mathcal{S}$ and $\mathcal{Q}$ are ϵ-isometric, then $d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}) \leq c\epsilon$, where c is some positive constant, independent of $\mathcal{S}$, $\mathcal{Q}$, and ϵ.

Property (F4) guarantees that d_{F} is a good similarity criterion, assigning large distances for dissimilar shapes and small distances for similar (nearly isometric) ones. A particular case of (F4) is the *isometry invariance* property: $d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}) = 0$ if and only if $\mathcal{S}$ and $\mathcal{Q}$ are isometric (note that our definition of similarity is not scale invariant). Together, (F1)–(F4) guarantee that d_{F} is a metric on the quotient space $\mathbb{M}\backslash\mathrm{Iso}(\mathbb{M})$ (equivalence class of all isometric shapes, in which a point represents a shape and all its isometries).

Since we want the distance to be computable in practice, we add another property:

(F5) *Consistency to sampling:* If $\mathcal{S}^r$ is a finite r-sampling of $\mathcal{S}$, then

$$\lim_{r \to 0} d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}^r) = d_{\mathrm{F}}(\mathcal{Q}, \mathcal{S}).$$

Property (F5) allows to discretize the continuous distance and approximate it on a finite sampling of points. It is tacitly assumed that the discrete distance can be efficiently computed or approximated.

If $\mathcal{S}$ is an ϵ-articulated shape, according to Proposition 1 we have that an articulation f is an ϵ-isometry. Therefore, a distance satisfying the set of properties (F) will guarantee that $d_{\mathrm{F}}(\mathcal{S}, f(\mathcal{S})) \leq c\epsilon$. Ideally, we would also like to be able to say the converse: if $d_{\mathrm{F}}(\mathcal{S}, \mathcal{Q}) \leq \epsilon$ and $\mathcal{S} \in \mathbb{M}_{c\epsilon}$, then there exists an articulation f of $\mathcal{S}$ such that $\mathcal{Q} = f(\mathcal{S})$. Yet, this is not true, since an ϵ-isometry is not necessarily an articulation. We formulate this as a weaker property:

(F4′) *Articulation-invariant dissimilarity:* If f is an articulation of $\mathcal{S} \in \mathbb{M}_{\epsilon}$, then $d_{\mathrm{F}}(\mathcal{S}, f(\mathcal{S})) \leq c\epsilon$, where c is some positive constant, independent of $\mathcal{S}$, f, and ϵ.

3.2 Canonical Forms Distance

Ling and Jacobs (2005) mention the possibility of using the method of *bending-invariant canonical forms*, proposed in Elad and Kimmel (2001) for the comparison of non-rigid surfaces. The key idea of this method consists of representing the intrinsic geometry of the shapes $\mathcal{S}$ and $\mathcal{Q}$ in some metric space $(\mathbb{X}, d_{\mathbb{X}})$, by means of minimum-distortion maps $\varphi : \mathcal{S} \to \mathbb{X}$ and $\psi : \mathcal{Q} \to \mathbb{X}$. The resulting metric subspaces $(\varphi(\mathcal{S}), d_{\mathbb{X}}|_{\varphi(\mathcal{S})})$ and $(\psi(\mathcal{Q}), d_{\mathbb{X}}|_{\psi(\mathcal{Q})})$ of $\mathbb{X}$, are called the *canonical forms* of $\mathcal{S}$ and $\mathcal{Q}$. In this manner, the intrinsic geometry of $\mathcal{S}$ and $\mathcal{Q}$ is replaced by the geometry of $\mathbb{X}$, allowing the reformulation of the distance between $\mathcal{S}$ and $\mathcal{Q}$ as the distance between two sets $\varphi(\mathcal{S})$ and $\psi(\mathcal{Q})$ in a common space $\mathbb{X}$. The process of comparing $\mathcal{S}$ and $\mathcal{Q}$ is done in two steps. First, the canonical forms are computed. Next, the canonical forms are compared using some distance on the subsets of $\mathbb{X}$, treating the canonical forms as rigid surfaces (see an illustration in Fig. 4).

As a particular setting of this approach, we assume here that the canonical form comparison is carried out by means of the *Hausdorff distance*,

$$d_{\mathrm{H}}^{\mathbb{X}}(\mathcal{S}, \mathcal{Q}) = \max\left\{\sup_{s \in \mathcal{S}} d_{\mathbb{X}}(s, \mathcal{Q}), \sup_{q \in \mathcal{Q}} d_{\mathbb{X}}(q, \mathcal{S})\right\}, \tag{4}$$

which acts as a measure of distance between two subsets of a metric space. Here, $d_{\mathbb{X}}(s, \mathcal{Q}) = \inf_{q \in \mathcal{Q}} d_{\mathbb{X}}(s, q)$ denotes the point-to-set distance in $\mathbb{X}$. Since the canonical forms are defined up to isometries in $(\mathbb{X}, d_{\mathbb{X}})$, we define

$$d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{S}) = \inf_{i \in \mathrm{Iso}(\mathbb{X})} d_{\mathrm{H}}^{\mathbb{X}}(i \circ \psi(\mathcal{Q}), \varphi(\mathcal{S})), \tag{5}$$

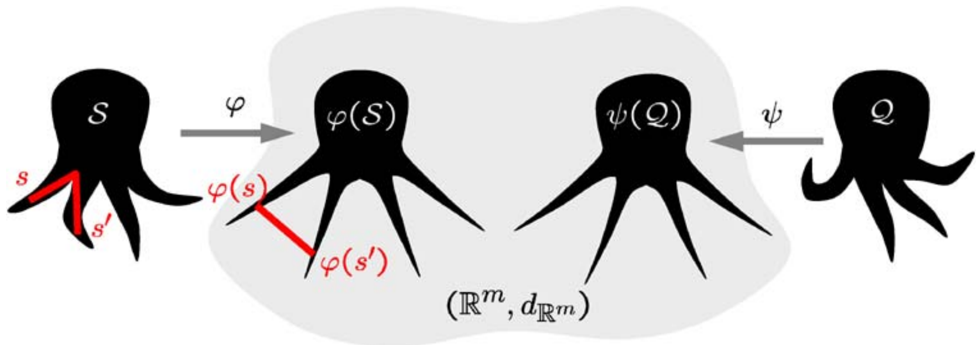

Fig. 4 (Color online) Illustration of the canonical form distance computation

by taking an infimum over all the isometries in the space $\mathbb{X}$. We refer to d_{CF} as the *canonical form distance*.

The embedding space $\mathbb{X}$ is usually chosen as $\mathbb{R}^m$, though other choices are possible (Elad and Kimmel 2002; Bronstein et al. 2005b, 2005c; Walter and Ritter 2002). In general, it is impossible to isometrically embed a non-trivial shape into a given metric space; therefore, the embeddings φ and ψ introduce some distortion. As we will see later, this fact has a fundamental impact on the discriminative power of d_{CF}.

3.3 Gromov-Hausdorff Distance

Instead of using a common embedding space $\mathbb{X}$, we can go one step further and let $\mathbb{X}$ be the best suitable space for the comparison of two given shapes $\mathcal{S}$ and $\mathcal{Q}$. Formally, we can write the following distance,

$$d_{\mathrm{GH}}(\mathcal{Q},\mathcal{S}) = \inf_{\substack{\mathbb{X} \\ \varphi:\mathcal{S}\to\mathbb{X} \\ \psi:\mathcal{Q}\to\mathbb{X}}} d_{\mathrm{H},\mathbb{X}}(\varphi(\mathcal{S}),\psi(\mathcal{Q})), \tag{6}$$

where the infimum is taken over all metric spaces $(\mathbb{X}, d_{\mathbb{X}})$ and isometric embeddings φ and ψ from $\mathcal{S}$ and $\mathcal{Q}$, respectively, to $\mathbb{X}$. d_{GH} is called the *Gromov-Hausdorff* distance (Gromov 1981) and can be thought of as an extension of the Hausdorff distance to arbitrary metric spaces. This distance was first used in the context of isometry-invariant matching of three-dimensional shapes by Mémoli and Sapiro (2005). Illustratively, we can think of the Gromov-Hausdorff distance as of trying all the possible isometries of $\mathcal{S}$ and $\mathcal{Q}$ and matching the resulting shapes using the Hausdorff distance.

Unfortunately, d_{GH} in (6) involving minimization over all metric spaces $\mathbb{X}$ is computationally infeasible, yet, for compact shapes, it can be reformulated in terms of distances in $\mathcal{S}$ and $\mathcal{Q}$, without resorting to the embedding space $\mathbb{X}$:

$$d_{\mathrm{GH}}(\mathcal{Q},\mathcal{S}) = \frac{1}{2}\inf_{\substack{f:\mathcal{S}\to\mathcal{Q} \\ g:\mathcal{Q}\to\mathcal{S}}} \max\{\mathrm{dis}\, f, \mathrm{dis}\, g, \mathrm{dis}(f,g)\}. \tag{7}$$

where the "mixed distortion" term

$$\mathrm{dis}(f,g) = \sup_{s\in\mathcal{S},q\in\mathcal{Q}} |d_{\mathcal{S}}(s,g(q)) - d_{\mathcal{Q}}(q,f(s))| \tag{8}$$

acts as a measure of surjectivity of f and g. For a proof of an equivalence between the two definitions, see Burago et al. (2001).

3.4 Canonical Forms Versus Gromov-Hausdorff

We start our comparison of d_{GH} and d_{CF} from the following result, stemming from the properties of the Gromov-Hausdorff distance.

Theorem 1 *d_{GH} satisfies properties* (F1)–(F5).

We do not provide a rigorous proof here. Properties (F1)–(F4) can be found in Burago et al. (2001). Property (F4) holds with the constant $c = 2$, namely, $d_{\mathrm{GH}}(\mathcal{S},\mathcal{Q}) \le \epsilon$ implies that $\mathcal{S}$ and $\mathcal{Q}$ are 2ϵ-isometric, and $\mathcal{S}$ and $\mathcal{Q}$ are ϵ-isometric implies that $d_{\mathrm{GH}}(\mathcal{S},\mathcal{Q}) \le 2\epsilon$. Property (F5) follows from the fact that given $\mathcal{S}^r$, an r-covering of $\mathcal{S}$, we can always construct a $2r$-isometry between $\mathcal{S}$ and $\mathcal{S}^r$. From (F4), it then follows that $|d_{\mathrm{GH}}(\mathcal{Q},\mathcal{S}^r) - d_{\mathrm{GH}}(\mathcal{Q},\mathcal{S})| \le r$, which in the limit $r \to 0$ gives us (F5).

The computation of the discrete Gromov-Hausdorff distance is an NP-complete combinatorial problem. Mémoli and Sapiro (2005) proposed an algorithm that heuristically approximates the Gromov-Hausdorff distance in polynomial time by computing a different distance related to d_{GH} by a probabilistic bound. Here, we use a different approach, according to which the computation of d_{GH} is formulated as a continuous minimization problem and solved using a local minimization algorithm. We defer this discussion to Sect. 6.

Compared to d_{GH}, the canonical forms distance is significantly weaker. Its properties can be summarized as follows:

Theorem 2 *d_{CF} satisfies properties* (F1)–(F3) *and the following relaxed version of the axiom* (F4):

Fig. 5 (Color online) Illustration of partial similarity intransitivity

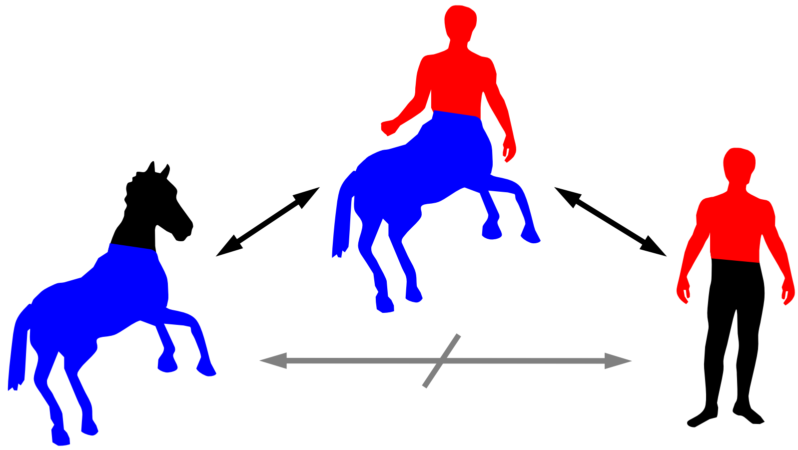

(F4w) Weak similarity: *Let $\mathcal{S}$ and $\mathcal{Q}$ be two shapes, whose canonical forms have the distortions δ and δ', respectively. If $d_{\mathrm{CF}}(\mathcal{Q},\mathcal{S}) \le \epsilon$, then $\mathcal{S}$ and $\mathcal{Q}$ are $2\epsilon + 4(\delta+\delta')$-isometric.*

Theorem 2 allows us consider d_{CF} as an upper bound on d_{GH}. If d_{CF} is small, we can conclude that $\mathcal{S}$ and $\mathcal{Q}$ are similar, but the converse is not guaranteed. Moreover, since the canonical forms have a usually inevitable distortion, the discriminative power of d_{CF} is limited. d_{CF} satisfies the isometry invariance property only approximately: if $d_{\mathrm{CF}}(\mathcal{Q},\mathcal{S}) = 0$, then $\mathcal{S}$ and $\mathcal{Q}$ are $2(\delta+\delta')$-isometric. If $\mathcal{S}$ and $\mathcal{Q}$ are isometric, we cannot say much about $d_{\mathrm{CF}}(\mathcal{Q},\mathcal{S})$. Particularly, two canonical forms of $\mathcal{S}$ may differ significantly. Also, due to the lack of symmetry in (F4w), d_{CF} does not satisfy the sampling consistency property (F5).

4 Partial Comparison of Shapes

So far, discussion the problem of shape similarity, we tacitly assumed that the two shapes were compared as a whole. Our criterion of dissimilarity was the distortion of the map from one shape to another, that is, how non-isometric the two shapes were. In a more general setting, two shapes are not necessarily similar if compared as a whole, yet, may have similar parts. A comparison of shapes taking into account such a possibility is referred to here as *partial comparison*.

In order to better understand the *partial similarity* relation, which we denote by $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q})$, we borrow the mythological creatures example from (Jacobs et al. 2000). A man and a centaur are dissimilar in the sense of a full similarity criterion, yet, parts of these shapes (the upper part of the centaur and the upper part of the man) are similar. Likewise, a horse and a centaur are similar because they share a common part (bottom part of the horse body). At the same time, a man and a horse are dissimilar (Fig. 5). We conclude from that example that the partial similarity relation differs significantly from the full similarity. Particularly, such a relation is *intransitive* (a man and a horse are similar to a centaur, but a man is dissimilar to a horse). This implies that partial similarity is not a metric, as the triangle inequality does not hold.

1. Divide the shapes $\mathcal{S}$ and $\mathcal{Q}$ into parts $\mathcal{S}_1,\ldots,\mathcal{S}_N$ and $\mathcal{Q}_1,\ldots,\mathcal{Q}_M$.
2. Compare each part separately using a full similarity criterion, $d_{\mathrm{F}}(\mathcal{S}_i,\mathcal{Q}_j)$, for all $i=1,\ldots,N$ and $j=1,\ldots,M$.
3. Compute the partial similarity as an aggregate of full similarities between the parts, $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) = \min_{i,j} d_{\mathrm{F}}(\mathcal{S}_i,\mathcal{Q}_j)$.

Algorithm 1 Recognition by parts

Trying to relate partial similarity to full similarity, we can come up with a simple theoretical algorithm for the computation of $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q})$ (see Algorithm 1) Since the parts of non-rigid shapes are non-rigid by themselves, we can use the Gromov-Hausdorff distance as the full similarity criterion.

Trying to implement this simplistic approach, we encounter two problems. First, it is not clear how to divide the shape into parts. Many works on shape partitioning exist in the literature on object recognition, including parts described as convex or near-convex subsets (Hoffman and Richards 1984; Koenderink and van Doorn 1981), primitive geometric objects (Binford 1987; Biderman 1985; Bajcsy and Solina 1987; Pentland 1987) or parametric description derived from a model of the shape class (Brooks 1981; Hel-Or and Werman 1994). The very existence of numerous shape partitioning approaches implies that there is no objective way to define a part, and therefore, the partial similarity criterion obtained in this way is subjective.

Fig. 6 An illustration of the potential danger of partial similarity

This problem can be overcome by considering *all* the possible partitions of the shapes, instead of favoring a specific one (Latecki et al. 2005). For this purpose, we denote by $\Sigma_{\mathcal{S}} \subseteq 2^{\mathcal{S}}$ and $\sum_{\mathcal{Q}} \subseteq 2^{\mathcal{Q}}$ the collections of all the parts of $\mathcal{S}$ and $\mathcal{Q}$, respectively. Here, $2^{\mathcal{S}}$ is the *power set* of $\mathcal{S}$ (the set of all the subsets of $\mathcal{S}$). Technically, we require $\Sigma_{\mathcal{S}}$ (or $\Sigma_{\mathcal{Q}}$, respectively) to be a σ-*algebra*, i.e., to satisfy the following properties:

(S1) *The whole shape is a part of itself*: $\mathcal{S} \in \Sigma_{\mathcal{S}}$.
(S2) *Closure under complement*: of $\mathcal{S}' \in \Sigma_{\mathcal{S}}$, then $\mathcal{S}'^c = (\mathcal{S} \setminus \mathcal{S}') \in \Sigma_{\mathcal{S}}$.
(S3) *Closure under countable union:* of $\mathcal{S}_i \in \Sigma_{\mathcal{S}}$, then $\bigcup_i \mathcal{S}_i \in \Sigma_{\mathcal{S}}$.

The metric on a part $\mathcal{S}' \in \Sigma_{\mathcal{S}}$ is assumed to be $d_{\mathcal{S}}|_{\mathcal{S}'}$. Using these definitions, the computation of the partial similarity can be formulated as the following problem:

$$d_{\mathrm{P}}(\mathcal{S}, \mathcal{Q}) = \min_{(\mathcal{S}', \mathcal{Q}') \in \Sigma_{\mathcal{S}} \times \Sigma_{\mathcal{Q}}} d_{\mathrm{F}}(\mathcal{S}', \mathcal{Q}'). \tag{9}$$

Another problem arises from the possible situation in which two different objects have small similar parts. Relying on the similarity of a such parts, the judgement about the entire shape similarity can be completely wrong. A potential danger of such a situation is depicted in the frivolous cartoon by Herluf Bidstrup (Fig. 6). Our conclusion from this example is that different parts have different importance, and that it is insufficient for the two shapes to have common similar parts in order to be partially similar—the parts must be *significant*. Our visual system appears to have the remarkable capability of recognizing shape form very small significant parts. Significant parts are usually such parts which our prior knowledge can clearly associate the entire object. For example, seing a human eye, we expect it to be part of the human face.

In the absence of additional information, the simplest way to define the significance of a part is by measuring its area: the larger is the part, the more significant it is. Using the measures $\mu_{\mathcal{S}}$ and $\mu_{\mathcal{Q}}$, we define the *partiality* of the parts $\mathcal{S}'$ and $\mathcal{Q}'$,

$$\begin{aligned}\lambda(\mathcal{S}', \mathcal{Q}') &= \mu_{\mathcal{S}}(\mathcal{S}'^c) + \mu_{\mathcal{Q}}(\mathcal{Q}'^c) \\ &= \mu_{\mathcal{S}}(\mathcal{S}) + \mu_{\mathcal{Q}}(\mathcal{Q}) - (\mu_{\mathcal{S}}(\mathcal{S}') + \mu_{\mathcal{Q}}(\mathcal{Q}')),\end{aligned} \tag{10}$$

as the area remaining from the shapes $\mathcal{S}$ and $\mathcal{Q}$ after $\mathcal{S}'$ and $\mathcal{Q}'$ are cropped. Large values of partiality corresponds to small (hence insignificant) parts.

4.1 Multicriterion Optimization and Set-Valued Distances

In order to quantify the partial similarity $d_{\mathrm{P}}(\mathcal{S}, \mathcal{Q})$, we are looking for the largest and the most similar parts of $\mathcal{S}$ and $\mathcal{Q}$. This translates into the simultaneous minimization of d_{F} and λ on all the possible parts of $\mathcal{S}$ and $\mathcal{Q}$, i.e., a *multicriterion* or *multiobjective* optimization problem (Salukwadze 1979),

$$d_{\mathrm{P}}(\mathcal{S}, \mathcal{Q}) = \min_{(\mathcal{S}', \mathcal{Q}') \in \Sigma_{\mathcal{S}} \times \Sigma_{\mathcal{Q}}} (\lambda(\mathcal{S}', \mathcal{Q}'), d_{\mathrm{F}}(\mathcal{S}', \mathcal{Q}')). \tag{11}$$

The objective function is vector-valued and contains two components: dissimilarity and partiality. It is crucial to realize that the two criteria are competing, and unless the shapes are fully similar, it is impossible to achieve both d_{F} and λ equal to zero.

Visualizing all the possible solutions of the problem as a planar region (Fig. 7), we see that at certain points, we arrive at the situation when by improving one criterion, we inevitably compromise the other, that is, we can obtain a smaller dissimilarity by taking smaller parts, and vice versa. Such solutions are called *Pareto optimal* (Pareto 1906). This notion is closely related to rate-distortion analysis in information theory (de Rooij and Vitanyi 2006) and to receiver operation characteristics (ROC) in pattern recognition (Everson and Fieldsend 2006). In our case, a Pareto optimum is achieved on $(\mathcal{S}^*, \mathcal{Q}^*)$, for which at least one of the following holds,

$$\begin{aligned}&d_{\mathrm{F}}(\mathcal{S}^*, \mathcal{Q}^*) \le d_{\mathrm{F}}(\mathcal{S}', \mathcal{Q}'), \quad \text{or} \\ &\lambda(\mathcal{S}^*, \mathcal{Q}^*) \le \lambda(\mathcal{S}', \mathcal{Q}'),\end{aligned} \tag{12}$$

for all $(\mathcal{S}', \mathcal{Q}') \in \Sigma_{\mathcal{S}} \times \Sigma_{\mathcal{Q}}$. The set of all the Pareto optimal solutions is called the *Pareto frontier* and is denoted in Fig. 7 by a solid curve. Solutions below this curve do not exist.

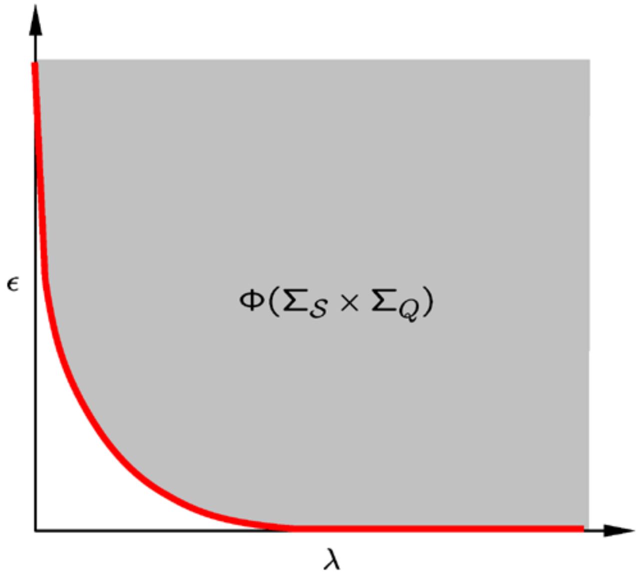

Fig. 7 (Color online) Visualization of the set of all the possible solutions of the multicriterion optimization problem. The Pareto frontier is denoted by *solid curve*

The partial similarity criterion $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q})$ obtained by solving problem (11) can be regarded as a generalized or *set-valued distance*, which is quite different from the traditional concept of similarity. While previously our criterion of similarity was the degree to which $\mathcal{S}$ and $\mathcal{Q}$ are not isometric, we now measure the optimal tradeoff between the isometry and the size of parts of $\mathcal{S}$ and $\mathcal{Q}$.

To formalize this idea, we introduce the notion of (λ,ϵ)-*isometry*. We say that $\mathcal{S}$ and $\mathcal{Q}$ are (λ,ϵ)-isometric if there exist parts $\mathcal{S}' \in \Sigma_{\mathcal{S}}$ and $\mathcal{Q}' \in \Sigma_{\mathcal{Q}}$, such that $\lambda(\mathcal{S},\mathcal{Q}) \le \lambda$ and $(\mathcal{S}', d_{\mathcal{S}'}|_{\mathcal{S}})$ and $(\mathcal{Q}', d_{\mathcal{Q}'}|_{\mathcal{Q}})$ are ϵ-isometric. Our distance can be represented as a non-increasing function of the form $\epsilon(\lambda)$. We will write $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) \le (\lambda_0,\epsilon_0)$, implying that the point (λ_0,ϵ_0) is above or on the graph of the function $\epsilon(\lambda)$; other strong and weak inequalities are defined in the same manner. To say that (λ_0,ϵ_0) is a Pareto optimum, we will write $(\lambda_0,\epsilon_0) \in d_{\mathrm{P}}(\mathcal{S},\mathcal{Q})$.

Given shapes $\mathcal{Q}$ and $\mathcal{S}$, and a full dissimilarity criterion d_{F} satisfying the set of axioms (F1)–(F5) with a constant c, the partial similarity criterion obtained by solving problem (11) satisfies the following properties:

(P1) *Non-negativity*: $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) \subset \mathbb{R}^2_+$.
(P2) *Symmetry*: $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) = d_{\mathrm{P}}(\mathcal{Q},\mathcal{S})$.
(P3) *Monotonicity*: If $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) \le (\lambda,\epsilon)$, then $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) \le (\lambda',\epsilon')$ for all $\lambda' \ge \lambda$ and $\epsilon' \ge \epsilon$.
(P4) *Pareto similarity*: (i) If $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) \le (\lambda,\epsilon)$, then $\mathcal{S}$ and $\mathcal{Q}$ are $(\lambda, c\epsilon)$-isometric; (ii) if $\mathcal{S}$ and $\mathcal{Q}$ are (λ,ϵ)-isometric, then $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q}) \le (\lambda, c\epsilon)$.

4.2 Scalar-Valued Partial Similarity

Though d_{P} encodes much information about the similarity of shapes, their main drawback is the inability to compare similarities. For example, given three shapes, $\mathcal{S}$, $\mathcal{Q}$ and $\mathcal{R}$, we can say that $\mathcal{S}$ is more similar to $\mathcal{Q}$ than $\mathcal{R}$ (which we would normally denote as $d_{\mathrm{P}}(\mathcal{Q},\mathcal{S}) < d_{\mathrm{P}}(\mathcal{Q},\mathcal{R})$) only when all the points of the Pareto frontier $d_{\mathrm{P}}(\mathcal{Q},\mathcal{S})$ are below those of $d_{\mathrm{P}}(\mathcal{Q},\mathcal{R})$. However, the two frontiers may intersect, such that for some values of λ we will have $\mathcal{S}$ more similar to $\mathcal{Q}$, and for some the opposite. This fundamental difference between scalar- and set-values distances stems from the fact that there is no total order relation between vectors.

In order to be able compare partial dissimilarities, we need to convert the set-valued distance into a traditional, scalar-valued one. The easiest way to do so is simply by considering a single point on the Pareto frontier. For example, we can fix the value of $\lambda = \lambda_0$ and use the distortion $\epsilon(\lambda_0)$ as the criterion of partial similarity. Alternatively, we can choose a point by fixing $\epsilon = \epsilon_0$. A scalar distance obtained in this way may be useful in a practical situation when we know *a priori* that the accuracy of geodesic distance measurement or the sampling radius is ϵ_0. This approach resembles the comparison of audio and video encoders, where the rate-distortion curves are compared in the proximity of some *operation point*, specifying either the designated bitrate, or the designated coding quality.

We should note that both of the above choices are rather arbitrary. A slightly more educated selection of a single point out of the set of Pareto optimal solutions was proposed by Salukwadze (1979) in the context of control theory. Salukwadze suggested choosing a Pareto optimum, which is the closest (in sense of some distance) to some optimal, usually non-achievable, *utopia point*. In our case, such an optimal point is $(0,0)$. Given a Pareto frontier $d(\mathcal{S},\mathcal{Q})$, we define the *scalar partial similarity* as

$$d_{\mathrm{S}}(\mathcal{S},\mathcal{Q}) = \inf_{(\lambda,\epsilon)\in d(\mathcal{S},\mathcal{Q})} \|(\lambda,\epsilon)\|_{\mathbb{R}^2_+}. \tag{13}$$

Depending on the choice of the norm $\|\cdot\|_{\mathbb{R}^2_+}$ in (13), we obtain different solutions, some of which have an explicit form. For instance, choosing the weighted L_1-norm, we arrive at the following problem,

$$d_{\mathrm{S}}(\mathcal{S},\mathcal{Q}) = \inf_{(\mathcal{S}',\mathcal{Q}')\in\Sigma_{\mathcal{S}}\times\Sigma_{\mathcal{Q}}} \alpha d_{\mathrm{F}}(\mathcal{S}',\mathcal{Q}') + \beta\lambda(\mathcal{S}',\mathcal{Q}'), \tag{14}$$

where α and β are some positive weights. For the particular choice of the Gromov-Hausdorff distance as d_{F}, in order to make the above expression meaningful in terms of units, α must have units of inverse distance, and β of inverse area. One possible choice is $\alpha = 1/\max\{\mathrm{diam}(\mathcal{S}),\mathrm{diam}(\mathcal{Q})\}$ and $\beta = 1/(\mu_{\mathcal{S}}(\mathcal{S}) + \mu_{\mathcal{Q}}(\mathcal{Q}))$.

4.3 Fuzzy Approximation

In the problems (11) and (14), the optimization was performed over all possible parts of the shapes, $\Sigma_{\mathcal{S}} \times \Sigma_{\mathcal{Q}}$. In the discrete setting, practical numerical solution of such problems is intractable, as the number of possible parts grows exponentially with the number of samples. In order to overcome this problem, we need to find a different way to represent the parts, and formulate the partiality and dissimilarity in these terms.

We begin with the obvious fact that a subset $\mathcal{S}'$ of $\mathcal{S}$ can be described by a *characteristic function*

$$m_{\mathcal{S}}(s) = \begin{cases} 1, & s \in \mathcal{S}', \\ 0, & \text{else}, \end{cases} \tag{15}$$

which indicates whether a point belongs to the subset $\mathcal{S}'$ or not. Using the characteristic functions $(m_{\mathcal{S}}, m_{\mathcal{Q}})$ to represent the parts $(\mathcal{S}', \mathcal{Q}')$ in our problems is still intractable, since the requirement that $m_{\mathcal{S}}$ and $m_{\mathcal{Q}}$ obtain the values of 0 or 1 leads to a combinatorial optimization problem. However, by relaxing this requirement and allowing $m_{\mathcal{S}}$ and $m_{\mathcal{Q}}$ to obtain the values in the entire interval [0, 1], we arrive at a computationally tractable problem, in which the optimization variables are continuous "weights". Sets represented by means of such "weights" are called *fuzzy sets* (Zadeh 1978, 1965; Klir and Yuan 1994; Zimmermann 2001); here, following our terminology, we refer to them as *fuzzy parts*. Parts characterized by binary-valued functions (corresponding to traditional definition of subset) are called *crisp*. The function $m_{\mathcal{S}} : \mathcal{S} \to [0, 1]$ is called a *membership function*.

The fuzzy set theory allows us to formulate a relaxed version of our optimization problem, which, in turn, requires us to extend the definition of the sets of parts, partiality and dissimilarity to the fuzzy setting, in a way consistent with the crisp ones. For this purpose, we make a few definitions. The complement of a fuzzy part is defined as $m^c_{\mathcal{S}} = 1 - m_{\mathcal{S}}$, coinciding with the standard definition on crisp sets. A membership function $m_{\mathcal{S}}$ is called *$\Sigma_{\mathcal{S}}$-measurable* if

$$\{s : m_{\mathcal{S}}(s) \le \delta\} \in \Sigma_{\mathcal{S}}, \tag{16}$$

for all $0 \le \delta \le 1$. We denote by $M_{\mathcal{S}}$ the set of all fuzzy parts of $\mathcal{S}$, defined as the set of all $\Sigma_{\mathcal{S}}$-measurable membership functions on $\mathcal{S}$. $M_{\mathcal{S}}$ replaces $\Sigma_{\mathcal{S}}$ in our relaxed problem.

The *fuzzy measure* is defined as

$$\tilde{\mu}_{\mathcal{S}}(m_{\mathcal{S}}) = \int_{\mathcal{S}} m_{\mathcal{S}}(s) d\mu_{\mathcal{S}}, \tag{17}$$

for all $m_{\mathcal{S}} \in M_{\mathcal{S}}$. For crisp parts, the fuzzy measure $\tilde{\mu}_{\mathcal{S}}$ boils down to the standard measure $\mu_{\mathcal{S}}$. As a matter of notation, we use the tilde to denote fuzzy quantities. We define the *fuzzy partiality* as

$$\tilde{\lambda}(m_{\mathcal{S}}, m_{\mathcal{Q}}) = \tilde{\mu}_{\mathcal{S}}(1 - m_{\mathcal{S}}) + \tilde{\mu}_{\mathcal{Q}}(1 - m_{\mathcal{Q}}), \tag{18}$$

using the fuzzy measure. Since the fuzzy measure coincides with the crisp one on crisp sets, so does the fuzzy partiality.

The definition of a fuzzy dissimilarity depends on the specific choice of d_{F} and may be more elaborate. For the Gromov-Hausdorff distance, it is possible to provide a fuzzy version, based on the following theorem.

Theorem 3 *Let $m_{\mathcal{S}}$ and $m_{\mathcal{Q}}$ be characteristic functions of crisp parts $\mathcal{S}'$ and $\mathcal{Q}'$. Then,*

$$d_{\mathrm{GH}}(\mathcal{S}', \mathcal{Q}') = \frac{1}{2} \inf_{\substack{f:\mathcal{S}\to\mathcal{Q} \\ g:\mathcal{Q}\to\mathcal{S}}} \max \left\{ \begin{array}{l} \sup\limits_{s,s'\in\mathcal{S}} m_{\mathcal{S}}(s) m_{\mathcal{S}}(s') \\ \quad \times |d_{\mathcal{S}}(s,s') - d_{\mathcal{Q}}(f(s), f(s'))| \\ \sup\limits_{q,q'\in\mathcal{Q}} m_{\mathcal{Q}}(q) m_{\mathcal{Q}}(q') \\ \quad \times |d_{\mathcal{Q}}(q,q') - d_{\mathcal{S}}(g(q), g(q'))| \\ \sup\limits_{\substack{s\in\mathcal{S} \\ q\in\mathcal{Q}}} m_{\mathcal{S}}(s) m_{\mathcal{Q}}(q) \\ \quad \times |d_{\mathcal{S}}(s, g(q)) - d_{\mathcal{Q}}(f(s), q)| \\ D \sup\limits_{s\in\mathcal{S}} (1 - m_{\mathcal{Q}}(f(s)))\, m_{\mathcal{S}}(s) \\ D \sup\limits_{q\in\mathcal{Q}} (1 - m_{\mathcal{S}}(g(q)))\, m_{\mathcal{Q}}(q) \end{array} \right\}, \tag{19}$$

where $D \ge \max\{\mathrm{diam}(\mathcal{S}), \mathrm{diam}(\mathcal{Q})\}$.

Employing Theorem 3 with generic membership functions $m_{\mathcal{S}}$ and $m_{\mathcal{Q}}$, it is possible to have a consistent fuzzy generalization of the Gromov-Hausdorff distance, $\tilde{d}_{\mathrm{GH}}(m_{\mathcal{S}}, m_{\mathcal{Q}})$, which by virtue of its definition coincides with the traditional d_{GH} on crisp sets.

Having all the above components, the *fuzzy partial dissimilarity* is computed by solving the relaxed multicriterion optimization problem,

$$\tilde{d}_{\mathrm{P}}(\mathcal{S}, \mathcal{Q}) = \min_{(m_{\mathcal{S}}, m_{\mathcal{Q}}) \in M_{\mathcal{S}} \times M_{\mathcal{Q}}} (\tilde{\lambda}(m_{\mathcal{S}}, m_{\mathcal{Q}}), \tilde{d}_{\mathrm{GH}}(m_{\mathcal{S}}, m_{\mathcal{Q}})). \tag{20}$$

5 Correspondence between Shapes

The last problem we address is the *deformation-invariant correspondence problem*, that is, how to find a map between two shapes that copies similar features to similar features. Implicitly, we have used a semantically vague definition, as the term "similar" is subject to different interpretations. For instance, there is no doubt how a correspondence between a cat and a dog should look like, since both have two ears, four legs and a tail. On the other hand, it would probably be much more difficult to agree about a correspondence between a dog and a bird (Bronstein et al. 2007a).

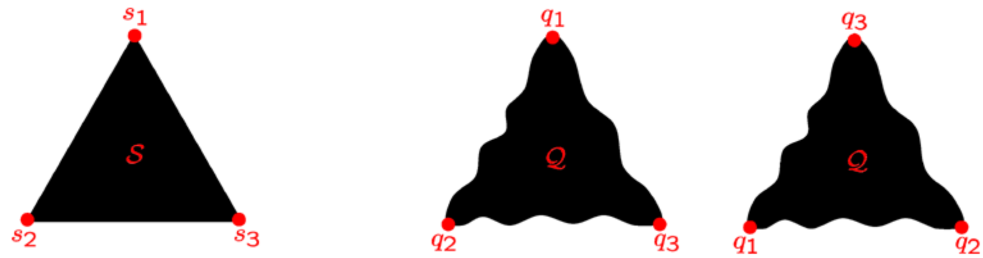

Fig. 8 (Color online) Ambiguity of correspondence in case when the shape has symmetries. Shown are two shapes $\mathcal{S}$ and $\mathcal{Q}$ and two possible correspondences

In our context of non-rigid shape analysis, the correspondence problem can be formulated in geometric terms, as we can use the notions of similarity introduced in Sects. 3 and 4. If two shapes $\mathcal{S}$ and $\mathcal{Q}$ are isometric, there exists a bijective map $f : \mathcal{S} \to \mathcal{Q}$ between them, which established a correspondence between intrinsically similar features. Note that such correspondence is defined up to self-isometries $i \in \mathrm{Iso}(\mathcal{S})$ and $j \in \mathrm{Iso}(\mathcal{Q})$, i.e., f and $j \circ f \circ i$ are both legitimate correspondences. If the shapes have symmetries, the isometry groups are non-trivial and consequently, the correspondence is ambiguous. Yet, most practically interesting shapes have a trivial isometry group, such that an ambiguity of this kind does not arise.

When the shapes $\mathcal{S}$ and $\mathcal{Q}$ are ϵ-isometric, we know that $\mathcal{Q}$ can be produced from $\mathcal{S}$ by means of an ϵ-isometry $f : \mathcal{S} \to \mathcal{Q}$, and, vice versa, $\mathcal{S}$ can be produced from $\mathcal{Q}$ by an ϵ-isometry $g : \mathcal{Q} \to \mathcal{S}$. We can say that for every s in $\mathcal{S}$, the corresponding point in $\mathcal{Q}$ is $f(s)$, and for every q in $\mathcal{Q}$, the corresponding point in $\mathcal{S}$ is $g(q)$. These correspondences can be found by solving

$$(f^*, g^*) = \arg \min_{\substack{f:\mathcal{S}\to\mathcal{Q} \\ g:\mathcal{Q}\to\mathcal{S}}} \max\{\mathrm{dis}\, f, \mathrm{dis}\, g, \mathrm{dis}(f, g)\}, \tag{21}$$

which can be thought of as a byproduct of the computation of the Gromov-Hausdorff distance $d_{\mathrm{GH}}(\mathcal{Q}, \mathcal{S})$. (Here, we tacitly assume that we can write minimum instead of infimum, which is not necessarily true in the continuous case. However, since in practice we work with discrete shapes consisting of a finite number of samples, the minimum is always achieved, therefore, we allow ourselves this relaxed notation.) It is guaranteed that $\mathrm{dis}\, f^*, \mathrm{dis}\, g^* \leq \epsilon$, and that both f^* and g^* are ϵ-surjective. Each of the maps f^*, g^* serves as the *minimum-distortion correspondence* (Bronstein et al. 2007a, 2006d). Since the correspondence is defined up to self-isometries, instead of f^* we may have $j \circ f^* \circ i$, and instead of g^* we may have $i \circ g^* \circ j$ (see Fig. 8).

5.1 Partial Correspondence

The minimum-distortion correspondence (21) matches the features of the entire shape $\mathcal{S}$ with similar features of the entire shape $\mathcal{Q}$, and can be therefore termed as *full correspondence*. Clearly, full correspondence is not applicable when $\mathcal{S}$ and $\mathcal{Q}$ are related by the partial similarity relation. In the latter case, we would like to establish a *partial correspondence*, relating the features of a part of $\mathcal{S}$ to similar features of a part of $\mathcal{Q}$.

Using the partial similarity from Sect. 4, we may define the *partial correspondence* between $\mathcal{S}$ and $\mathcal{Q}$ as the map between the parts $(\mathcal{S}', \mathcal{Q}') \in \Sigma_{\mathcal{S}} \times \Sigma_{\mathcal{Q}}$, which is obtained by solving

$$(f^*, g^*) = \arg \min_{\substack{f:\mathcal{S}'\to\mathcal{Q}' \\ g:\mathcal{Q}'\to\mathcal{S}' \\ (\mathcal{S}', \mathcal{Q}')\in\Sigma_{\mathcal{S}}\times\Sigma_{\mathcal{Q}}}} \Big(\lambda(\mathcal{S}', \mathcal{Q}'), \frac{1}{2} \max\{\mathrm{dis}\, f, \mathrm{dis}\, g, \mathrm{dis}(f, g)\}\Big), \tag{22}$$

where the minimum is interpreted in the Pareto sense. The correspondence (f^*, g^*) is obtained between the parts $(\mathcal{S}^*, \mathcal{Q}^*)$, and can be considered as a byproduct of problem (11).

Similarly, in the fuzzy setting, the correspondence is obtained by solving

$$(f^*, g^*) = \arg \min_{\substack{f:\mathcal{S}\to\mathcal{Q} \\ g:\mathcal{Q}\to\mathcal{S} \\ (m_{\mathcal{S}}, m_{\mathcal{Q}})\in M_{\mathcal{S}}\times M_{\mathcal{Q}}}} \left(\tilde{\lambda}(m_{\mathcal{S}}, m_{\mathcal{Q}}), \frac{1}{2} \max \left\{ \begin{array}{c} \sup\limits_{s,s'\in\mathcal{S}} m_{\mathcal{S}}(s) m_{\mathcal{S}}(s') |d_{\mathcal{S}}(s, s') - d_{\mathcal{Q}}(f(s), f(s'))| \\ \sup\limits_{q,q'\in\mathcal{Q}} m_{\mathcal{Q}}(q) m_{\mathcal{Q}}(q') |d_{\mathcal{Q}}(q, q') - d_{\mathcal{S}}(g(q), g(q'))| \\ \sup\limits_{\substack{s\in\mathcal{S} \\ q\in\mathcal{Q}}} m_{\mathcal{S}}(s) m_{\mathcal{Q}}(q) |d_{\mathcal{S}}(s, g(q)) - d_{\mathcal{Q}}(f(s), q)| \\ D \sup\limits_{s\in\mathcal{S}} (1 - m_{\mathcal{Q}}(f(s)))\, m_{\mathcal{S}}(s) \\ D \sup\limits_{q\in\mathcal{Q}} (1 - m_{\mathcal{S}}(g(q)))\, m_{\mathcal{Q}}(q) \end{array} \right\} \right), \tag{23}$$

i.e., as a byproduct of problem (20), in which the optimal fuzzy parts $(m^*_{\mathcal{S}}, m^*_{\mathcal{Q}})$ are also found. Note that here, unlike the crisp case, (f^*, g^*) are maps between the entire shapes $\mathcal{S}$ and $\mathcal{Q}$. Thresholding $(m^*_{\mathcal{S}}, m^*_{\mathcal{Q}})$ at some level $0 \le \delta \le 1$, we convert the fuzzy parts into crisp ones,

$$\begin{aligned} \mathcal{S}^*_\delta &= \{s \in \mathcal{S} : m^*_{\mathcal{S}}(s) \ge \delta\}, \\ \mathcal{Q}^*_\delta &= \{q \in \mathcal{Q} : m^*_{\mathcal{Q}}(q) \ge \delta\}, \end{aligned} \tag{24}$$

and define the δ-*partial correspondences* $f^*_\delta : \mathcal{S}^*_\delta \to \mathcal{Q}^*_\delta$ and $g^*_\delta : \mathcal{Q}^*_\delta \to \mathcal{S}^*_\delta$ as f^* and g^* restricted to $\mathcal{S}^*_\delta$ and $\mathcal{Q}^*_\delta$, respectively. The partial correspondence is the collection of (f^*_δ, g^*_δ) for all $0 \le \delta \le 1$.

6 Numerical Framework

6.1 Discretization

The discretization of a shape $\mathcal{S}$ involves three components: discretization of the set $\mathcal{S}$ itself, the metric $d_{\mathcal{S}}$ and the measure $\mu_{\mathcal{S}}$. The set $\mathcal{S} \subset \mathbb{R}^2$ is represented as a finite r-sampling $\mathcal{S}_N = \{s_1, \dots, s_N\}$. Triangulating the points s_i in the plane, we obtain a flat polyhedral (first-order) approximation of $\mathcal{S}$. Representing $\mathcal{S}$ as a triangular mesh $T(\mathcal{S}_N)$ allows us to work with a finite discrete set of points on one hand, while preserving the continuous nature of the set $\mathcal{S}$ on the other.

The metric on $\mathcal{S}$ is discretized by numerically approximating the geodesic distances between the samples s_i on the triangular mesh $T(\mathcal{S}_N)$. For this purpose, we use the *fast marching method* (FMM) (Sethian 1996; Kimmel and Sethian 1998). The distances are arranged into an $N \times N$ matrix denoted by $\mathbf{D}_{\mathcal{S}} = (d_{\mathcal{S}}(s_i, s_j))$. Fast marching computes the matrix $\mathbf{D}_{\mathcal{S}}$ in $\mathcal{O}(N^2 \log N)$; parametric versions of FMM (Spira and Kimmel 2004) can work in $\mathcal{O}(N^2)$ and are highly-parallelizable with only a slight degradation of accuracy (Bronstein et al. 2007c).

The measure on $\mathcal{S}$ is discretized by constructing a discrete measure $\mu_{\mathcal{S}_N} = (\mu_1, \dots, \mu_N\}$, assigning to each s_i on $\mathcal{S}_N$ the area of the corresponding Voronoi cell. In practice, when the sampling is sufficiently uniform, selecting $\mu_i = 1/N$ constitutes a reasonable approximation.

6.2 Generalized Multidimensional Scaling

The basic computation involved in the problems we defined is finding a minimum-distortion embedding of a shape $\mathcal{S}$ into $\mathcal{Q}$. In order to avoid optimization over all the maps $f : \mathcal{S} \to \mathcal{Q}$ (which is untractable in practice), we minimize over the images $q'_i = f(s_i)$, where q'_i are represented in continuous coordinates on the triangular mesh $T(\mathcal{Q}_M)$,

$$\min_{q'_1,\dots,q'_N \in T(\mathcal{Q}_M)} \max_{j>i} |d_{\mathcal{Q}}(q'_i, q'_j) - d_{\mathcal{S}}(s_i, s_j)|. \tag{25}$$

Optimization problem (25) is similar in its spirit to multidimensional scaling (MDS), and is referred to as the *generalized MDS* (GMDS) problem (Bronstein et al. 2006a, 2006b). It can be reformulated as the following constrained minimization,

$$\begin{aligned} &\min_{\epsilon \ge 0, q'_1,\dots,q'_N \in T(\mathcal{Q}_M)} \epsilon \\ &\text{s.t.} \quad |d_{\mathcal{Q}}(q'_i, q'_j) - d_{\mathcal{S}}(s_i, s_j)| \le \epsilon, \end{aligned} \tag{26}$$

for $i, j = 1, \dots, N$, with $N + 1$ variables and $2N^2$ inequality constraints. An alternative approach adopted here, is to replace the min-max problem by a weighted least-squares formulation,

$$\min_{q'_1,\dots,q'_N \in T(\mathcal{Q}_M)} \sum_{j>i} w_{ij} \cdot \mu_i \mu_j (d_{\mathcal{Q}}(q'_i, q'_j) - d_{\mathcal{S}}(s_i, s_j))^2, \tag{27}$$

where $\{w_{ij}\}$ is a set of non-negative weights. In the sequel, we show an iterative reweighting scheme, which allows to approximate the solution of the GMDS problem in its original L_∞ formulation.

The main distinction of GMDS from the traditional MDS problem (Borg and Groenen 1997) is the fact that the geodesic distances in the target space $\mathcal{Q}$ have no analytic expression. We have the numerically approximated geodesic distances $\mathbf{D}_{\mathcal{Q}}$, but since the q'_i usually fall inside the triangular faces of the mesh $T(\mathcal{Q}_M)$, one has to compute the geodesic distances $d_{\mathcal{Q}}$ between any two arbitrary points on $T(\mathcal{Q}_M)$. For this purpose, we use a variation of the *three-point geodesic distance approximation*, proposed in (Bronstein et al. 2006b). Let us assume without loss of generality that we need to approximate the geodesic distance $d_{\mathcal{Q}}(q'_1, q'_2)$, where q'_1 and q'_2 are two points on the mesh $T(\mathcal{Q}_M)$. Let us furthermore assume that q'_1 and q'_2 are located on the faces t_1 and t_2, whose vertices are $q_{t_1,1}, q_{t_1,2}, q_{t_1,3}$ and $q_{t_2,1}, q_{t_2,2}, q_{t_2,3}$, respectively. The location of q'_i on the mesh can be unequivocally described by the index t_i of the enclosing triangle, and the position inside the triangle. The latter can be expressed as the convex combination

$$q'_i = u_i q_{t_i,1} + v_i q_{t_i,2} + (1 - u_i - v_i) q_{t_i,3} \tag{28}$$

of the triangle vertices, where the pair of non-negative coefficients (u_i, v_i) satisfying $u_i + v_i = 1$ is referred to as the *barycentric coordinates* of q'_i. In what follows, we will switch freely between q'_i and its barycentric representation (t_i, u_i, v_i).

We first approximate the three distances $d_{\mathcal{Q}}(q'_1, q_{t_2,1})$, $d_{\mathcal{Q}}(q'_1, q_{t_2,2})$, and $d_{\mathcal{Q}}(q'_1, q_{t_2,3})$ using linear interpolation in the triangle t_1,

$$\begin{aligned} d_{\mathcal{Q}}(q'_1, q_{t_2,k}) &= u_1\, d_{\mathcal{Q}}(q_{t_1,1}, q_{t_2,k}) + v_1\, d_{\mathcal{Q}}(q_{t_1,2}, q_{t_2,k}) \\ &\quad + (1 - u_1 - v_1)\, d_{\mathcal{Q}}(q_{t_1,3}, q_{t_2,k}). \end{aligned} \tag{29}$$

Note that all geodesic distance terms in the above expression are between fixed vertices of the mesh $T(\mathcal{Q}_M)$, and can be therefore pre-computed or computed on demand and cached. Thus, the evaluation of $d_{\mathcal{Q}}(q_1', q_i)$ has constant complexity independent of the sample size M. The linear interpolation step is repeated again, this time in the triangle t_2, yielding

$$d_{\mathcal{Q}}(q_1', q_2') = u_2\, d_{\mathcal{Q}}(q_1', q_{t_2,1}) + v_2\, d_{\mathcal{Q}}(q_1', q_{t_2,2}) + (1 - u_2 - v_2)\, d_{\mathcal{Q}}(q_1', q_{t_2,3}). \quad (30)$$

The first-order derivatives of $d_{\mathcal{Q}}(q_1', q_2')$ with respect to the coordinates (u_1, v_1) and (u_2, v_2) are evaluated in a similar manner.

Plugging the former result into (27), we observe that the cost function

$$\sigma(u_1, \ldots, u_N, v_1, \ldots, v_N, t_1, \ldots, t_N) = \sum_{j>i} w_{ij} \cdot \mu_i \mu_j (d_{\mathcal{Q}}(q_i', q_j') - d_{\mathcal{S}}(s_i, s_j))^2 \quad (31)$$

is a fourth-order polynomial containing second-order terms of u_i and v_i. Consequently, considering σ as a function of a pair (u_i, v_i) and fixing the rest of the optimization variables, results in a convex quadratic function with respect to (u_i, v_i), whose minimum (u_i^*, v_i^*) can be found analytically by solving the 2×2 Newton system

$$H_i(u_i, v_i)(u_i^* - u_i, v_i^* - v_i)^{\mathrm{T}} = -g_i(u_i, v_i), \quad (32)$$

where g_i and H_i are, respectively, the gradient and the Hessian of σ with respect to u_i and v_i. However, the solution (u_i^*, v_i^*) obtained this way may not be a valid pair of barycentric coordinates, meaning that the point q_i' may be displaced outside the triangle t_i. To disallow such a case, we find analytically the solution to the *constrained* quadratic problem

$$\min_{u_i \ge 0, v_i \ge 0} \sigma(u_i, v_i) \quad \text{s.t.} \quad u_i + v_i \le 1. \quad (33)$$

If (u_i^*, v_i^*) happens to lie on an edge or a vertex of the triangle t_i (that is, at least one constrain is active), the need to update the triangle index t_i may arise. If (u_i^*, v_i^*) lies on a triangle edge shared with some other triangle t_i', we translate the barycentric representation (u_i^*, v_i^*) in the coordinate system of t_i to (u_i', v_i') with respect to the coordinate system of t_i'. This translation does not change the value of σ, yet, as the cost function is not $\mathcal{C}^1$ at on the triangle boundaries, the gradient direction may change. We evaluate the new gradient direction in the triangle t_i', and update t_i to be t_i' only if the negative gradient direction points inside t_i'. In this case, subsequent minimization of σ with respect to the updated (u_i, v_i) will guarantee cost decrease. If the triangle edge is not shared with another triangle (i.e., the edge is part of the shape boundary), no index update is performed. A similar procedure is performed in the case where (u_i^*, v_i^*) lies on a triangle vertex.

1. **for** $k = 0, 1, 2, \ldots$ **do**
2. Evaluate the gradients g_i and the Hessian matrices H_i of the cost function $\sigma(u_i^{(k)}, v_i^{(k)}, t_i^{(k)})$ with respect to the variables $u_i^{(k)}$ and $v_i^{(k)}$.
3. Select i corresponding to $\max \|g_i\|$.
4. **if** $\|g_i\|$ is sufficiently small **then** Stop
5. Solve the constrained quadratic problem
$$(u_i^*, v_i^*) = \arg \min_{u_i \ge 0, v_i \ge 0} \sigma(u_i, v_i) \text{ s.t. } u_i + v_i \le 1$$
with the rest of u_j and v_j fixed to $u_j^{(k)}$ and $v_j^{(k)}$.
6. **if** (u_i^*, v_i^*) is on an edge of t_i **then**
7. Set T' to be the set containing the triangle sharing the edge with t_i, or $\emptyset$ in case the edge is on the shape boundary.
8. **else if** (u_i^*, v_i^*) is on a vertex of t_i **then**
9. Set T' to be the list of triangles sharing the vertex with t_i.
10. **else** Set $T' = \emptyset$
11. **forall** $t' \in T'$ **do**
12. Translate (u_i^*, v_i^*) to the coordinates of the triangle t'.
13. Evaluate the gradient g_i of σ at (u_i^*, v_i^*) in t'.
14. **if** $-g_i$ is directed inside the triangle t' **then**
15. Update $t_i^{(k+1)} = t'$.
16. Go to Step 19.
17. **end**
18. **end**
19. Update $(u^{(k+1)}, v^{(k+1)}) = (u^*, v^*)$.
20. **end**

Algorithm 2 Weighted least squares GMDS

The entire minimization procedure is summarized in Algorithm 2. The described minimization algorithm can be viewed as a block-coordinate descent, where at each iteration the block of two coordinates corresponding to the point with the largest gradient is selected (Step 3). The constrained Newton descent performed in Step 7 guarantees monotonicity of the sequence of values of σ produced by the algorithm.

6.3 Iteratively Reweighted Least Squares

The proposed weighted least squares minimization procedure can be employed for solving GMDS problems with arbitrary norms. Let us consider a cost function of the form

$$\sigma_\rho(q_1', \ldots, q_N') = \sum_{i>j} \mu_i \mu_j \rho(d_{\mathcal{Q}}(q_i', q_j') - d_{\mathcal{S}}(s_i, s_j)), \quad (34)$$

where $\rho(t)$ is some norm. For example, setting $\rho(t) = |t|^p$ gives the L_p norm, with L_∞ in the limit $p \to \infty$. Other *robust norms* are preferable in practical applications, with the notable examples of the German-McLure function

$$\rho(t) = \frac{t^2}{t^2 + \epsilon^2}, \tag{35}$$

and the quadratic-linear Huber function

$$\rho(t) = \begin{cases} \frac{t^2}{2\epsilon}, & |t| \le \epsilon, \\ |t| - 0.5\epsilon, & |t| > \epsilon, \end{cases} \tag{36}$$

where ϵ is a positive constant. These norms exhibit good properties in the presence of noise.

The necessary condition for $q_1'^*, \ldots, q_N'^*$ to be a local minimizer of σ_ρ is

$$\nabla\sigma_\rho(q_1'^*, \ldots, q_N'^*) = \sum_{i>j} \mu_i \mu_j \rho'(d_{\mathcal{Q}}(q_i'^*, q_j'^*) - d_{\mathcal{S}}(s_i, s_j))\nabla d_{\mathcal{Q}}(q_i'^*, q_j'^*) = 0.$$

Instead of minimizing σ_ρ, we can solve the weighted least squares problem (27), whose solution has to satisfy

$$\nabla\sigma(q_1'^*, \ldots, q_N'^*) = \sum_{i>j} 2w_{ij}\mu_i \mu_j (d_{\mathcal{Q}}(q_i'^*, q_j'^*) - d_{\mathcal{S}}(s_i, s_j))\nabla d_{\mathcal{Q}}(q_i'^*, q_j'^*) = 0.$$

If we could select the weights in $\sigma(q_1'^*, \ldots, q_N'^*)$ according to

$$w_{ij} = \frac{\rho'(d_{\mathcal{Q}}(q_i'^*, q_j'^*) - d_{\mathcal{S}}(s_i, s_j))}{2(d_{\mathcal{Q}}(q_i'^*, q_j'^*) - d_{\mathcal{S}}(s_i, s_j))}, \tag{37}$$

the two minimizers would coincide and we could reduce the minimization of σ_ρ to the solution of the weighted least squares problem. However, such a selection of the weights requires the knowledge of the minimizer of σ_ρ, which is, of course, unknown. A possible remedy is to start by solving the uniformly weighted least squares problem (all $w_{ij} = 1$), use the solution to update the weights, and iterate the process until convergence. Such *iteratively reweighted least squares* (IRLS) techniques are often used in statistics to approximate the solution of problems with robust norms (Hampel et al. 1986; Forsyth and Ponce 2003; Huber 2004) and in computer vision (Black and Anandan 1993).

6.4 Multiresolution Optimization

Despite the fact that the GMDS problem is convex with respect to each pair of coordinates (u_i, v_i), like the traditional MDS, it is non-convex with respect to all the minimization variables together. Therefore, it is prone to converge to local minima rather than to the global one (Borg and Groenen 1997). Nevertheless, convex optimization is widely used in the MDS community if some precautions are taken in order to prevent local convergence. Here, we use a multiresolution optimization scheme that in practical applications shows good global convergence (Bronstein et al. 2006b, 2006e).

The key idea of a multiresolution optimization scheme is to work with a hierarchy of problems, starting from a coarse version of the problem containing a small number of variables (points). The coarse level solution is interpolated to the next resolution level, and is used as an initialization for the optimization at that level. The process is repeated until the finest level solution is obtained. Such a multi-scale scheme can be thought of as a smart way of initializing the optimization problem. Small local minima tend to disappear at coarse resolution levels, thus reducing the risk of local convergence which is more probable when working at a single resolution.

The main components of a multiresolution scheme are the hierarchy of data which defines optimization problems at different resolution levels, and the interpolation procedure, which allows to pass from coarse level to a finer one. Such a data hierarchy can be constructed using the *holographic sampling* (Bruckstein et al. 1998) or the *farthest point sampling* (FPS) strategies (Eldar et al. 1997). For passing from one resolution level to another we use the *geodesic interpolation* technique, detailed in Bronstein et al. (2006b).

6.5 Initialization

Though the multiresolution scheme reduces the probability of local convergence, in order that the solutions at finer resolution levels be in the basin of attraction of the global minimum, the coarse resolution problem has to be initialized sufficiently close to it. Given $\mathcal{S}_N$ and $\mathcal{Q}_M$ sampled with the radius r, we can sub-sample them with a larger radius R, producing sparser sampling $\mathcal{S}_{N'} \subset \mathcal{S}_N$ and $\mathcal{Q}_{M'} \subset \mathcal{Q}_M$ containing $N' \ll N$ and $M' \ll M$ points, respectively. We denote by $\mathbb{F}$ the space of all discrete mappings $\pi : \mathcal{S}_{N'} \to \mathcal{Q}_{M'}$, which can be represented as a correspondence between N' indices, $(1, \ldots, N') \mapsto (\pi_1, \ldots, \pi_{N'})$, $\pi_i \in \{1, \ldots, M'\}$. A mapping π with the minimum distortion is an approximation to the global minimum of the GMDS problem, and as such, it is a good candidate for coarse resolution initialization. Unfortunately, the space $\mathbb{F}$ is very large even for modest sample sizes, containing $M'^{N'}$ mappings, and exhaustively searching for the best mapping in it is impractical. However, the search space can be significantly reduced by ruling out mappings that are unlikely to have low distortion.

We observe that in order for π to be a good candidate for a global minimum, the intrinsic properties of the shape $\mathcal{S}$, such as the behavior of the metric $d_{\mathcal{S}}$ around every s_i should be similar to those of $\mathcal{Q}$ around q_{π_i}. In order to quantify this behavior, for each $s_i \in \mathcal{S}_{N'}$ we compute the histogram $h(s_i) = \text{hist}(\{d_{\mathcal{S}}(s_i, s_j) : d_{\mathcal{S}}(s_i, s_j) \le R'\})$ of the geodesic distances in a R'-ball centered s_i. In the same manner, the set of histograms $h(q_i)$ is computed in $\mathcal{Q}_{M'}$. Using the vectors $h(s_i)$ and $h(q_j)$ as local descriptors of the points in $\mathcal{S}_{N'}$ and $\mathcal{Q}_{M'}$, respectively, we compute the dissimilarity of two points $s_i \in \mathcal{S}_{N'}, q_j \in \mathcal{Q}_{M'}$ as the Euclidean distance $\|h(s_i) - h(q_j)\|_2$ between their descriptors. For each point s_i in $\mathcal{S}_{N'}$, we construct a set $C_i \subset \{1, \ldots, M'\}$ of indices of K points in $\mathcal{Q}_{M'}$ having the most similar descriptors. K is selected to be a small number, typically significantly smaller than N'. We define the reduced search space $\mathbb{F}_{\text{init}} = C_1 \times C_2 \times \cdots \times C_{N'}$. Mappings copying any s_i to q_{π_i} with $\pi_i \notin C_i$ are excluded from the search space.

Even though the coarse sample sizes N' and M' and the number of initial matches for every point are relatively small, $\mathbb{F}_{\text{init}}$ has still $\mathcal{O}(K^{N'})$ mappings, making an exhaustive search prohibitively expensive. However, adopting the spirit of Gelfand et al. (2005), we can use the following hierarchical greedy algorithm for selecting a reasonably good mapping from $\mathbb{F}_{\text{init}}$.

1. **Pairing:** For each pair $(i, j) \in \{1, \ldots, M\}^2$, choose the best pair $(m, n) \in C_i \times C_j$ minimizing the distortion $\left|d_{\mathcal{S}}(s_i, s_j) - d_{\mathcal{Q}}(q_{\pi_i}, q_{\pi_j})\right|$. This establishes a two-point correspondence $(i, j) \mapsto (m, n)$. The outcome of this step is the set of $\mathcal{O}(N'^2)$ two-point correspondences E_2, which we sort in increasing order of distortion.
2. **Merging:** The pairs are merged into four-point correspondences. Taking the first two-point correspondence $e \in E_2$, we find another two-point correspondence having a disjoint domain and minimizing the distortion of the obtained four-point correspondence. We remove all correspondences sharing the same domain from E_2 and continue until E_2 becomes empty. The merging continues hierarchically, producing E_{2k} from E_k, stopping typicall at E_8 or E_{16}.
3. **Completion:** We select the minimum distortion correspondence $(i_1, \ldots, i_k) \mapsto (\pi_{i_1}, \ldots, \pi_{i_k})$ from the last produced E_k, and complete it to a full N'-point correspondence by adding the missing indices $\{i_{k+1}, \ldots, i_{N'}\} = \{1, \ldots, N'\} \setminus \{i_1, \ldots, i_k\}$ and their images $\pi_{i_{k+1}}, \ldots, \pi_{i_{N'}}$. For each added point j, we select
$$\pi_j = \arg \min_{\pi_j \in \{1,\ldots,M'\}} \max_{i \in \{i_1,\ldots,i_k\}} |d_{\mathcal{S}}(s_i, s_j) - d_{\mathcal{Q}}(q_{\pi_i}, q_{\pi_j})|.$$
The returned results are the mapping π and its distortion $\epsilon_{\min}$.

Since the algorithm never backtracks, it may produce a suboptimal mapping π. However, practice shows that if some good pairs are found at Step 1, the algorithm tends to produce a very good estimate for the minimum distortion mapping on $\mathbb{F}_{\text{init}}$.

A guaranteed global minimum on $\mathbb{F}_{\text{init}}$ can be computed by using a branch and bound algorithm similar in spirit to that presented in Gelfand et al. (2005) for improving convergence of iterative closest point-based extrinsic surface alignment. The idea of the algorithm is based on the fact that if a good estimate for π is found using the greedy matching, a large set of mappings in $\mathbb{F}_{\text{init}}$ can be further eliminated efficiently. The algorithm is initialized by $\pi_{\min}$ and $\epsilon_{\min}$ found by greedy matching, and proceeds as follows:

1. Given a correspondence of $k-1$ feature points $(1, \ldots, k-1) \mapsto (\pi_1, \ldots, \pi_{k-1})$, we would like to establish $k \mapsto \pi_k$.
2. **Prune:** For each potential correspondence $\pi_k \in C_k$, evaluate
$$\max_{i=1,\ldots,k} |d_{\mathcal{S}}(s_i, s_k) - d_{\mathcal{Q}}(q_{\pi_i}, q_{\pi_k})|.$$
If the obtained distortion is larger than $\epsilon_{\min}$, discard the potential correspondence.
3. **Branch:** For each remaining π_k, recursively invoke Step 1 with $(1, \ldots, k) \mapsto (\pi_1, \ldots, \pi_k)$.
4. **Bound:** If $k = N'$, compute the distortion $\text{dis}(\pi)$. If $\text{dis}(\pi) < \epsilon_{\min}$, set $\epsilon_{\min} = \text{dis}(\pi)$ and $\pi_{\min} = \pi$.

6.6 Computation of d_{GH} and the Full Correspondence

So far, our focus was on finding the minimum distortion embedding of a shape $\mathcal{S}$ into another shape $\mathcal{Q}$. However, the GMDS framework can be straightforwardly adapted for computation of the Gromov-Hausdorff distance between two shapes. In fact, Definition 7 suggests that d_{GH} can be formulated as two minimum-distortion embedding problems, coupled together by the third distortion term $\text{dis}(f, g)$:

$$d_{\text{GH}}(\mathcal{Q}_M, \mathcal{S}_N) = \frac{1}{2} \min_{\substack{q'_1,\ldots,q'_N \in T(\mathcal{Q}_M) \\ s'_1,\ldots,s'_M \in T(\mathcal{S}_N)}} \max_{\substack{i,j=1,\ldots,N \\ k,l=1,\ldots,M}} \max \left\{ \begin{array}{l} |d_{\mathcal{S}}(s_i, s_j) - d_{\mathcal{Q}}(q'_i, q'_j)|, \\ |d_{\mathcal{Q}}(q_k, q_l) - d_{\mathcal{S}}(s'_k, s'_l)|, \\ |d_{\mathcal{S}}(s_i, s'_k) - d_{\mathcal{Q}}(q_k, q'_i)| \end{array} \right\}, \quad (38)$$

where the minimization is performed over two sets of continuous variables $q'_i = f(s_i)$, and $s'_k = g(q_k)$. This problem, in turn, can be cast as the following constrained minimization problem

$$d_{\text{GH}}(\mathcal{Q}_M, \mathcal{S}_N) = \min_{\substack{\epsilon \geq 0 \\ q'_1,\ldots,q'_N \in T(\mathcal{Q}_M) \\ s'_1,\ldots,s'_M \in T(\mathcal{S}_N)}} \frac{\epsilon}{2} \quad \text{s.t.} \left\{ \begin{array}{l} |d_{\mathcal{S}}(s_i, s_j) - d_{\mathcal{Q}}(q'_i, q'_j)| \leq \epsilon, \\ |d_{\mathcal{Q}}(q_k, q_l) - d_{\mathcal{S}}(s'_k, s'_l)| \leq \epsilon, \\ |d_{\mathcal{S}}(s_i, s'_k) - d_{\mathcal{Q}}(q_k, q'_i)| \leq \epsilon, \end{array} \right. \quad (39)$$

with $M+N+1$ variables and $2(M^2+N^2+MN)$ inequality constraints. The maps $s_i \mapsto q_i'$ and $q_k \mapsto s_k'$ at the minimum define the minimum-distortion full correspondence between $T(\mathcal{S}_N)$ and $T(\mathcal{Q}_M)$. Alternatively, one can resort to the weighted least squares formulation

$$\min_{\substack{q_1',\ldots,q_N' \in T(\mathcal{Q}_M)\\ s_1',\ldots,s_M' \in T(\mathcal{S}_N)}} \sum_{j>i} \alpha_{ij} \cdot \mu_i \mu_j (d_{\mathcal{S}}(s_i,s_j) - d_{\mathcal{Q}}(q_i',q_j'))^2 + \sum_{l>k} \beta_{kl} \cdot \nu_k \nu_l (d_{\mathcal{Q}}(q_k,q_l) - d_{\mathcal{S}}(s_k',s_l'))^2 + \sum_{i,k} \gamma_{ik} \cdot \mu_i \nu_k (d_{\mathcal{S}}(s_i,s_k') - d_{\mathcal{Q}}(q_k,q_i'))^2, \tag{40}$$

where $\nu = \{\nu_1,\ldots,\nu_M\}$ denotes the discretized measure of $\mathcal{Q}$, and $\{\alpha_{ij}\}$, $\{\beta_{kl}\}$, and $\{\gamma_{ik}\}$ are sets of non-negative weights. Using iterative reweighting, the GMDS problem can be solved with an arbitrary norm.

6.7 Computation of $\tilde{d}_{\mathrm{P}}$ and the Partial Correspondence

In the discrete version of problem (20), the membership functions $m_{\mathcal{S}}$ and $m_{\mathcal{Q}}$ are replaced by vectors $m_{\mathcal{S}_N} = (m_{\mathcal{S}}(s_1),\ldots,m_{\mathcal{S}}(s_N))$ and $m_{\mathcal{Q}_M} = m_{\mathcal{Q}}(q_1),\ldots,m_{\mathcal{Q}}(q_M))$. The fuzzy partiality $\tilde{\lambda}(m_{\mathcal{S}}, m_{\mathcal{Q}})$ is discretized as

$$\tilde{\lambda}(m_{\mathcal{S}_N}, m_{\mathcal{Q}_M}) = m_{\mathcal{S}_N}^{\mathrm{T}} \mu_{\mathcal{S}_N} + m_{\mathcal{Q}_M}^{\mathrm{T}} \mu_{\mathcal{Q}_M}. \tag{41}$$

The computation of $\tilde{d}_{\mathrm{P}}(\mathcal{S}_N, \mathcal{Q}_M)$ is performed by computing a finite set of points on the Pareto frontier, by fixing a value of λ and computing the corresponding dissimilarity, which can be posed as the following optimization problem,

$$\min_{\substack{\epsilon \geq 0\\ m_{\mathcal{S}_N}, m_{\mathcal{Q}_M}\\ q_1',\ldots,q_N' \in T(\mathcal{Q}_M)\\ s_1',\ldots,s_M' \in T(\mathcal{S}_N)}} \epsilon \quad \text{s.t.} \begin{cases} m_{\mathcal{S}_N}(s_i) m_{\mathcal{S}_N}(s_j) |d_{\mathcal{S}}(s_i,s_j) - d_{\mathcal{Q}}(q_i',q_k')| \leq \epsilon, \\ m_{\mathcal{Q}_M}(q_k) m_{\mathcal{Q}_M}(q_l) |d_{\mathcal{Q}}(q_k,q_l) - d_{\mathcal{S}}(s_k',s_l')| \leq \epsilon, \\ m_{\mathcal{S}_N}(s_i) m_{\mathcal{Q}_M}(q_k) |d_{\mathcal{S}}(s_i,s_k') - d_{\mathcal{Q}}(q_k,q_i')| \leq \epsilon, \\ D(1 - m_{\mathcal{Q}_M}(q_i')) m_{\mathcal{S}_N}(s_i) \leq \epsilon, \\ D(1 - m_{\mathcal{S}_N}(s_k')) m_{\mathcal{Q}}(q_k) \leq \epsilon, \\ m_{\mathcal{S}_N}^{\mathrm{T}} \mu_{\mathcal{S}_N} + m_{\mathcal{Q}_M}^{\mathrm{T}} \mu_{\mathcal{Q}_M} \geq 1 - \lambda. \end{cases} \tag{42}$$

If we assume that $m_{\mathcal{S}_N}, m_{\mathcal{Q}_M}$ in problem (42) are fixed, we can compute $\tilde{d}_{\mathrm{GH}}(m_{\mathcal{S}_N}, m_{\mathcal{Q}_M})$ in a manner similar to the Gromov-Hausdorff distance computation using a GMDS-like numerical scheme,

$$\min_{\substack{\epsilon \geq 0\\ q_1',\ldots,q_N' \in T(\mathcal{Q}_M)\\ s_1',\ldots,s_M' \in T(\mathcal{S}_N)}} \epsilon \quad \text{s.t.} \begin{cases} m_{\mathcal{S}_N}(s_i) m_{\mathcal{S}_N}(s_j) |d_{\mathcal{S}}(s_i,s_j) - d_{\mathcal{Q}}(q_i',q_k')| \leq \epsilon, \\ m_{\mathcal{Q}_M}(q_k) m_{\mathcal{Q}_M}(q_l) |d_{\mathcal{Q}}(q_k,q_l) - d_{\mathcal{S}}(s_k',s_l')| \leq \epsilon, \\ m_{\mathcal{S}_N}(s_i) m_{\mathcal{Q}_M}(q_k) |d_{\mathcal{S}}(s_i,s_k') - d_{\mathcal{Q}}(q_k,q_i')| \leq \epsilon, \\ D(1 - m_{\mathcal{Q}_M}(q_i')) m_{\mathcal{S}_N}(s_i) \leq \epsilon, \\ D(1 - m_{\mathcal{S}_N}(s_k')) m_{\mathcal{Q}}(q_k) \leq \epsilon, \end{cases} \tag{43}$$

where $i,j = 1,\ldots,N$ and $k,l = 1,\ldots,M$, the geodesic distances $d_{\mathcal{S}}(s_i,s_j)$ and $d_{\mathcal{Q}}(q_k,q_l)$ are pre-computed by FMM and the distances $d_{\mathcal{Q}}(q_i',q_k')$, $d_{\mathcal{S}}(s_k',s_l')$, $d_{\mathcal{S}}(s_i,s_k')$ and $d_{\mathcal{Q}}(q_k,q_i')$ are interpolated. On the other hand, if we fix $s_1',\ldots,s_M'$ and $q_1',\ldots,q_N'$, we can solve problem (42) with respect to $m_{\mathcal{S}_N}, m_{\mathcal{Q}_M}$ only, where the values $m_{\mathcal{S}_N}(s_i')$ and $m_{\mathcal{Q}_M}(q_i')$ are computed by interpolation.

The computation of $\tilde{d}_{\mathrm{P}}(\mathcal{S}_N, \mathcal{Q}_M)$, for every value of λ, is performed by alternating minimization in two steps: first, we fix $m_{\mathcal{S}_N}, m_{\mathcal{Q}_M}$ and solve (43). Second, we fix $s_1',\ldots,s_M'$ and $q_1',\ldots,q_N'$ and solve (42) by optimizing over $m_{\mathcal{S}_N}, m_{\mathcal{Q}_M}$. The process is repeated until convergence, which gives us a single point on the Pareto frontier corresponding to the selected value of λ. The whole scheme is repeated for another value of λ.

The entire computation is summarized in Algorithm 3.

Selection of larger values of D results in crisper parts.

initialization: $\tilde{d}_{\mathrm{P}}(\mathcal{S}_N, \mathcal{Q}_M) = \emptyset$.
1. **for** $\lambda_0 = 0, \Delta\lambda, \ldots, 1^{\mathrm{T}} \mu_{\mathcal{S}_N} + 1^{\mathrm{T}} \mu_{\mathcal{Q}_M}$ **do**
   initialization: $k = 0$; $m_{\mathcal{S}_N}^{(0)} = 1$, $m_{\mathcal{Q}_M}^{(0)} = 1$; $s_1'^{(0)},\ldots,s_M'^{(0)}$; $q_1',\ldots,q_N'$.
2. **repeat**
3. Compute the $(k+1)$st iteration solution $s_1'^{(k+1)},\ldots,s_M'^{(k+1)}, q_1'^{(k+1)},\ldots,q_N'^{(k+1)}$ by solving problem (43) with fixed $m_{\mathcal{S}_N}^{(k)}, m_{\mathcal{Q}_M}^{(k)}$.
4. Compute the $(k+1)$st iteration solution $m_{\mathcal{S}_N}^{(k+1)}, m_{\mathcal{Q}_M}^{(k+1)}$ by solving problem (42) with fixed $s_1'^{(k)},\ldots,s_M'^{(k)}, q_1'^{(k)},\ldots,q_N'^{(k)}$.
5. Set $k \longleftarrow k+1$.
6. **until** *convergence*
7. Set $m_{\mathcal{S}_N}^* = m_{\mathcal{S}_N}^{(k)}, m_{\mathcal{Q}_M}^* = m_{\mathcal{Q}_M}^{(k)}$.
8. Add a point to the Pareto frontier,
$$\tilde{d}_{\mathrm{P}}(\mathcal{S}_N, \mathcal{Q}_M) = \tilde{d}_{\mathrm{P}}(\mathcal{S}_N, \mathcal{Q}_M) \cup \{(\lambda_0, \tilde{d}_{\mathrm{GH}}(m_{\mathcal{S}_N}^*, m_{\mathcal{Q}_M}^*))\}.$$
9. **end**

Algorithm 3 Fuzzy partial dissimilarity computation

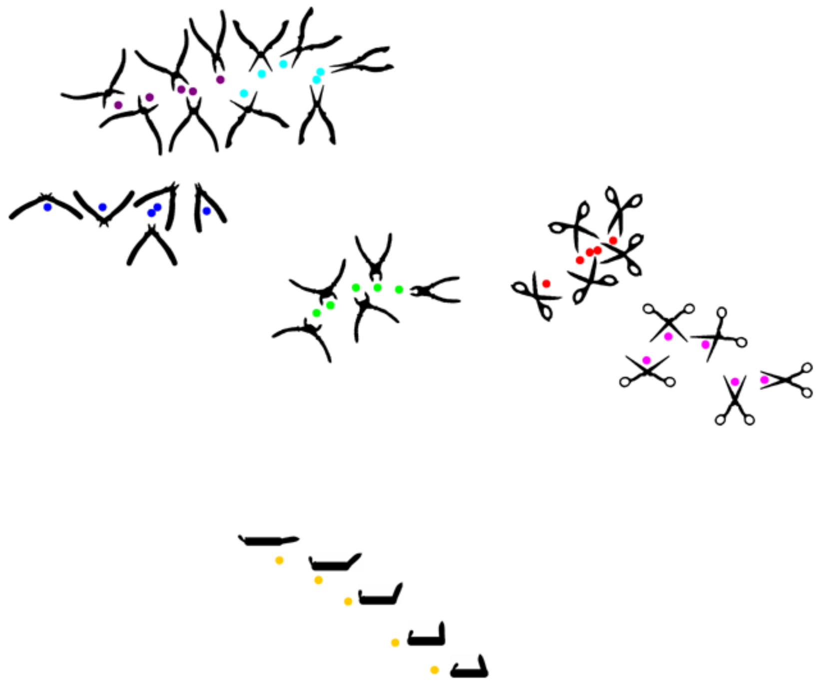

Fig. 9 (Color online) Visualization of full similarity between the Tool shapes. Each point represents a shape, and the Euclidean distance between a pair of points approximates the computed Gromov-Hausdorff distance between the corresponding shapes

7 Results

In order to evaluate our approach, we performed three experiments. The first two experiments demonstrate full and partial matching between articulated shapes. In the third experiment, we show the correspondence problem. In all the experiments, shapes were represented as binary images and triangulated using Delaunay triangulation. A typical shape contained about 2500 points. The inner geodesic distances were computed using an efficient parallel version of FMM optimized for the Intel SSE2 architecture (using our implementation, a matrix of distances of size 2500×2500 can be computed in about 1.5 seconds on a PC workstation).

The similarities between the shapes were computed using GMDS. We used a multiresolution optimization scheme, initialized at 5 points at the coarsest resolution. A total of $N = 50$ points were used in all the experiments. Note that such a relatively small number of points is still sufficient for accurate recognition of shapes.

All the data and codes will be available for academic use at `http://tosca.cs.technion.ac.il` after the approval of the associated patent. Additional experimental results can be found in Bronstein et al. (2006f).

7.1 Full Comparison

In the first experiment, we used the Tools data set (Bronstein et al. 2006f) to exemplify the comparison of articulated shapes. The data set contained seven shapes of different tools, each in five articulations. The tools were classified into four groups: scissors, pliers, pincers, cutters and knife. The knife had three parts and two joints; all the rest of the tools had four parts and one joint.

Figure 9 visualizes the shape space with d_{GH} using a Euclidean similarity pattern. Semantically similar shapes are clearly distinguishable as clusters in this plot. For example, the two different types of pliers form two close clusters, and two types of scissors form another two close clusters. On the other hand, dissimilar shapes like the knife form a distant cluster.

7.2 Partial Comparison

In the second experiment, we used the Mythological Creatures data set in order to demonstrate partial matching. The data set consisted of fifteen shapes of horses, humans and centaurs, which appeared in different articulations (e.g. different positions of hands and legs), as well as with different modifications (e.g. centaurs holding a spear, a sword and a whip).

Figures 10 and 11 depict the Gromov-Hausdorff and the scalar partial dissimilarity between the shapes. The results demonstrate the difference between full and partial matching, and show the advantage of the latter. In terms of full similarity, a horse and a winged Pegasus are dissimilar, since they are not isometric. However, in terms of partial similarity, these shapes are similar as they have a similar large part (the equine body).

The difference between full and partial similarity criteria can be clearly seen in Fig. 12, depicting the set-valued distances (Pareto frontiers) between the shapes of a man and a spear-bearer (solid curve), and a centaur (dotted curve).

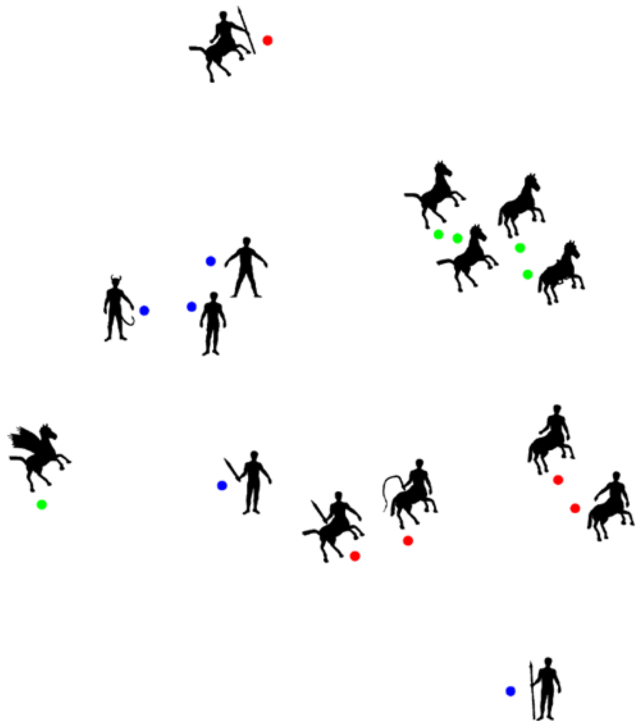

Fig. 10 (Color online) Visualization of full similarity (Gromov-Hausdorff distance) between the Mythological Creatures

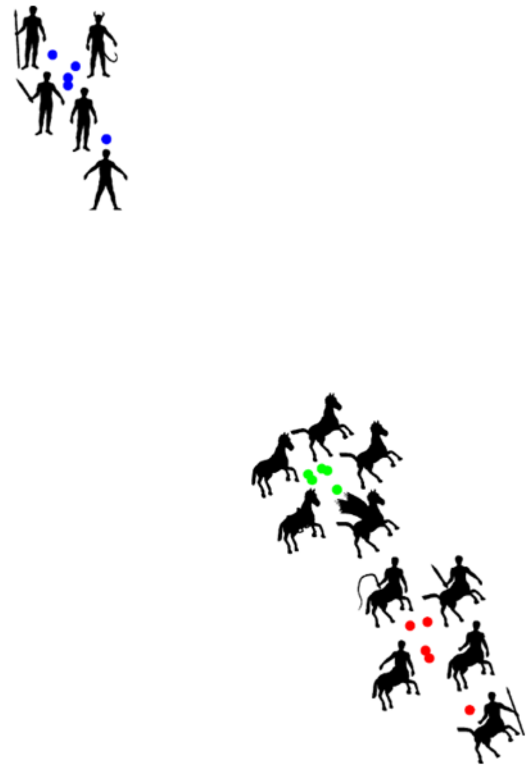

Fig. 11 (Color online) Visualization of the scalar partial similarity between the Mythological Creatures

The values of ϵ at $\lambda = 0$ correspond to the values of d_{GH}; it follows that the man—centaur dissimilarity ($\sim$0.65) is nearly 1.5 times larger than the man—spear-bearer dissimilarity ($\sim$0.45). However, if we look at the Pareto frontiers, we see that the first curve decays significantly faster. This implies that by removing a small part from the spear-bearer, we can make it similar to the man's shape. This information is captured by the scalar partial dissimilarity (Salukwadze distance), which differs approximately by an order of magnitude.

7.3 Correspondence

In the third experiment, we used GMDS to solve the correspondence problem. Figure 13 depicts full correspondence between two articulated horse shapes; the Voronoi cells are used to represent corresponding points. We can see that the correspondence is accurate despite strong deformations of the shapes. Figure 14 depicts partial correspondence between horse and Pegasus shapes (the crisp parts shown are obtained by thresholding). The correspondence is accurate, despite large dissimilar parts.

8 Conclusions

We presented a general framework for the analysis of non-rigid two-dimensional shapes based on their intrinsic geometric properties. Using an axiomatic construction, we defined similarity criteria for shape comparison and studied similarity criteria proposed in prior works. We thus gave a theoretical justification to the use of the Gromov-Hausdorff distance, and also showed that the canonical forms method (Elad and Kimmel 2001; Ling and Jacobs 2005) has somewhat weaker properties.

As the numerical framework for the efficient computation of our similarity criteria, we used the GMDS algorithm. Being a convex optimization method, this algorithm by its nature is prone to converge to a local minimum. We showed an efficient scheme for initializing the GMDS in order to ensure global convergence. The same numerical methods were also used for solving the correspondence problem between non-rigid shapes.

For the problem of partial shape comparison, we introduced the Pareto framework and showed how this idea leads to a new concept of set-valued distances. Such an approach is generic, and can be applied to measuring partial similarity of different objects, such as text sequences.

The presented approach can be extended to finding similarity and correspondence between grayscale and color images by augmenting the geometric similarity criteria with photometric information.

Acknowledgements This research was supported by the Center for Security Science and Technology and by the FP-6 MUSCLE program.

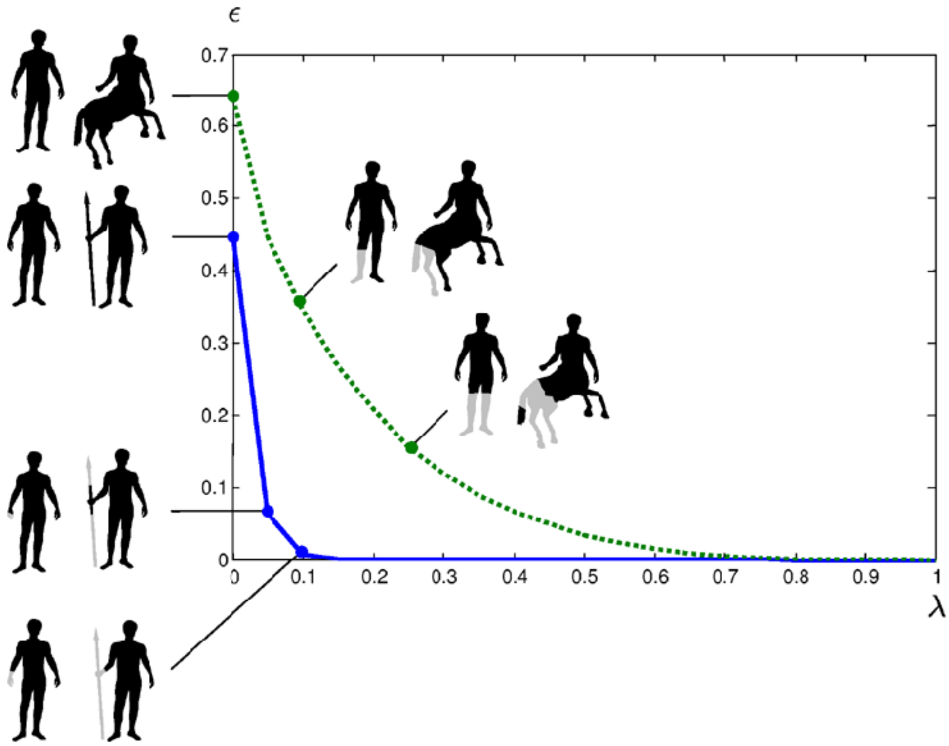

Fig. 12 (Color online) Set-valued partial dissimilarities between mythological creatures. The optimal parts corresponding to points on the Pareto frontier are shown in black

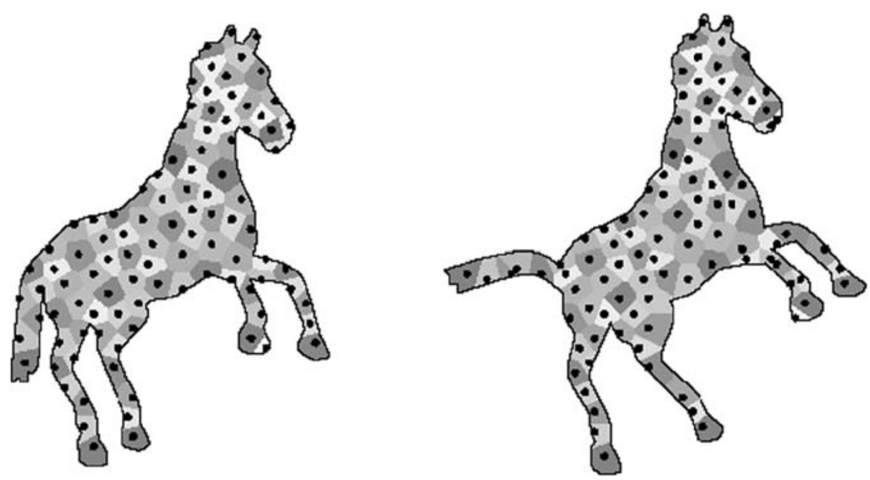

Fig. 13 Full correspondence between two articulated objects

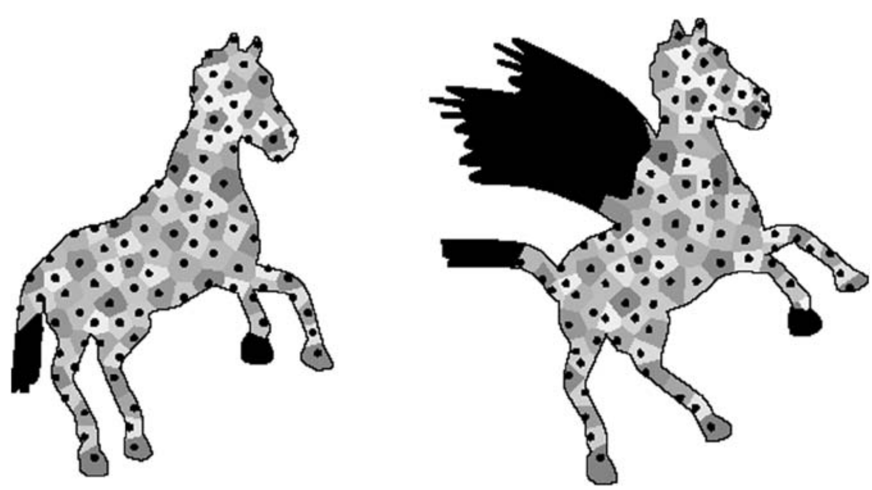

Fig. 14 Partial correspondence between two articulated objects (parts obtained by thresholding at 0.2)

Appendix

Proof of Theorem 2 Properties (F1) and (F2) hold by definition of d_{CF}. To show (F3), let $\mathcal{S}$, $\mathcal{Q}$ and $\mathcal{R}$ be shapes, embeddable into $(\mathbb{X}, d_{\mathbb{X}})$ by the maps $\varphi : \mathcal{S} \to \mathbb{X}$, $\psi : \mathcal{Q} \to \mathbb{X}$ and $\eta : \mathcal{R} \to \mathbb{X}$. Let $i, j \in \mathrm{Iso}(\mathbb{X})$ be isometries in the embedding space. Since the Hausdorff distance satisfies the triangle inequality, we have

$$\begin{aligned} d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{S}) &= \inf_{i \in \mathrm{Iso}(\mathbb{X})} d_{\mathrm{H}}^{\mathbb{X}}(\psi(\mathcal{Q}), i \circ \varphi(\mathcal{S})) \\ &\le d_{\mathrm{H}}^{\mathbb{X}}(\psi(\mathcal{Q}), j \circ \varphi(\mathcal{S})) \\ &\le d_{\mathrm{H}}^{\mathbb{X}}(\psi(\mathcal{Q}), i \circ \eta(\mathcal{R})) \\ &\quad + d_{\mathrm{H}}^{\mathbb{X}}(i \circ \eta(\mathcal{R}), j \circ \varphi(\mathcal{S})). \end{aligned} \tag{44}$$

We define a sequence of isometries $\{i_1, i_2, \ldots\} \subset \mathrm{Iso}(\mathbb{X})$ and $\{j_1, j_2, \ldots\} \subset \mathrm{Iso}(\mathbb{X})$ such that

$$\lim_{n \to \infty} d_{\mathrm{H}}^{\mathbb{X}}(\psi(\mathcal{Q}), i_n \circ \eta(\mathcal{R}))$$

$$= \inf_{i \in \mathrm{Iso}(\mathbb{X})} d_{\mathrm{H}}^{\mathbb{X}}(\psi(\mathcal{Q}), i \circ \eta(\mathcal{R})) \tag{45}$$

$$= d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{R}), \tag{46}$$

and

$$\lim_{n \to \infty} d_{\mathrm{H}}^{\mathbb{X}}(i \circ \eta(\mathcal{R}), j_n \circ \varphi(\mathcal{S}))$$

$$= \inf_{i \in \mathrm{Iso}(\mathbb{X})} d_{\mathrm{H}}^{\mathbb{X}}(\eta(\mathcal{R}), i \circ \varphi(\mathcal{S})) \tag{47}$$

$$= d_{\mathrm{CF}}(\mathcal{R}, \mathcal{S}). \tag{48}$$

Using $i = i_n$ and $j = j_n$ in (49) and taking the limit $n \to \infty$ on the right hand side, we obtain the triangle inequality,

$$d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{S}) \le d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{R}) + d_{\mathrm{CF}}(\mathcal{R}, \mathcal{S}). \tag{49}$$

To show (F4w), we establish a relation between the Gromov-Hausdorff and the canonical forms distance. Trivially, $d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{S}) \ge d_{\mathrm{GH}}(\psi(\mathcal{Q}), \varphi(\mathcal{S}))$. Assuming that the embeddings $\varphi : \mathcal{S} \to \mathbb{X}$ and $\psi : \mathcal{Q} \to \mathbb{X}$ have distortions $\mathrm{dis}\,\varphi \le \delta$ and $\mathrm{dis}\,\psi \le \delta'$, respectively, $\mathcal{S}$ and $\varphi(\mathcal{S})$ are δ-isometric. Using property (F4) satisfied by the Gromov-Hausdorff distance, this implies that $d_{\mathrm{GH}}(\mathcal{S}, \varphi(\mathcal{S})) \le 2\delta$. Similarly, $d_{\mathrm{GH}}(\mathcal{Q}, \varphi(\mathcal{Q})) \le 2\delta'$. Using the triangle inequality, we have

$$\begin{aligned} d_{\mathrm{GH}}(\mathcal{Q}, \mathcal{S}) &\le d_{\mathrm{GH}}(\psi(\mathcal{Q}), \varphi(\mathcal{S})) \\ &\quad + d_{\mathrm{GH}}(\mathcal{S}, \varphi(\mathcal{S})) d_{\mathrm{GH}}(\mathcal{Q}, \psi(\mathcal{Q})) \\ &\le d_{\mathrm{GH}}(\psi(\mathcal{Q}), \varphi(\mathcal{S})) + 2(\delta + \delta'), \end{aligned} \tag{50}$$

from which it follows that

$$\begin{aligned} d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{S}) &\ge d_{\mathrm{GH}}(\psi(\mathcal{Q}), \varphi(\mathcal{S})) \\ &\ge d_{\mathrm{GH}}(\mathcal{Q}, \mathcal{S}) - 2(\delta + \delta'). \end{aligned} \tag{51}$$

Hence, if $d_{\mathrm{CF}}(\mathcal{Q}, \mathcal{S}) \le \epsilon$, then $d_{\mathrm{GH}} \le \epsilon + 2(\delta + \delta')$, from which we conclude that $\mathcal{S}$ and $\mathcal{Q}$ are $2\epsilon + 4(\delta + \delta')$-isometric. This completes the proof. □

Proof of Theorem 3 In order to show the equivalence, we have to show that though the maps $f : \mathcal{S} \to \mathcal{Q}$ and $g : \mathcal{Q} \to \mathcal{S}$ are defined on the entire shapes, their ranges and images are $\mathcal{S}'$ and $\mathcal{Q}'$. Given a crisp part $\mathcal{S}'$, we denote by $m_{\mathcal{S}}$ its characteristic function. The characteristic functions in the infima terms restrict the ranges,

$$\frac{1}{2} \inf_{\substack{f:\mathcal{S}\to\mathcal{Q} \\ g:\mathcal{Q}\to\mathcal{S}}} \max \left\{ \begin{array}{l} \sup\limits_{s,s' \in \mathcal{S}} m_{\mathcal{S}}(s) m_{\mathcal{S}}(s') \\ \quad \times |d_{\mathcal{S}}(s, s') - d_{\mathcal{Q}}(f(s), f(s'))| \\ \sup\limits_{q,q' \in \mathcal{Q}} m_{\mathcal{Q}}(q) m_{\mathcal{Q}}(q') \\ \quad \times |d_{\mathcal{Q}}(q, q') - d_{\mathcal{S}}(g(q), g(q'))| \\ \sup\limits_{\substack{s \in \mathcal{S} \\ q \in \mathcal{Q}}} m_{\mathcal{S}}(s) m_{\mathcal{Q}}(q) \\ \quad \times |d_{\mathcal{S}}(s, g(q)) - d_{\mathcal{Q}}(f(s), q)| \\ D \sup\limits_{s \in \mathcal{S}} (1 - m_{\mathcal{Q}}(f(s))) m_{\mathcal{S}}(s) \\ D \sup\limits_{q \in \mathcal{Q}} (1 - m_{\mathcal{S}}(g(q))) m_{\mathcal{Q}}(q) \end{array} \right\}$$

$$= \frac{1}{2} \inf_{\substack{f:\mathcal{S}'\to\mathcal{Q} \\ g:\mathcal{Q}'\to\mathcal{S}}} \max \left\{ \begin{array}{l} \sup\limits_{s,s' \in \mathcal{S}'} |d_{\mathcal{S}}(s, s') - d_{\mathcal{Q}}(f(s), f(s'))| \\ \sup\limits_{q,q' \in \mathcal{Q}'} |d_{\mathcal{Q}}(q, q') - d_{\mathcal{S}}(g(q), g(q'))| \\ \sup\limits_{\substack{s \in \mathcal{S}' \\ q \in \mathcal{Q}'}} |d_{\mathcal{S}}(s, g(q)) - d_{\mathcal{Q}}(f(s), q)| \\ D \sup\limits_{s \in \mathcal{S}'} (1 - m_{\mathcal{Q}}(f(s))) \\ D \sup\limits_{q \in \mathcal{Q}'} (1 - m_{\mathcal{S}}(g(q))) \end{array} \right\},$$

assuming $D = \max\{\mathrm{diam}(\mathcal{S}), \mathrm{diam}(\mathcal{Q})\}$.

If $f(\mathcal{S}') \not\subseteq \mathcal{Q}'$ or $g(\mathcal{Q}') \not\subseteq \mathcal{S}'$, we have $\sup_{s \in \mathcal{S}'} (1 - m_{\mathcal{Q}}(f(s))) = 1$ (respectively, $\sup_{q \in \mathcal{Q}'} (1 - m_{\mathcal{S}}(g(q))) = 1$); hence, the values of the above expression will be at least D. Since the other terms are bounded above by D, it follows that for $f(\mathcal{S}') \subseteq \mathcal{Q}'$ and $g(\mathcal{Q}') \subseteq \mathcal{S}'$, the above expression will be at most D. As a result, solutions with $f(\mathcal{S}') \not\subseteq \mathcal{Q}'$ or $g(\mathcal{Q}') \not\subseteq \mathcal{S}'$ are always suboptimal, which implies that the images of f and g are $\mathcal{Q}'$ and $\mathcal{S}'$, respectively. It follows that we can rewrite the above expressions as

$$\frac{1}{2} \inf_{\substack{f:\mathcal{S}'\to\mathcal{Q}' \\ g:\mathcal{Q}'\to\mathcal{S}'}} \max \left\{ \begin{array}{l} \sup\limits_{s,s' \in \mathcal{S}'} |d_{\mathcal{S}}(s, s') - d_{\mathcal{Q}}(f(s), f(s'))| \\ \sup\limits_{q,q' \in \mathcal{Q}'} |d_{\mathcal{Q}}(q, q') - d_{\mathcal{S}}(g(q), g(q'))| \\ \sup\limits_{\substack{s \in \mathcal{S}' \\ q \in \mathcal{Q}'}} |d_{\mathcal{S}}(s, g(q)) - d_{\mathcal{Q}}(f(s), q)| \end{array} \right\}$$

$$= d_{\mathrm{GH}}(\mathcal{S}', \mathcal{Q}'),$$

which completes the proof. □

References

Bajcsy, R., & Solina, F. (1987) Three dimensional object representation revisited. In *Proceedings of ICCV* (pp. 231–240).

Belongie, S., Malik, J., & Puzicha, J. (2002). Shape matching and object recognition using shape context. *IEEE Transactions on PAMI*, *24*, 509–522.

Berg, A.C., Berg, T. L., & Malik, J. (2005). Shape matching and object recognition using low distortion correspondences. In *Proceedings of CVPR*.

Biderman, I. (1985). Human image understanding: recent research and a theory. *Computer Graphics, Vision and Image Processing*, *32*, 29–73.

Binford, T. O. (1987). Visual perception by computer. In *Proceedings of IEEE conference on systems and control* (pp. 231–240).

Black, M. J., & Anandan, P. (1993). A framework for the robust estimation of optical flow. In *Proceedings of ICCV* (pp. 231–236).

Borg, I., & Groenen, P. (1997). *Modern multidimensional scaling—theory and applications*. Berlin: Springer.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2003). Expression-invariant 3D face recognition. In *Lecture notes on computer science: Vol. 2688*. *Proceedings of audio and video-based biometric person authentication* (pp. 62–69). Berlin: Springer.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2005a). Three-dimensional face recognition. *International Journal of Computer Vision*, *64*(1), 5–30.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2005b). Expression-invariant face recognition via spherical embedding. In *Proceedings of IEEE international conference on image processing (ICIP)* (Vol. 3, pp. 756–759).

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2005c). On isometric embedding of facial surfaces into $\mathbb{S}^3$. In *Lecture notes on computer science: Vol. 3459*. *Proceedings of international conference on scale space and PDE methods in computer vision*. (pp. 622–631). Berlin: Springer.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006a). Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, *103*(5), 1168–1172.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006b). Efficient computation of isometry-invariant distances between surfaces, *SIAM Journal of Scientific Computing*. *28*(5), 1812–1836.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006c). Robust expression-invariant face recognition from partially missing data. In *Proceedings of European conference on computer vision (ECCV)* (pp. 396–408).

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006d). Face2face: an isometric model for facial animation. In *Proceedings of conference on articulated motion and deformable objects (AMDO)* (pp. 38–47).

Bronstein, M. M., Bronstein, A. M., Kimmel, R., & Yavneh, I. (2006e). Multigrid multidimensional scaling. *Numerical Linear Algebra with Applications (NLAA)*, *13*, 149–171.

Bronstein, A. M., Bronstein, M. M., Bruckstein, A. M., & Kimmel, R. (2006f). Matching two-dimensional articulated shapes using generalized multidimensional scaling. In *Proceedings of conference on articulated motion and deformable objects (AMDO)* (pp. 48–57).

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2007a). Calculus of non-rigid surfaces for geometry and texture manipulation. *IEEE Transactions on Visualization and Computer Graphics*, *13*(5), 902–913.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2007b). *Numerical geometry of non-rigid shapes*. Berlin: Springer.

Bronstein, A. M., Bronstein, M. M., Devir, Y. S., Kimmel, R., & Weber, O. (2007c). *Parallel algorithms for approximation of distance maps on parametric surfaces*. (Technical Report CIS-2007-03). Department of Computer Science, Technion, Israel, April 2007. *ACM Transactions on Graphics* (in press).

Brooks, R. (1981). Symbolic reasoning among 3-D models and 2-D images. *Artificial Intelligence*, *17*, 285–348.

Bruckstein, A. M., Holt, R. J., & Netravali, A. N. (1998). Holographic representations of images. *IEEE Transactions on Image Processing*, *7*(11), 1583–1597.

Burago, D., Burago, Y., & Ivanov, S. (2001). A course in metric geometry. In *Graduate studies in mathematics,* (Vol. 33). Providence: American Mathematical Society.

Burr, D. (1981). Elastic matching of line drawings. *IEEE Transactions on PAMI*, *3*(6), 708–713.

Charpiat, G., Maurel, P., Pons, J. P., Keriven, R., & Faugeras, O. (2007). Generalized gradients: priors on minimization flows. *International Journal of Computer Vision*, *73*(3), 325–344.

Cheng, S. W., Edelsbrunner, H., Fu, P., & Lam, K. P. (2001). Design and analysis of planar shape deformation. *Computational Geometry: Theory and Applications*, *19*(2–3), 205–218.

Cohen, I., Ayache, N., & Sulger, P. (1992). Tracking points on deformable objects using curvature information. In *Proceedings of ECCV* (pp. 458–466).

Connell, J. H., & Brady, M. (1987). Generating and generalizing models of visual objects. *Artificial Intelligence*, *31*, 159–183.

de Rooij, S., & Vitanyi, P. (2006). Approximating rate-distortion graphs of individual data: Experiments in lossy compression and denoising. *IEEE Transactions on Information Theory* (submitted).

Elad, A., & Kimmel, R. (2001). Bending invariant representations for surfaces. In *Proceedings of CVPR* (pp. 168–174).

Elad, A., & Kimmel, R. (2002). Spherical flattening of the cortex surface. In *Lecture notes in computer Science: Vol. 2191*. *Geometric methods in bio-medical image processing* (pp. 77–89). Berlin: Springer.

Eldar, Y., Lindenbaum, M., Porat, M., & Zeevi, Y. Y. (1997). The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, *6*(9), 1305–1315.

Everson, R. M., & Fieldsend, J. E. (2006). Multi-class ROC analysis from a multi-objective optimization perspective. *Pattern Recognition Letters*, *27*(8), 918–927.

Felzenszwalb, P. F. (2005). Representation and detection of deformable shapes. *IEEE Transactions on PAMI*, *27*(2), 208–220.

Felzenszwalb, P. F., & Huttenlocher, D. P. (2005). Pictorial structures for object recognition. *International Journal of Computer Vision*, *61*(1), 55–79.

Felzenszwalb, P. F., & Schwartz, J. (2007). Hierarchical matching of deformable shapes. In *Proceedings of CVPR*.

Fischler, M. A., & Elschlager, R. A. (1973). The representation and matching of pictorial structures. *IEEE Transactions on Computers*, *100*(22), 67–92.

Forsyth, D. A., & Ponce, J. (2003). *Computer vision: A modern approach*. Englewood Cliffs: Prentice Hall.

Gdalyahu, Y., & Weinshall, D. (1999). Flexible syntactic matching of curves and its application to automatic hierarchical classification of silhouettes. *IEEE Transactions on PAMI*, *21*, 1312–1328.

Geiger, D., Liu, T. L., & Kohn, R. (2003). Representation and self-similarity of shapes. *IEEE Transactions on PAMI*, *25*(1), 86–99.

Gelfand, N., Mitra, N. J., Guibas, L., & Pottmann, H. (2005). Robust global registration. In *Proceedings of symposium on geometry processing (SGP)*.

Grenander, U., Chow, Y., & Keenan, D. M. (1991). *Hands: A pattern theoretic study of biological shapes*. Berlin: Springer.

Gromov, M. (1981). *Textes mathématiques*: *Vol. 1*. *Structures métriques pour les variétés riemanniennes*.

Jacobs, D., & Ling, H. (2005). Deformation invariant image matching. In *Proceedings of ICCV*.

Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., & Stahel, W. A. (1986). *Robust statistics: The approach based on influence functions*. New York

Hel-Or, Y., & Werman, M. (1994). Constraint-fusion for interpretation of articulated objects. In *Proceedings of CVPR* (pp. 39–45).

Hildreth, E. (1983). *The measurement of visual motion*. Cambridge: MIT Press.

Hoffman, D., & Richards, W. (1984). Visual cognition. In *Parts of recognition*. Cambridge: MIT Press.

Huber, P. J. (2004). *Robust statistics*. New York: Wiley.

Jacobs, D., Weinshall, D., & Gdalyahu, Y. (2000). Class representation and image retrieval with non-metric distances. *IEEE Transactions on PAMI*, *22*, 583–600.

Jain, A., Zhong, Y., & Lakshmanan, S. (1996). Object matching using deformable templates. *IEEE Transactions on PAMI*, *3*(18), 267–278.

Kass, M., Witkin, A., & Terzopoulos, D. (1988). Snakes: active contour models. *International Journal of Computer Vision*, *1*(4), 321–331.

Kimmel, R., & Sethian, J. A. (1998). Computing geodesic on manifolds. In *Proceedings of US National Academy of Science* (Vol. 95, pp. 8431–8435).

Klir, G. J., & Yuan, B. (1994). *Fuzzy sets and fuzzy logic: theory and applications*. Englewood Cliffs: Prentice Hall.

Koenderink, J. J., & van Doorn, A. J. (1981). The shape of smooth objects and the way contours end. *Perception*, *11*, 129–137.

Kupeev, K., & Wolfson, H. (1994). On shape similarity. In *Proceedings of international conference on pattern recognition* (pp. 227–237).

Lades, M., Vorbruggen, J. C., Buhmann, J., Lange, J., von der Malsburg, C., Wurtz, R. P., & Konen, W. (1993). Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers*, *42*(3), 300–311.

Latecki, L. J., & Lakamper, R. (2000). Shape similarity measure based on correspondence of visual parts. *IEEE Transactions on PAMI*, *22*(10), 1185–1190.

Latecki, L. J., Lakamper, R., & Wolter, D. (2005). Optimal partial shape similarity. *Image and Vision Computing*, *23*, 227–236.

Ling, H., & Jacobs, D. (2005). Using the inner-distance for classification of articulated shapes. In: *Proceedings of CVPR*.

Mémoli, F., & Sapiro, G. (2005). A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, *5*(3), 313–347.

Michor, P. W., & Mumford, D. (2003). Riemannian geometries on spaces of plane curves. Arxiv preprint math.DG/0312384.

Mio, W., Srivastava, A., & Joshi, S. (2007). On shape of plane elastic curves. *International Journal of Computer Vision*, *73*(3), 307–324.

Mumford, D. (1991). Mathematical theories of shape: do they model perception? In *Proceedings of SPIE* (pp. 2–10).

Pareto, V. (1906). *Manuale di economia politica*.

Pentland, A. (1987). Recognition by parts. In *Proceedings of ICCV* (pp. 612–620).

Platel, B., Balmachnova, E., Florack, L. M. J., Kanters, F. M. W., & ter Haar Romeny, B. M. (2005). Using top-points as interest points for image matchng. In *Lecture notes in computer science* (Vol. 3753, p. 211).

Geiger, D., Basri, R., Costa, L., & Jacobs, D. (1998). Determining the similarity of deformable shapes. *Vision Research*, *38*, 2365–2385.

Rivlin, E., Dickenson, S., & Rosenfeld, A. (1992). Recognition by functional parts. In *Proceedings of CVPR* (pp. 267–275).

Salukwadze, M. E. (1979). *Vector-valued optimization problems in control theory*. Orlando: Academic Press.

Sethian, J. A. (1996). A review of the theory, algorithms, and applications of level set method for propagating surfaces. *Acta numerica* 309–395.

Siddiqi, K., & Kimia, B. (1996). A shock grammar for recognition. In *Proceedings of CVPR* (pp. 507–513).

Spira, A., & Kimmel, R. (2004). *An efficient solution to the eikonal equation on parametric manifolds* (Vol. 6, pp. 315–327).

Stark, L., & Bowyer, K. (1991). Achieving generalized object recognition through reasoning about association of function to structure. *IEEE Transactions on PAMI*, *10*(13), 992–1006.

Tappert, C. (1982). Cursive script recognition by elastic matching. *IBM Journal of Research and Development*, *26*(6), 765–771.

Ullman, S. (1989). Aligning pictorial descriptions: an approach to object recognition. *Cognition*, *3*(32), 193–254.

Walter, J., & Ritter, H. (2002). On interactive visualization of high-dimensional data using the hyperbolic plane. In *Proceedings of ACM SIGKDD int. conference on knowledge discovery and data mining* (pp. 123–131).

Yezzi, A., & Mennucci, A. (2004). *Metrics in the space of curves*. Arxiv preprint math.DG/0412454.

Yuille, A., Cohen, D., & Hallinan, P. (1989). Feature extraction from faces using deformable templates. In *Proceedings of CVPR* (pp. 104–109).

Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, *8*, 338–353.

Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, *1*(1), 3–28.

Zhang, J., Collins, R., & Liu, Y. (2004). Representation and matching of articulated shapes. In *Proceedings of CVPR* (Vol. 2, pp. 342–349).

Zimmermann, H. J. (2001). *Fuzzy set theory and its applications*. Dordrecht: Kluwer Academic.