Topology-invariant similarity of nonrigid shapes

Alexander M. Bronstein, Michael M. Bronstein and Ron Kimmel

June 10, 2008

Abstract

This paper explores the problem of similarity criteria between nonrigid shapes. Broadly speaking, such criteria are divided into intrinsic and extrinsic, the first referring to the metric structure of the object and the latter to how it is laid out in the Euclidean space. Both criteria have their advantages and disadvantages: extrinsic similarity is sensitive to nonrigid deformations, while intrinsic similarity is sensitive to topological noise. In this paper, we approach the problem from the perspective of metric geometry. We show that by unifying the extrinsic and intrinsic similarity criteria, it is possible to obtain a stronger topology-invariant similarity, suitable for comparing deformed shapes with different topology. We construct this new joint criterion as a tradeoff between the extrinsic and intrinsic similarity and use it as a set-valued distance. Numerical results demonstrate the efficiency of our approach in cases where using either extrinsic or intrinsic criteria alone would fail.

1 Introduction

In the childhood game Rock, Paper, Scissors, the players bend their fingers in different ways to make the hand resemble one of three objects: a rock, a sheet of paper and scissors. Looking at these shapes, we can recognize the hand postures as the objects they intend to imitate. The rock is represented by a closed fist, the paper by an open palm and the scissors by the extended index and middle fingers (Figure 1). At the same time, we can say that all these shapes are just different articulations of the same hand.

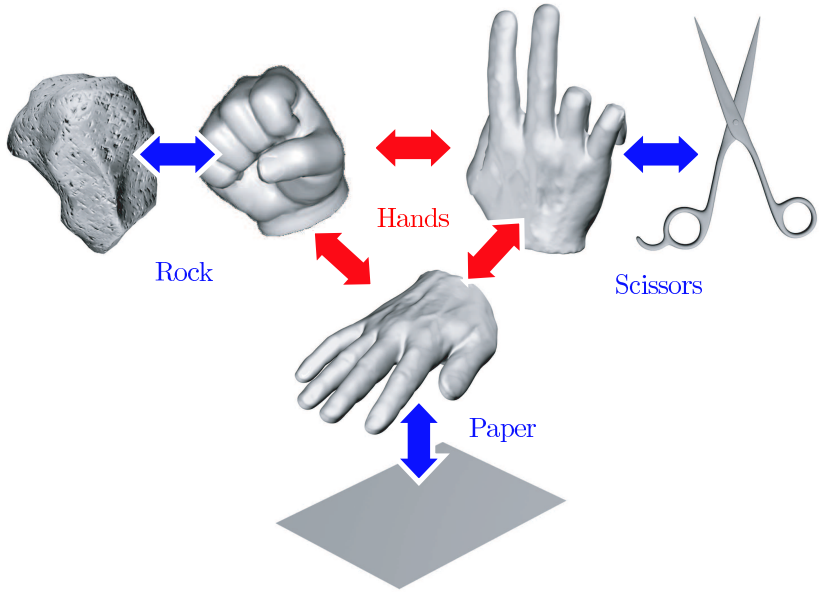


Figure 1: Rock, Paper, and Scissors: this childhood game is based on the similarity of the postures of the hand to different objects (blue arrows). Though being a valid similarity criterion, in many other applications, we would rather be interested in saying that the different postures are all similar, being nonrigid deformations of the same object (red arrows).

This example demonstrates the difficulty of defining the similarity of nonrigid shapes. On one hand, instances of nonrigid objects can be considered as stand-alone rigid shapes. On the other hand, these shapes can be regarded as nonrigid deformations of the same object. Using geometric terminology, the first similarity criterion is *extrinsic*, i.e., considers the properties of the shape related to the particular way it is laid out. The second criterion, looks at the *intrinsic* properties of the shape, *invariant* to the object deformations.

Extrinsic similarity of shapes has been widely addressed in computer vision, pattern recognition, and computational geometry literature [71, 18]. Most of the papers in these fields, either implicitly or explicitly, look at the problem of shape similarity from the extrinsic point of view (see, for example, [20, 47, 40]). A classical method for rigid object matching, introduced by Chen and Medioni [22] and Besl and McKay [5], is the *iterative closest point* (ICP) algorithm. ICP and its numerous flavors [80, 49, 56, 30] try to find a rigid transformation between two shapes, minimizing an extrinsic distance between them, usually a variant of the Hausdorff distance.

Another important class of extrinsic similarity methods is based on high-order moments [72, 39, 33, 70], where the shape's extrinsic geometry is represented by a vector of coefficients obtained from the decomposition of shape properties in some basis. Conceptually, one can think of such methods as of a Fourier-like representation [77, 78], though other bases like wavelets [60], spherical harmonics [74, 76, 44], and Zernike descriptors [58] have been explored in the literature as well. Besides "global" methods, there exist other families of extrinsic shape similarity methods based on histograms of *local* shape properties like curvatures, distances, angles and areas [59, 69]. Such local methods can be, to some extent, insensitive to shape deformations. For a comprehensive survey of these approaches, the reader is referred to [71].

At the other end, methods for the computation of intrinsic similarity of shapes started penetrating into the computer vision and pattern recognition communities relatively late. As a precursor, we consider the paper by Schwartz *et al.* [68], in which a method for the representation of the intrinsic geometry of the cortical surface of the brain using *multidimensional scaling*

(MDS) was presented. MDS is a family of algorithms [6] commonly used in dimensionality reduction and data analysis [64, 26, 73] and graph representation [53]. The idea of Schwartz *et al.* was extended by Elad and Kimmel [29], who proposed a nonrigid shape recognition method based on Euclidean embeddings. Elad and Kimmel mapped the metric structure of the surfaces to a low-dimensional Euclidean space and compared the resulting images (called *canonical forms*) in this space.[1] Canonical forms were applied to the problem of three-dimensional face recognition, where this method was proved to be efficient in recognizing the identity of people while being insensitive to their facial expressions [11]. Ling and Jacobs used this method for recognition of articulated two-dimensional shapes [52]. Other applications were in texture mapping and object morphing [81, 35, 15], mesh segmentation [43] and image matching [51].

One of the main disadvantages of the canonical forms is the fact that they can represent the intrinsic geometry of the objects only approximately, as it is generally impossible to isometrically embed a non-flat surface into a low-dimensional (and in general, finite-dimensional) Euclidean space. It was shown empirically in [75, 16] that using spaces with non-Euclidean geometry, it is possible to obtain more accurate representations. Mémoli and Sapiro [55] showed how the representation error can be theoretically avoided by using to the *Gromov-Hausdorff distance*, introduced by Mikhail Gromov in [34]. Theoretically, the computation of the discrete Gromov-Hausdorff distance is an NP-hard problem. In order to overcome this difficulty, Mémoli and Sapiro proposed an approximation, related to the Gromov-Hausdorff distance by a probabilistic bound.

Using the fact that the Gromov-Hausdorff distance can be related to the distortion of embedding one surface into another, we proposed in [13] a relaxation of the discrete Gromov-Hausdorff distance, yielding a continuous

[1] The measurement of pairwise geodesic distances and the solution of the underlying MDS problem are the two most computationally-intensive components of the canonical forms method. The measurement of geodesic distance can be performed very efficiently using the recently proposed parallel fast marching method [10]. The solution of an MDS problem can be carried out efficiently using the multigrid framework [19] or the vector extrapolation methods [63]. This allows for real-time applications like face recognition.

optimization problem similar to MDS. This algorithm, named *generalized MDS* (GMDS) [13, 12], can be thought of as a natural extension of previous works on isometric embedding into non-Euclidean spaces. GMDS was used for the comparison of two-dimensional [9] and three-dimensional [13] shapes, face recognition [14], and in a particular setting of intrinsic self-similarity for symmetry detection [61, 9]. It was also shown that GMDS can be used to find deformation-invariant correspondence between non-rigid shapes in computer graphics applications such as texture mapping and shape manipulation [15]. In these problems, GMDS is closely related to the method of Litke *et al.* [54]

Another class of intrinsic similarity methods is based on the analysis of spectral properties of the Laplace-Beltrami operator of the shape. Reuter *et al.* used Laplace-Beltrami spectra (eigenvalues), referring to them as "shape DNA", for the characterization of surfaces [62]. Since the Laplace-Beltrami operator is an intrinsic characteristic of the surface, it is insensitive to isometric deformations [24, 57, 50]. It appears, however, that such criteria are able to identify *isospectral* rather than isometric surfaces, and isospectrality is a weaker property (two surfaces can be isospectral but not isometric [42, 32, 4]). In [65], Rustamov introduced the *global point signature* (GPS) embedding, based on eigenfunctions and eigenvalues of the Laplace-Beltrami operator. Such a descriptor is exact and theoretically allows to represent the shape up to isometric deformations. GPS embeddings are intimately related to methods used in manifold learning and data analysis (see e.g. [64, 79, 3, 26, 25]) and can be thought of as infinite-dimensional canonical forms.

Though apparently completely different, from the viewpoint of metric geometry both the intrinsic and extrinsic similarity criteria can be formalized using the notion of *isometry invariance*. Regarding a shape as a *metric space*, its extrinsic properties are described by using the Euclidean metric, while intrinsic ones using the geodesic metric, which measures distances between points as the lengths of the shortest paths on the shape. Shape transformations preserving the metric are called *isometries*; extrinsic isometries are rigid motions and intrinsic isometries are inelastic deformations. Two

shapes can be thus said to be similar if they are isometric. Depending on whether we choose the Euclidean or the geodesic metric, we obtain the extrinsic or the intrinsic similarity, respectively. This perspective is recurrent in the present paper, allowing us to consider intrinsic and extrinsic similarity using the same framework.

The choice of whether to use intrinsic or extrinsic similarity depends significantly on the application. The drawback of extrinsic similarity is its sensitivity to nonrigid deformations. Using our previous example, a gesture of the human hand can be extrinsically more similar to a rock or scissors rather than another hand. This makes extrinsic criteria usually unsuitable for the analysis of nonrigid objects with significant deformations (see Figure 2, left and middle; see also Figure 4). On the other hand, extrinsic similarity is insensitive to deformation changing the *topology* of the shape (such as "gluing" the fingers of the hand in Figure 2, right; see also Figure 5). The intrinsic similarity criterion, in a sense, behaves as the opposite of the extrinsic one: it is insensitive to inelastic deformations, but is sensitive to topology changes (Figure 2, right). In practical situation, such changes can arise due to acquisition imperfections (the so-called *topological noise* often encountered in surfaces acquired using three-dimensional scanners, or reconstructed from volumetric data) or partially missing data ("holes") [8].

In the case shown in Figure 2, neither extrinsic nor intrinsic similarity *per se* is good enough, as none of these criteria is capable of saying that the three hand shapes are similar, which is the semantically desired result. In this example, we can distinguish between two types of deformations: *geometric* and *topological*. Geometric transformation change the coordinates of the points of the shape, and include, for example, rigid motions (to which the extrinsic similarity criterion is invariant) and inelastic deformations (to which intrinsic similarity is invariant). Topological transformations, on the other hand, change the "connectivity" of the shape. The extrinsic geometry does not change significantly as a result of such transformations, yet, the intrinsic one does: by modifying the connectivity, the paths between points can change as well, which can significantly alter the geodesic metric.

The criterion we need in order to capture correctly the similarity of

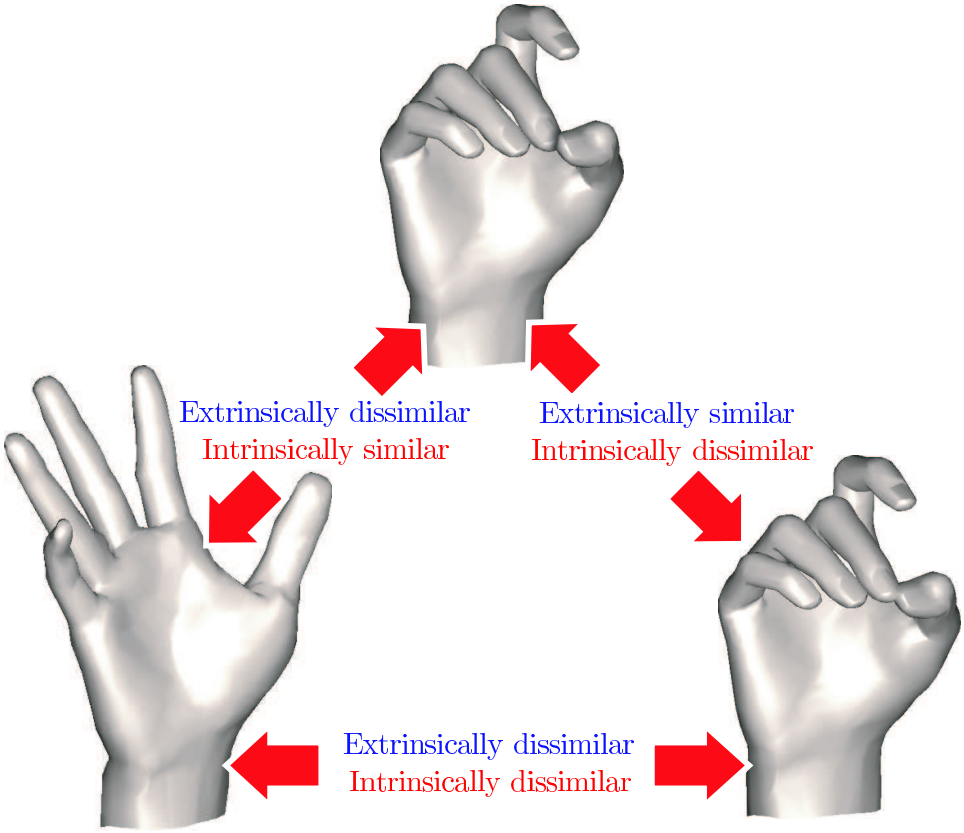


Figure 2: Illustration of the difference between intrinsic and extrinsic similarity.

shapes in Figure 2 must be insensitive to both topological and geometric deformations. We call such a criterion *topology-invariant similarity*. In [17], we proposed an approach for computing topology-invariant similarity between nonrigid shapes by combining the advantages of intrinsic and extrinsic similarity criteria, while avoiding their shortcomings. This paper presents an extended discussion and experimental validation of this method. The main idea can be visualized by looking at the example of fitting a rubber glove onto a hand. The extent to which the rubber surface is stretched represents the intrinsic geometry distortion. The fit quality, or in other words, how close the glove is to the hand surface, represents the extrinsic distance between the two objects (Figure 3). A "virtual" glove fitting is performed by moving the points of the glove with respect to the shape of the hand, trying to simultaneously minimize the misfit (extrinsic dissimilarity) and the stretching (intrinsic dissimilarity). Note that unlike real glove fitting, topological changes like glued fingers do not pose an obstacle in our case: the "virtual" glove can intersect the hand.

Finding an optimal tradeoff between the misfit and stretching can be posed as a multicriterion optimization problem and related to the notion of Pareto optimality. From this point of view, our approach is close in its spirit to the method of [48, 9, 7] for partial shape matching, using a tradeoff between the size of the parts cropped out of the shapes and the similarity between them. In our case, the set of all Pareto optimal solutions can also be represented as a set-valued similarity criterion, which contains much richer information than each of the intrinsic and extrinsic criteria separately.

Our approach can also be thought of as a generalization or "hybridization" of ICP and GMDS methods. ICP methods compute extrinsic similarity of shapes by means of finding a rigid isometry minimizing the extrinsic distance between them. Our method extends the class of transformations, allows for nonrigid isometries and near-isometries, thus relating to the recent works on nonrigid ICP [23, 36, 1]. Using GMDS, intrinsic similarity of shapes is computed as the degree of distortion when trying to embed one shape into another. This problem can be regarded as a particular case of our approach in which we constraint the extrinsic distance between the two

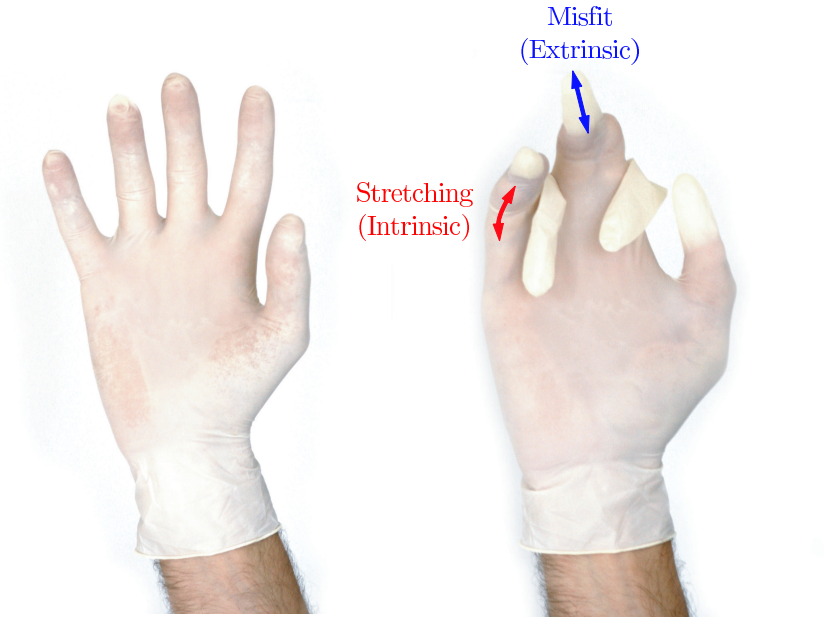


Figure 3: Fitting a glove onto one's hand may be thought of as an optimal trade-off between the stretching of the rubber surface (intrinsic dissimilarity), and the "amount of air" left between the glove and the hand (extrinsic dissimilarity).

surfaces to be zero.

The rest of this paper is organized as follows. In Section 2, we introduce the mathematical background and formulate the shape similarity problem from the perspective of metric geometry. We review standard approaches to measuring intrinsic and extrinsic similarity. In Section 3, we define topology-invariant similarity and present an approach for computing it using a joint intrinsic-extrinsic criterion. We discuss this approach using the formalism of Pareto optimality. In Section 4, we show the numerical framework for computing the proposed distance. Section 5 is dedicated to experimental results. In Section 6, we conclude the paper and discuss the relation of our approach to recent results in computer graphics literature on nonrigid ICP and "as isometric as possible" morphing.

2 Isometry-invariant similarity

We model a shape as a *metric space* (X, d), where X is a two-dimensional smooth compact connected surface (possibly with boundary) embedded in the three-dimensional Euclidean space $\mathbb{E} = \mathbb{R}^3$, and $d : X \times X \to \mathbb{R}_+$, is a *(semi) metric*[2] measuring the distance between each pair of points on X. There exist two natural choices for the metric on X. The first one is the restriction of the Euclidean metric denoted here by $d_{\mathbb{E}}(x, x') = \|x - x'\|_2$, measuring the distance between any $x, x' \in X$ along "straight lines" in $\mathbb{E}$. The second natural choice is the *geodesic metric*,

$$d_X(x, x') \quad = \quad \min_{\gamma \in \Gamma(x,x')} L(\gamma), \tag{1}$$

measuring the length of the shortest path on the surface connecting a pair of points x and x'. In order to give a formal definition of the path length $L(\gamma)$, we express γ as the limit of piece-wise linear segments connecting the points $\{x_1, ..., x_n\}$ on X,

$$L(\gamma) \quad = \quad \lim_{n\to\infty} \sum_{i=1}^{n-1} d_{\mathbb{E}}(x_i, x_{i+1}). \tag{2}$$

[2]A semi-metric does not require the property $d(x, x') = 0$ if and only if $x = x'$ to hold.

The geodesic metric $d_X(x, x')$ is given as the minimum over the set of all *admissible* paths between x and x', denoted by $\Gamma(x, x')$. A path is admissible if the edges of all the infinitesimal segments x_i, x_{i+1} are connected. We defer a more precise definition to the next section. We broadly refer to properties described in terms of $d_{\mathbb{E}}$ as the *extrinsic geometry* of X, and to properties associated with d_X as the *intrinsic geometry* of X.

In order to determine whether two shapes X and Y are similar, we compare them as metric spaces. From the point of view of metric geometry, two metric spaces are equivalent if their corresponding metric structures are equal. Such metric spaces are said to be *isometric*. More formally, a bijective map $f : (X, d) \to (Y, \delta)$ is called an *isometry* if

$$\delta \circ (f \times f) = d. \tag{3}$$

In other words, an isometry is a metric-preserving map between two metric spaces, such that $d(x_1, x_2) = \delta(f(x_1), f(x_2))$. We call such (X, d) and (Y, δ) *isometric* and denote this property by $(X, d) \sim (Y, \delta)$.

This definition of equivalence obviously depends on the choice of the metric. A bijection $f : (X, d_X) \to (Y, d_Y)$ satisfying $d_Y \circ (f \times f) = d_X$ is called an *intrinsic isometry*. Saying that (X, d_X) and (Y, d_Y) are isometric is synonymous to X and Y being *intrinsically equivalent*. On the other hand, if we consider the extrinsic geometry of the shapes (i.e., look at the shapes endowed with the Euclidean rather than geodesic metric), we notice that $(X, d_{\mathbb{E}})$ and $(Y, d_{\mathbb{E}})$ are subsets of the same metric space $(\mathbb{E}, d_{\mathbb{E}})$. As a result, an extrinsic isometry is a bijection between subsets of the Euclidean space rather than between two different metric spaces. In Euclidean geometry, the only possible isometries are rigid motions, which include rotation, translation and reflection transformations; we denote the family of such transformations by $\mathrm{Iso}(\mathbb{E})$. Thus, X and Y are extrinsically isometric if there exists $f \in \mathrm{Iso}(\mathbb{E})$ such that $d_{\mathbb{E}} = d_{\mathbb{E}} \circ (f \times f)$ on $X \times X$. This means that two shapes are extrinsically isometric if one can be obtained by a rigid transformation of the other, which is often expressed by saying that X and Y are *congruent*.

To avoid confusion, in the following, we say that X and Y are isometric

implying intrinsic isometry, and that X and Y are congruent when referring to extrinsic isometry. The class of intrinsic isometries is usually richer than that of congruences, since any congruence is by definition also an intrinsic isometry. However, for some objects these two classes coincide, meaning that they have no incongruent isometries. Such shapes are called *rigid*, and their extrinsic geometry is completely defined by the intrinsic one.

2.1 Similarity

In practice, perfect equivalence rarely exists, and we are usually limited to speaking about *similarity* of shapes rather than their equivalence in the strict sense. In order to account for this, we need to relax the notion of isometry. Two metric spaces (X, d) and (Y, δ) are said to be ϵ-*isometric* if there exists an ϵ-*surjective* map $f : (X, d) \to (Y, \delta)$ (i.e., $\delta(y, f(X)) \leq \epsilon$ for all $y \in Y$), which has the *distortion*

$$\operatorname{dis} f \quad = \quad \sup_{x,x' \in X} \left| d(x, x') - \delta(f(x), f(x')) \right| \leq \epsilon. \tag{4}$$

Such an f is called an ϵ-*isometry*. ϵ-isometries are quite different from their true counterparts. Particularly, an isometry is always bi-Lipschitz continuous [21], which is not necessarily true for an ϵ-isometry. If we further relax the requirement of ϵ-surjectivity by demanding that f has only $\operatorname{dis} f \leq \epsilon$, we refer to such f as an ϵ-*isometric embedding*.

A way to quantify the degree of shape dissimilarity is by defining a *shape distance* $d_{\text{shape}} : \mathbb{S} \times \mathbb{S} \to \mathbb{R}_+$ on the space of shapes $\mathbb{S}$. Note that d_{shape} is a function of (X, d) and (Y, δ). In the following, we will write $d_{\text{shape}}(X, Y)$ omitting explicit reference to the metrics for notation brevity. It is common to require d_{shape} to satisfy the following set of properties for any X, Y, and Z in $\mathbb{S}$ and a constant $c \geq 1$ independent of X, Y and Z:

(I1) *Equivalence:* $d_{\text{shape}}(X, Y) = 0$ if and only if (X, d) and (Y, δ) are isometric.

(I2) *Similarity:* if $d_{\text{shape}}(X, Y) \leq \epsilon$, then (X, d) and (Y, δ) are $c\epsilon$-isometric; and if (X, d) and (Y, δ) are ϵ-isometric, then $d_{\text{shape}}(X, Y) \leq c\epsilon$.

(I3) *Symmetry:* $d_{\text{shape}}(X, Y) = d_{\text{shape}}(Y, X)$.

(I4) *Triangle inequality:* $d_{\text{shape}}(X, Z) \leq d_{\text{shape}}(X, Y) + d_{\text{shape}}(Y, Z)$.

The first two properties guarantee that the shape distance d_{shape} reflects the degree of shape dissimilarity, i.e. d_{shape} is small for similar (isometric) shapes, and large for dissimilar ones. The third property is synonymous to the *reflexivity* of the shape similarity relation, while the fourth reflects its *transitivity*: if X is similar to Y and Y is similar to Z, then X and Z cannot be too much dissimilar. Properties (I1), (I3), and (I4) can be expressed equivalently by saying that d_{shape} is a metric on the quotient space of $\mathbb{S}$ under the isometric equivalence relation, denoted by $\mathbb{S}\backslash\sim$. We refer to a shape distance satisfying (I1)–(I4) as to *isometry-invariant shape distance*. Note, however, that in the case of a *partial* shape similarity relation, metric axioms are usually too restrictive [41, 8].

Since the definitions of equivalence and similarity depend on the choice of the metrics on the shapes, our notion of isometry-invariant distance between shapes also depends on them. In the remainder of this section, we will consider two particular cases of extrinsic and intrinsic shape distances, defining, respectively, the extrinsic and intrinsic similarity.

2.2 Extrinsic distances

Extrinsic similarity is a simple case of the general problem of metric space comparison, since two shapes with the Euclidean metric, $(X, d_{\mathbb{E}})$ and $(Y, d_{\mathbb{E}})$, are a subset of the same metric space $(\mathbb{E}, d_{\mathbb{E}})$. Consequently, we can use the *Hausdorff distance* measuring the distance between two sets X and Y in $\mathbb{E}$,

$$d_{\text{H}}^{\mathbb{E}}(X, Y) \quad = \quad \max\left\{\sup_{x\in X} d_{\mathbb{E}}(x, Y), \sup_{y\in Y} d_{\mathbb{E}}(y, X)\right\}, \tag{5}$$

where $d_{\mathbb{E}}(y, X) = \inf_{x\in X} \|y - x\|_2$ denotes the distance between the set X and the point y. In practice, a non-symmetric version of $d_{\text{H}}^{\mathbb{E}}$,

$$d_{\text{NH}}^{\mathbb{E}}(X, Y) \quad = \quad \sup_{x\in X} d_{\mathbb{E}}(x, Y), \tag{6}$$

is often preferred since it allows for partial comparison of surfaces. L_2 approximations are also often preferred, as the original L_∞ formulation is sensitive to outliers.

Extrinsic equivalence of shapes implies that they are congruent, i.e., can be brought into ideal correspondence by means of a rigid transformation. In other words, there exists $f \in \mathrm{Iso}(\mathbb{E})$ such that $d_{\mathrm{H}}^{\mathbb{E}}(f(X), Y) = 0$. We can use the same idea in order to measure extrinsic similarity: find the minimum possible value of the Hausdorff distance over all possible rigid motions,

$$d_{\mathrm{ext}}(X, Y) \quad = \quad \inf_{f \in \mathrm{Iso}(\mathbb{E})} d_{\mathrm{H}}^{\mathbb{E}}(f(X), Y) \tag{7}$$

Since $\mathrm{Iso}(\mathbb{E})$ can be easily parametrized, the resulting extrinsic distance can be computed as

$$d_{\mathrm{ext}}(X, Y) \quad = \quad \inf_{R,t} d_{\mathrm{H}}^{\mathbb{E}}(RX + t, Y), \tag{8}$$

where R denotes the rotation matrix and t is the translation vector.[3] Practical methods to solve problem (7) are the ICP algorithms, which use an alternating minimization scheme consisting of two stages. First, the closest-point correspondence between X and Y is computed. Next, the rigid transformation is found between X and Y that minimizes the Hausdorff distance. The process is repeated until convergence. Though theoretically a global minimum should be searched for, in practice, local optimization methods are usually employed, and thus, ICP algorithms are often prone to convergence to a suboptimal solution (local minimum).

The distance d_{ext} satisfies Properties (I1)–(I4), being a metric on the space of shapes modulo $\mathrm{Iso}(\mathbb{E})$, and is a good way to measure extrinsic similarity.

2.3 Intrinsic distances

Computation of intrinsic shape dissimilarity is a more complex task compared to its extrinsic counterpart due to two main reasons. First, intrinsic isometries usually constitute a significantly richer class of transformations

[3] Reflection is usually excluded as having no physical realization in $\mathbb{E}$.

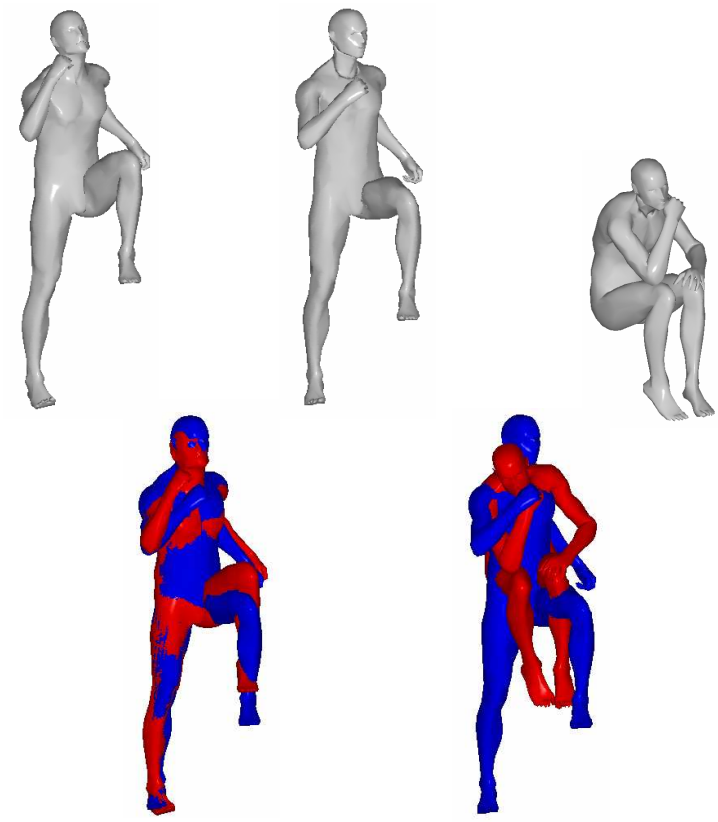


Figure 4: The weakness of extrinsic similarity. Top row shown deformations of a three-dimensional shape of the nonrigid human body. Weak deformations (left and middle) can be compared using ICP algorithm, which finds a meaningful alignment between the shapes (bottom row, left), despite topological noise (simulated here by welding hand and leg at a point denoted by a red circle). However, in the case of strong deformations (top row, right) ICP produces meaningless matching (bottom row, right).

than Euclidean congruences, and there often exists no simple parametrization comparable to the six degrees of freedom that are sufficient to represent a rigid motion. Second, unlike in the rigid case where $(X, d_{\mathbb{E}})$ and $(Y, d_{\mathbb{E}})$ are compared as subsets of a common metric space $(\mathbb{E}, d_{\mathbb{E}})$, the comparison of (X, d_X) and (Y, d_Y) does not allow using the Hausdorff distance, and hence d_{ext} in (7) cannot be straightforwardly generalized to the intrinsic case.

In [29], Elad and Kimmel reduced the problem of intrinsic shape similarity to the more tractable problem of extrinsic similarity, by first computing an *extrinsic representation* of the intrinsic geometry of the shapes in some common metric space $(\mathbb{Q}, d_{\mathbb{Q}})$. Such a representation, dubbed as the *canonical form*, is constructed by minimum-distortion embedding, attempting to find two maps $\varphi : (X, d_X) \to (\mathbb{Q}, d_{\mathbb{Q}})$ and $\psi : (Y, d_Y) \to (\mathbb{Q}, d_{\mathbb{Q}})$ with minimum distortions dis φ and dis ψ. The embedding, in a sense, allows to "undo" all the isometric deformations of the shapes (though, some degree of ambiguity stemming from isometries in $\mathbb{Q}$ still remains). Once the canonical forms $\varphi(X)$ and $\psi(Y)$ are computed, they are compared extrinsically using, for example, ICP. In other words, the intrinsic distance between two shapes X and Y is computed as the extrinsic distance between their canonical forms,

$$d_{\text{int}}(X, Y) \quad = \quad d_{\text{ext}}(\varphi(X), \psi(Y)). \tag{9}$$

If the Euclidean space $(\mathbb{E}^m, d_{\mathbb{E}^m})$ is used as the embedding space, the minimum distortion embedding of (X, d_X) can be found by solving a *multidimensional scaling* (MDS) problem. Given the shape X sampled at points $\hat{X} = \{x_1, ..., x_N\}$ and the $N \times N$ matrix of distances $\mathbf{D}_{\hat{X}} = (d_X(x_i, x_j))$, MDS algorithms try to find a configuration of points in $\mathbb{R}^m$ such that the Euclidean distances between these points are as close as possible in some sense to the elements of $\mathbf{D}_{\hat{X}}$. Typically, the L_2 distortion criterion is used, in which case the MDS problem is referred to as *least-squares* MDS (LSMDS). The canonical form is given by

$$Z \quad = \quad \underset{Z \in \mathbb{R}^{N \times m}}{\operatorname{argmin}} \sum_{i>j} |\|x_i - x_j\|_2 - d_X(x_i, x_j)|^2, \tag{10}$$

where Z is an $N \times m$ matrix representing the coordinates of the points

in $\mathbb{E}^m$. There exist numerically simple and efficient algorithms for solving problem (10) [6, 19].

The choice of $\mathbb{Q}$ has an important influence on the computation of canonical forms. Though the Euclidean space is the simplest and most convenient metric space for this purpose, other choices are possible as well [28, 75, 16]. There are a few criteria for choosing the embedding space. First, it is desired that $\mathbb{Q}$ is homogeneous and its isometry group is simple, in order to reduce the number of degrees of freedom in the definition of the canonical forms. Secondly, an analytic expression for $d_{\mathbb{Q}}$ allows using MDS algorithms [6]. Finally, it appears that for some classes of shapes, certain embedding spaces are generally more suitable (see, for example, experimental results in [16]).

However, achieving a true isometric embedding (i.e., the minimum in problem (10) equal to zero) is usually impossible [53] with a metric space $\mathbb{Q}$ satisfying the above criteria. Hence, the canonical forms are only an approximate representation of the intrinsic geometry of the shapes. The problem of inaccuracy introduced by the embedding into $\mathbb{Q}$ can be resolved if we do not assume a given embedding space, but instead, include $\mathbb{Q}$ as a variable into the optimization problem. We can always find a sufficiently complicated metric space into which both X and Y can be embedded isometrically, and compare the images using the Hausdorff distance,

$$d_{\mathrm{GH}}((X, d_X), (Y, d_Y)) = \inf_{\substack{\mathbb{Q} \\ \varphi: X \to \mathbb{Q} \\ \psi: Y \to \mathbb{Q}}} d_{\mathrm{H}}^{\mathbb{Q}}(\varphi(X), \psi(Y)), \tag{11}$$

(here φ and ψ are assumed to be isometric embeddings). The resulting distance is referred to as *Gromov-Hausdorff distance* [34]. In the following, we use a brief notation $d_{\mathrm{GH}}(X, Y)$ when the implied metrics are clear.

While the computation of d_{GH} as defined in (11) is hardly tractable due to the minimization over all embedding spaces $\mathbb{Q}$, it appears that for compact surfaces, the Gromov-Hausdorff distance can be expressed in terms of the distortion obtained by embedding one surface into the other,

$$d_{\mathrm{GH}}(X, Y) = \frac{1}{2} \inf_{\substack{\varphi: X \to Y \\ \psi: X \to Y}} \max\{\mathrm{dis}\,\varphi, \mathrm{dis}\,\psi, \mathrm{dis}\,(\varphi, \psi)\}, \tag{12}$$

where,

$$\operatorname{dis}(\varphi,\psi) \quad = \quad \sup_{x\in X, y\in Y} |d_X(x,\psi(y)) - d_Y(y,\varphi(x))|. \tag{13}$$

The computation of the distortions is performed using the generalized multidimensional scaling [13, 12], a procedure similar in its spirit to MDS, but not limited to spaces with analytically expressed geodesic distances. Using the Gromov-Hausdorff distance, the intrinsic shape distance can be computed as

$$d_{\mathrm{int}}(X,Y) \quad = \quad d_{\mathrm{GH}}((X,d_X),(Y,d_Y)). \tag{14}$$

2.4 A unified view of extrinsic and intrinsic similarity

It is worthwhile to note that the Gromov-Hausdorff distance is a generic isometry-invariant shape distance satisfying properties (I1)–(I4) with $c = 2$. Particularly, this implies that if $d_{\mathrm{GH}}((X,d),(Y,\delta)) \leq \epsilon$, then (X,d) and (Y,δ) are 2ϵ-isometric and conversely, if (X,d) and (Y,δ) are ϵ-isometric, then $d_{\mathrm{GH}}((X,d),(Y,\delta)) \leq 2\epsilon$ [21, 12]. One can think of the Gromov-Hausdorff distance as a general framework for isometry-invariant shape comparison, unifying both extrinsic and intrinsic similarity. Selecting d and δ as the geodesic metrics (d_X and d_Y, respectively), we obtain the intrinsic shape distance d_{int}. Considering the shapes with the Euclidean metric $d_{\mathbb{E}}$ yields an equivalent (but not equal) way to express the extrinsic shape distance as

$$d_{\mathrm{ext}}(X,Y) \quad = \quad d_{\mathrm{GH}}((X,d_{\mathbb{E}}),(Y,d_{\mathbb{E}})), \tag{15}$$

instead of the definition we have seen before. In this case, we interpret isometry as congruence, while in the former case, it implies the existence of a geodesic distance-preserving deformation of X into Y.

3 Topology-invariant similarity

Let us now return to our example of glove fitting. Assume that Y is the hand surface, and X is the glove we wish to fit. A perfect fit is achieved when

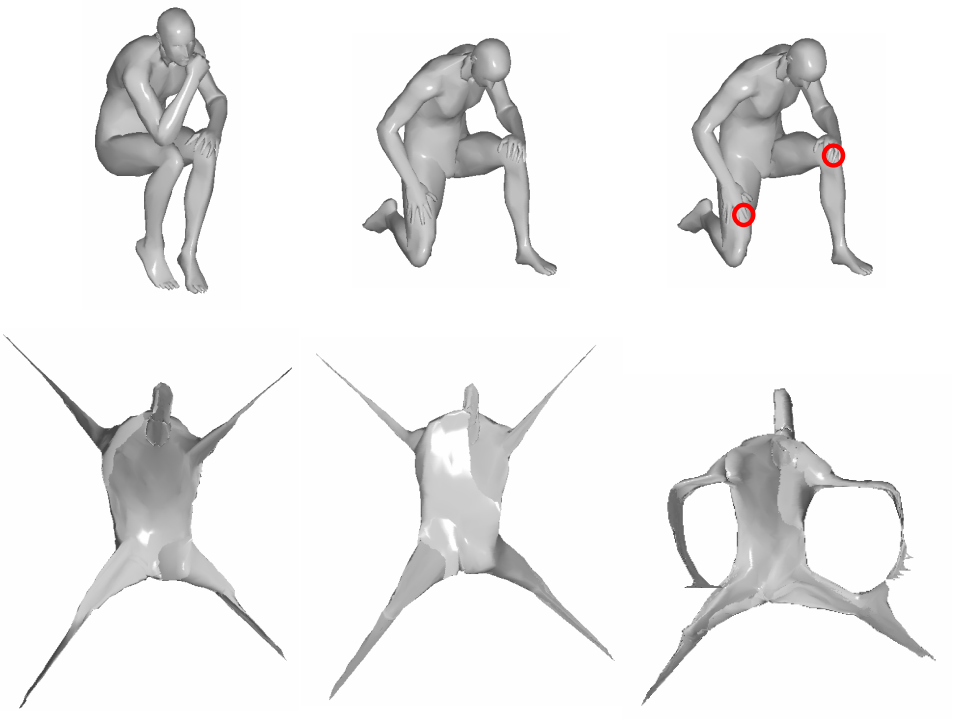


Figure 5: The weakness of intrinsic similarity. Shown are deformations of a three-dimensional nonrigid human body shape (top row) and the corresponding canonical forms (bottom row) computed using classical MDS. The canonical forms appear to be insensitive to near-isometric deformations of the shape (left and middle columns). However, topological noise (simulated here by welding the hands to the legs at points indicated by red circles) results in a completely different canonical form.

X and Y are equivalent. Since both the glove and the hand are nonrigid shapes, we understand equivalence in the intrinsic sense, i.e., as the existence of an isometry between (X, d_X) and (Y, d_Y). An alternative way to express this fact is by saying that there exists an isometric deformation f of X, such that $f(X) = Y$. In other words, we may say that the glove X perfectly fits the hand Y if there exists Z intrinsically equivalent to X and extrinsically equivalent to Y.

Let us now assume that the hand is given in a posture where the fingers are glued, while the fingers of the glove are disconnected. Such a topological difference would make distinct the intrinsic geometries of the hand and the glove, preventing X and Y from being intrinsically equivalent (for an illustration, see Figure 5). However, our alternative definition of equivalence would still hold, as we can still find Z intrinsically equivalent to X, which will fit Y extrinsically (though, unlike the previous case, Z and Y are no longer intrinsically equivalent).

This example visualizes the limitations of the notion of intrinsic equivalence of shapes, and brings forth the need to construct a more general notion of equivalence, insensitive to topological changes. Such a construction requires several additions to our mathematical machinery. We define a *topology* T of X as a family of subsets of X (i.e., T is a subset of the power set 2^X) closed under finite intersection and union of arbitrarily many elements of T. By definition, both the empty set and X itself are members of T. Conventionally, a subset of X is called *open* if it belongs to T, and *closed* if its complement belongs to T. From now on, when speaking about a shape X, we will actually imply the triplet (X, T, d), and omitting T will imply that the topology is induced by the metric. For example, for the choice $d = d_{\mathbb{E}}$, the topology of the Euclidean space $\mathbb{E}$ is assumed. Topology defines properties of the shape which are "coarser" than metric geometry, such as connectivity of points.

Our glove fitting example suggests that there may exist two shapes (X, T, d_X) and (X, T', d'_X) with identical realization X in the Euclidean space, yet different topology and intrinsic geometry. This is possible due to the fact that our definition of the geodesic metric in (1) was based on the

length of the shortest admissible path $\gamma \in \Gamma(x, x')$ connecting two points x and x' on X.

However, the definition of the set of all admissible paths connecting x and x' is itself a function of the topology, as it depends on the connectivity. For example, in case of a hand with glued fingers, there exists a path between two finger tips crossing the fingers. When the fingers are disconnected, the shortest path between the finger tips goes along the fingers, since the path directly connecting the finger tips is no longer admissible (for a more formal discussion, the reader is referred to Chapter 2 in [21]).

Two topological spaces (X, T) and (Y, T') are said to be *homeomorphic* or *topologically equivalent* if there exists a continuous bijection $g : (X, T) \to (Y, T')$, whose inverse is also continuous (the continuity of g and g^{-1} is understood in the sense of T and T', respectively). Since isometries are also homeomorphisms, it follows straightforwardly that both T and d remain invariant under an isometry.

In order to distinguish between purely geometric and purely topological deformations of a shape, we will say that a map $f : (X, T, d) \to (Y, T', \delta)$ is a *topological deformation* if it leaves X unchanged (i.e., $X = Y$)[4]. Using this notion, we can finally define two shapes (X, T_X, d_X) and (Y, T_Y, d_Y) to be equivalent in the sense of topology-invariant similarity if there exists a topological deformation $g : (X, T_X, d_X) \to (X, T'_X, d'_X)$, and an isometry $f : (X, T'_X, d'_X) \to (Z, T'_X, d'_X)$ such that $Y = Z = (f \circ g)(X)$. Note that though we gave a simple example of topological deformations that change point-wise connectivity, other, more generic deformations (such as opening holes) are also possible.

Such a definition of topology-invariant equivalence is more general than extrinsic or intrinsic similarity we had before. In our example, a glove (Y, T_Y, d_Y) shown in Figure 2 (left) and a hand with glued fingers (X, T_X, d_X) shown in Figure 2 (right) are extrinsically dissimilar since X and Y are incongruent. On the other hand, (X, T_X, d_X) and (Y, T_Y, d_Y) are intrinsically dissimilar since (X, T_X) and (Y, T_Y) are not topologically equivalent and as a result, (X, d_X) and (Y, d_Y) are not isometric. Yet, using the above crite-

[4]Or more generally, X and Y are congruent.

rion, there is a topology-invariant equivalence between shapes (X, T_X, d_X) and (Y, T_Y, d_Y): we can first "unglue" the fingers of the hand by means of a topological deformation $g : (X, T_X, d_X) \to (X, T'_X, d'_X)$, and then bend the hand by means of an isometry $f : (X, T'_X, d'_X) \to (Z, T_Z, d_Z)$, such that $Z = Y$, T'_X is equivalent to T_Z, and d'_X is equivalent to d_Z.

In order to relax this notion of equivalence into a topology-invariant similarity relation, a slightly more complicated construction is required. We will say that (X, T_X, d_X) and (Y, T_Y, d_Y) are $(\epsilon_{\text{int}}, \epsilon_{\text{ext}})$-*similar* if there exists a topological deformation $g : (X, T_X, d_X) \to (X, T'_X, d'_X)$, and an ϵ_{int}-isometry $f : (X, T'_X, d'_X) \to (Z, T'_X, d'_X)$ such that $d_{\text{H}}^{\mathbb{E}}((f \circ g)(X), Y) \leq \epsilon_{\text{ext}}$.

In many practical situations, an *asymmetric* definition of topology-invariant similarity is sufficient. In such cases, we distinguish between a *probe* X that is fit to a *model* Y. We say that (X, T_X, d_X) and (Y, T_Y, d_Y) are $(\epsilon_{\text{int}}, \epsilon_{\text{ext}})$-similar if there exists an ϵ_{int}-isometry $f : (X, T'_X, d'_X) \to (Z, T'_X, d'_X)$ such that $d_{\text{H}}^{\mathbb{E}}((f \circ g)(X), Y) \leq \epsilon_{\text{ext}}$. In what follows, we will limit our attention to the latter definition of similarity, and will construct a distance based on a combination of extrinsic and intrinsic distances to reflect it.

3.1 Topology-invariant distance

Another way to express our notion of topology-invariant similarity is by saying that the probe X admits a topology-preserving deformation Z such that $d_{\text{int}}(X, Z) \leq \epsilon_{\text{int}}$, and $d_{\text{ext}}(Z, Y) \leq \epsilon_{\text{ext}}$. Since it is usually impossible to say which of the two criteria is more important, we judge the similarity as a *tradeoff* between them. Such a joint intrinsic-extrinsic similarity between two shapes can be quantitatively represented by the extent to which we have to modify the extrinsic geometry in order to make the two shapes intrinsically similar, or alternatively, the extent to which we have to modify the intrinsic geometry in order to make the two shapes extrinsically similar. This, in turn, can be formulated as a multicriterion optimization problem, in which we bring to minimum the vector objective function $\Phi(Z) = (d_{\text{int}}(Z, X), d_{\text{ext}}(Z, Y))$ with respect to Z.

Unlike optimization with a scalar objective, we cannot define unambiguously the minimum of Φ, since there does not exist a total order relation

between the criteria – we cannot say, for example, whether it is better to have $\Phi = (0.5, 1)$ or $\Phi = (1, 0.5)$. At the same time, we have no doubt that $\Phi = (0.5, 0.5)$ is better than $\Phi = (1, 1)$, since both criteria have smaller values. This allows us to define a minimizer of our vector objective Φ as a Z^*, such that there is no other Z for which $\Phi(Z)$ is better that $\Phi(Z^*)$ (which can be expressed as a vector inequality $\Phi(Z^*) > \Phi(Z)$). Such an Z^* is called *non-inferior* or *Pareto optimal*. Pareto optimum is not unique; we denote by Ω^* the set of all Z^* that satisfy the above relation. The corresponding values of $\Phi(\Omega^*)$ are referred to as a *Pareto frontier* and can be visualized as a planar curve (Figure 6).

In [8], Bruckstein and the authors proposed considering the entire Pareto frontier as a criterion of similarity, of which we can think as a generalized, *set-valued distance*.[5] In our case, this distance measures the tradeoff between the intrinsic and extrinsic similarity, which we denote by $d_{\mathrm{P}}(X, Y) = \Phi(\Omega^*)$ and refer to as the *Pareto distance*. The Pareto distance quantifies the degree of asymmetric topology-invariant dissimilarity; particularly, $(0, 0) \in d_{\mathrm{P}}(X, Y)$ if and only if X and Y are equivalent in the sense that X has an isometry congruent to Y. The Pareto distance also generalizes the similarity criteria based on purely extrinsic or intrinsic geometry. Asserting $d_{\mathrm{int}}(X, Z) = 0$, the other criterion $d_{\mathrm{ext}}(Z, Y)$ in $d_{\mathrm{P}}(X, Y)$ will measure how close a perfectly isometric deformation of X can bring X to Y. Since in practice we work with meshes which are known to be almost always rigid, the only possible deformations are congruences of X. This means that with $d_{\mathrm{int}}(X, Z)$ fixed to zero, $d_{\mathrm{P}}(X, Y)$ is equivalent to ICP. On the other hand, if we require $d_{\mathrm{ext}}(Y, Z) = 0$, the probe is forced to be attached to the model surface, which boils down to a GMDS problem, where we are trying to find a deformation of X minimizing $d_{\mathrm{int}}(Z, Y)$. We conclude that the two extreme cases of $d_{\mathrm{P}}(X, Y)$ are $(0, d_{\mathrm{ext}}(X, Y))$ and $(d_{\mathrm{int}}(X, Y), 0)$. Other points on the Pareto frontier represent different tradeoffs between the extrinsic and

[5] Set-valued distances arise from similarity relations based on more than one criterion such as in the case of partial similarity, which can be regarded as the tradeoff between similarity and significance of the parts. We believe that it is more natural to use set-valued distances for such relations, rather than trying to fit them into the Procrustean bed of scalar-valued similarity.

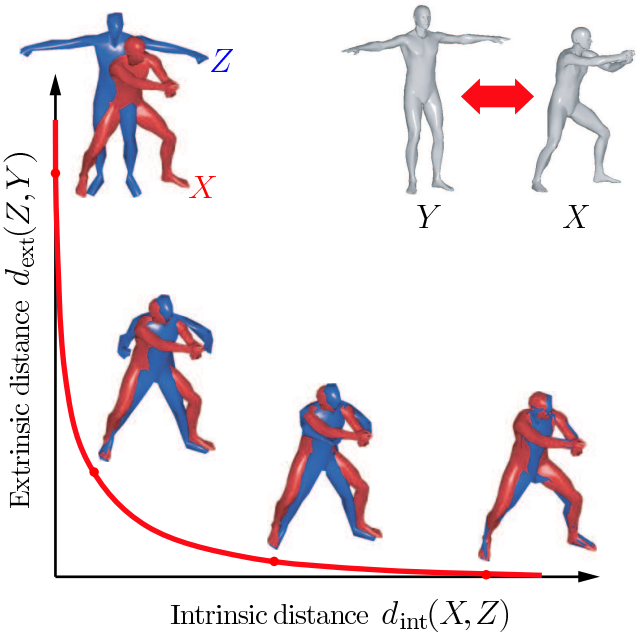


Figure 6: Visualization of the concept of set-valued distance. Different points on the Pareto frontier represent different tradeoffs between intrinsic and extrinsic similarity.

the intrinsic criteria, as depicted in Figure 6.

In general, by saying that $d_{\text{P}}(X, Y) \leq (\epsilon_{\text{int}}, \epsilon_{\text{ext}})$ (the vector inequality implies that the point $(\epsilon_{\text{int}}, \epsilon_{\text{ext}})$ is above or on the Pareto frontier), we mean that there exists a deformation Z of X distorting its intrinsic geometry by less than ϵ_{int}, while making its extrinsic geometry less than ϵ_{ext}-dissimilar from that of Y. Saying that $d_{\text{P}}(X, Y) < d_{\text{P}}(X, Y')$, implies that bringing X to a certain proximity of Y requires a smaller intrinsic distortion than bringing X to the same proximity of Y'. Using our glove fitting example, we would say that the glove X fits the hand Y better than the hand Y'. Geometrically, this fact is manifested by the curve $d_{\text{P}}(X, Y)$ being entirely below $d_{\text{P}}(X, Y')$.

3.2 Joint intrinsic and extrinsic similarity

The set-valued distance d_{P} may be inconvenient to use since only a partial order relation exists between the Pareto frontiers – given two set-valued distances, we cannot in general say which of them is "smaller", unless one curve is entirely above or below the other. In order to compare shape distances, we have to convert them into a single scalar value.

Intuitively, the "speed of decay" of the Pareto frontier indicates how similar two shape are (the faster, the more similar). There are multiple ways to represent this information as a single number. First, one can measure the area under the Pareto frontier. Smaller area corresponds to higher similarity. Second, it is possible to select a single point on the Pareto frontier by fixing a pre-set value of one of the dissimilarities. For example, if we know that the shapes are inelastic to a certain degree, i.e. they can stretch or shrink by no more than ϵ, we will fix the intrinsic dissimilarity $d_{\mathrm{int}} = \epsilon$ and will use the value extrinsic dissimilarity at this point as the shape distance. A third alternative is to require that both dissimilarities are equal and compute the point at which $d_{\mathrm{int}} = d_{\mathrm{ext}}$. This way, we obtain a shape distance similar in its spirit to the *equal error rate* (EER) used in receiver operating characteristic analysis.

A more generic approach to converting a set-valued distance into a scalar-valued one comes from the multicriterion optimization theory [66]. Among all the Pareto optima, we cannot prefer any since they are non-comparable: we cannot say which Pareto optimum is better. However, ideally we want to bring both of our criteria to zero, that is, achieve the "utopia point" $(0, 0)$. We can choose a single point on the Pareto frontier the closest to the utopia point, in the sense of some distance.

Using this idea, we define the *joint distance* as

$$d_{\mathrm{joint}}(X, Y) \quad = \quad \inf_{(\epsilon_{\mathrm{int}}, \epsilon_{\mathrm{ext}}) \in d_{\mathrm{P}}(X,Y)} \|(\epsilon_{\mathrm{int}}, \epsilon_{\mathrm{ext}})\|_{\mathbb{R}^2_+}, \tag{16}$$

where $\|\cdot\|_{\mathbb{R}^2_+}$ is some norm on $\mathbb{R}^2_+$. In the following, we will consider a family of norms $\|(\epsilon_{\mathrm{int}}, \epsilon_{\mathrm{ext}})\|_\lambda = \epsilon_{\mathrm{int}} + \lambda\epsilon_{\mathrm{ext}}$, for $\lambda > 0$. The joint distance in this

case can be written as

$$d_{\text{joint}}(X, Y) \quad = \quad \min_{Z} \, d_{\text{int}}(X, Z) + \lambda d_{\text{ext}}(Z, Y). \tag{17}$$

Different selections of the multiplier λ attribute importance either to d_{int} or d_{ext}, and give us different points on the Pareto frontier.

3.3 Simplification

The intrinsic distance term $d_{\text{int}}(X, Z)$ can be simplified by observing that the deformation $Z = f(X)$ gives a one-to-one correspondence between X and Z. We can therefore fix φ and ψ in the Gromov-Hausdorff distance to f and f^{-1}, respectively, obtaining the following intrinsic distance,

$$d_{\text{int}}(X, Z) \quad = \quad \text{dis } f \quad = \quad \sup_{x,x' \in X} |d_X(x, x') - d_Z(f(x), f(x'))|. \tag{18}$$

Note that the correspondence between the surfaces is now fixed and does not participate anymore in the minimization. $d_{\text{int}}(X, Z)$ defined this way measures the distortion in the intrinsic geometry of X introduced by the extrinsic deformation f. The main difficulty in its computation stems from the fact the the geodesic metric d_Z has to be re-computed every time the deformation changes. We will defer the discussion of this issue to the following section.

A simplification of the extrinsic distance term $d_{\text{ext}}(Z, Y)$ in $d_{\text{joint}}(X, Y)$ is possible due to the fact that the deformation Z of X already accounts for all possible congruences. Consequently, there is no need to minimize over $\text{Iso}(\mathbb{E})$ when computing $d_{\text{ext}}(Z, Y)$ – we can simply use $d_{\text{H}}^{\mathbb{E}}(Y, Z)$ or $d_{\text{NH}}^{\mathbb{E}}(Y, Z)$ instead. Like in ICP, the main computational challenge is the need to re-compute the set of closest points every time the deformation changes.

4 Numerical framework

For practical computations, we work with discretized shapes. The surface X is sampled at N points $\hat{X} = \{x_1, ..., x_N\} \subseteq X$, constituting an *r-covering*,

i.e., $X = \bigcup_{n=1}^{N} B_X(x_n, r)$ (here, $B_X(x, r)$ is a metric ball of radius r around x). The extrinsic coordinates of $\hat{X}$ can be represented as an $N \times 3$ matrix $\mathbf{X}$, each row of which corresponds to $x_i \in \mathbb{E}$. The discrete shape is represented as a triangular mesh; each triangle is a triplet of indices of vertices belonging to it. The maximum length of an edge is r. Vertices connected by an edge (or in other words, belonging to the same triangle) are said to be *adjacent*; we describe the adjacency by the set E of all adjacent pairs of vertices in $\hat{X}$. The geodesic distances on $\hat{X}$ are approximated using the fast marching method [46] or the Dijkstra's algorithm, forming an $N \times N$ matrix $\mathbf{D}_{\hat{X}}$.

Assuming the deformed surface $\hat{Z} = f(\hat{X})$ maintains the connectivity of $\hat{X}$, we can formulate the following minimization problem with respect to the $N \times 3$ matrix $\mathbf{Z}$ of the extrinsic coordinates of $\hat{Z}$:

$$d_{\text{joint}}(\hat{X}, \hat{Y}) \quad = \quad \min_{X} \frac{1}{N^2} \sum_{i,j=1}^{N} (d_{\hat{X}}(x_i, x_j) - d_{ij}(\mathbf{Z}))^2 + \frac{\lambda}{N} \sum_{i=1}^{N} d^2(z_i, \hat{Y}) \tag{19}$$

where $d_{ij}(\mathbf{Z}) = d_{\hat{Z}}(z_i, z_j)$ denote the geodesic distances on $\hat{Z}$, and $d(z_i, \hat{Y})$ denotes the Euclidean distance from the point z_i to the discretized surface $\hat{Y}$. The first term of the above cost function is the discretization of d_{int}, whereas the second term is the discretization of d_{ext}. In the sequel, we show how to efficiently compute these two terms and their derivatives with respect to $\mathbf{Z}$, required for the minimization of (19).

4.1 Intrinsic distance computation

The main challenge in the computation of the intrinsic distance term d_{int} is the need to evaluate the geodesic distances on $\hat{Z}$ and their derivatives with respect to the extrinsic geometry of $\hat{Z}$ changing at each iteration of the minimization algorithm. The simplest remedy would be to modify the intrinsic term by restricting i and j to the neighboring points only,

$$d_{\text{int}} \quad \approx \quad \frac{1}{|E|} \sum_{(i,j) \in E} \left(d_{\hat{X}}(x_i, x_j) - d_{\hat{Z}}(z_i, z_j) \right)^2 . \tag{20}$$

This way, we use only the local distances on $\hat{Z}$, which can be approximated as the Euclidean distances $d_{\hat{Z}}(z_i, z_j) = \|z_i - z_j\|$. However, such a modification

makes d_{int} significantly less sensitive to large deformations of $\hat{X}$. Indeed, many deformations change the local distance only slightly, while introducing large distortion to larger distances.

In order to penalize for such deformations of $\hat{X}$, we need to approximate the full matrix of geodesic distances on $\hat{Z}$. We first define the matrix $\mathbf{D}(\mathbf{Z})$ of *local distances*, whose elements are, as before,

$$d_{ij}(\mathbf{Z}) \quad = \quad \begin{cases} \|z_i - z_j\| & : \quad (i,j) \in E \\ 0 & : \quad (i,j) \notin E. \end{cases} \tag{21}$$

Using the Dijkstra algorithm,[6] we compute the set of shortest paths between all pairs of points (i,j). For example, let $\mathcal{P}_{ij} = \{(i,i_1),(i_1,i_2),...,(i_{n-1},i_n),(i_n,j)\} \subset E$ be the shortest path between the points i and j. Its length is given by $L(\mathcal{P}_{ij}) = d_{i,i_1} + d_{i_1,i_2} + ... + d_{i_n,j}$, which is a linear combination of the elements of $\mathbf{D}(\mathbf{Z})$. We can therefore "complete" the missing entries in the matrix $\mathbf{D}(\mathbf{Z})$ by defining the matrix of *global distances*

$$\hat{\mathbf{D}}(\mathbf{Z}) \quad = \quad \mathcal{I}(\mathbf{Z})\mathbf{D}(\mathbf{Z}), \tag{22}$$

where $\mathcal{I}$ is a sparse fourth order tensor, with the elements $\mathcal{I}_{ijkl} = 1$ if the edge (k,l) is contained in the shortest path $\mathcal{P}_{ij}$, and 0 otherwise. Note that $\mathcal{I}$ depends on the connectivity E, which is assumed to be fixed, and the matrix of local distances $\mathbf{D}$, which, in turn, depends on $\mathbf{Z}$.

It is straightforward to verify that the entries of $\hat{\mathbf{D}}(\mathbf{Z})$ and $\mathbf{D}(\mathbf{Z})$ coincide for all $(i,j) \in E$. The global distance matrix $\hat{\mathbf{D}}(\mathbf{Z})$ constitutes an approximation of the geodesic distances on the surface $\hat{Z}$, $\hat{d}_{ij} \approx d_{\hat{Z}}(z_i, z_j)$, while having a simple linear form in terms of the local distances. A consistent and more accurate estimate of $d_{\hat{Z}}$ can be produced by replacing the Dijkstra algorithm with fast marching and allowing the shortest paths to pass on the faces of the mesh representing $\hat{Z}$. However, such an approach results in more elaborate expressions and will not be discussed here.

[6] Dijkstra's algorithm (unlike e.g. Fast Marching) is known to produce a triangulation-dependent approximation of the geodesic distances due to the so-called metrication error. However, since in our case we are comparing the shape to a deformed version of itself, the triangulation remains the same and thus triangulation-independence is not required.

In order to compute the derivative of $\hat{\mathbf{D}}(\mathbf{Z})$ with respect to $\mathbf{Z}$, we assume that a small perturbation $d\mathbf{Z}$ of $\mathbf{Z}$ does not change the connectivity of the points on $\hat{Z}$, and as a result, the trajectory of the shortest paths between the points on $\hat{Z}$ remains constant (though their length may change). Thus, we may write

$$\hat{\mathbf{D}}(\mathbf{Z}+d\mathbf{Z}) \;=\; \mathcal{I}(\mathbf{Z}+d\mathbf{Z})\mathbf{D}(\mathbf{Z}+d\mathbf{Z}) \;=\; \mathcal{I}(\mathbf{Z})\mathbf{D}(\mathbf{Z}+d\mathbf{Z}), \tag{23}$$

and compute the derivative of $\hat{\mathbf{D}}(\mathbf{Z})$ as the derivative of the linear form $\mathcal{I}(\mathbf{Z})\mathbf{D}(\mathbf{Z})$. If $\mathcal{I}(\mathbf{Z}) \neq \mathcal{I}(\mathbf{Z}+d\mathbf{Z})$, the assumption does not hold, and the derivative of $\hat{\mathbf{D}}$ usually does not exist. Yet, the derivative of $\mathcal{I}(\mathbf{Z})\mathbf{D}(\mathbf{Z})$ belongs to the sub-gradient set of $\hat{\mathbf{D}}$ at the point $\mathbf{Z}$. This is sufficient for many minimization algorithms to work correctly.

The intrinsic distance term can be readily written in terms of $\hat{\mathbf{D}}$ as the Frobenius norm

$$\begin{aligned} d_{\text{int}}(Z) &= \frac{1}{N^2}\,\|\hat{\mathbf{D}}(\mathbf{Z}) - \mathbf{D}_{\hat{X}}\|_{\text{F}}^2 \\ &= \frac{1}{N^2}\,\text{trace}((\hat{\mathbf{D}}(\mathbf{Z}) - \mathbf{D}_{\hat{X}})^{\text{T}}(\hat{\mathbf{D}}(\mathbf{Z}) - \mathbf{D}_{\hat{X}})). \end{aligned} \tag{24}$$

Its derivative with respect to $\mathbf{Z}$ is given by

$$\frac{\partial d_{\text{int}}(\mathbf{Z})}{\partial \mathbf{Z}} \;=\; \frac{2}{N^2}(\hat{\mathbf{D}}(\mathbf{Z}) - \mathbf{D}_{\hat{X}})^{\text{T}}\frac{\partial \hat{\mathbf{D}}(\mathbf{Z})}{\partial \mathbf{Z}}, \tag{25}$$

where

$$\frac{\partial \hat{d}_{ij}(\mathbf{Z})}{\partial \mathbf{Z}} \;=\; \sum_{k,l} \mathcal{I}_{ijkl}\frac{\partial d_{kl}(\mathbf{Z})}{\partial \mathbf{Z}}, \tag{26}$$

and the elements of $\frac{\partial d_{kl}(\mathbf{Z})}{\partial \mathbf{Z}}$ are given by

$$\frac{\partial d_{kl}}{\partial z_n^m} \;=\; \frac{1}{d_{kl}}\begin{cases} z_k^m - z_l^m & : \quad n = k \\ z_l^m - z_k^m & : \quad n = l \\ 0 & : \quad n \neq k, l \end{cases} \tag{27}$$

for $m = 1, 2, 3$.

4.2 Extrinsic distance computation

The computation of the extrinsic distance is similar to the one used in ICP algorithms, where the main difficulty arises from the need to re-compute the closest points each time the extrinsic geometry of $\hat{Z}$ changes. The extrinsic distance term can be written as

$$d_{\text{ext}}(\mathbf{Z}) \quad = \quad \frac{1}{N} \operatorname{trace}\left((\mathbf{Z} - \mathbf{Y}^*(\mathbf{Z}))(\mathbf{Z} - \mathbf{Y}^*(\mathbf{Z}))^{\mathrm{T}}\right) \tag{28}$$

where $\mathbf{Y}^*(\mathbf{Z})$ denotes the $N \times 3$ matrix, whose i-th row y_i^* is the closest point on Y corresponding to z_i. The closest points y_i^* are computed as a weighted average of the points on Y, which are the closest to z_i. The weights are selected in inverse proportion to the distance from z_i.

In ICP algorithms, it is common to assume that $\mathbf{Y}^*(\mathbf{Z} + d\mathbf{Z}) \approx \mathbf{Y}^*(\mathbf{Z})$. By fixing $\mathbf{Y}^*$, $d_{\text{ext}}(\mathbf{Z})$ becomes a simple quadratic function, and its derivative can be written as

$$\frac{\partial d_{\text{ext}}(\mathbf{Z})}{\partial \mathbf{Z}} \quad = \quad \frac{2}{N}(\mathbf{Z} - \mathbf{Y}^*(\mathbf{Z}))^{\mathrm{T}}. \tag{29}$$

4.3 Iterative minimization algorithm

For $\mathbf{Z}'$ in the neighborhood of some $\mathbf{Z}$, the cost function that needs to be minimized can be approximated as

$$\begin{aligned} \sigma(\mathbf{Z}') \quad \approx \quad & \frac{1}{N^2} \operatorname{trace}\left(\hat{\mathbf{D}}(\mathbf{Z}')^{\mathrm{T}}\hat{\mathbf{D}}(\mathbf{Z}') - 2\mathbf{D}_{\hat{X}}^{\mathrm{T}}\hat{\mathbf{D}}(\mathbf{Z}') + \mathbf{D}_{\hat{X}}^{\mathrm{T}}\mathbf{D}_{\hat{X}}\right) \\ & + \frac{\lambda}{N} \operatorname{trace}\left(\mathbf{Z}'\mathbf{Z}'^{\mathrm{T}} - 2\mathbf{Y}^*(\mathbf{Z})\mathbf{Z}'^{\mathrm{T}} + \mathbf{Y}^{*\mathrm{T}}(\mathbf{Z})\mathbf{Y}^*(\mathbf{Z})\right), \end{aligned} \tag{30}$$

where $\hat{\mathbf{D}}(\mathbf{Z}') = \mathcal{I}(\mathbf{Z})\mathbf{D}(\mathbf{Z}')$. Like in ICP, after finding a new $\mathbf{Z}'$ which decreases $\sigma(\mathbf{Z}')$, the closest points $\mathbf{Y}^*$ and the operator $\mathcal{I}$ are updated. The iterative minimization algorithm can be summarized as follows:

1 Compute the closest points $\mathbf{Y}^*(\mathbf{Z})$.
2 Compute the shortest paths between all pairs of points on $\hat{Z}$ and assemble $\mathcal{I}$.
3 Find $\mathbf{Z}'$ sufficiently decreasing $\sigma(\mathbf{Z}')$ in (30).
4 If the change $\|\mathbf{Z} - \mathbf{Z}'\|$ is small, stop. Else, set $\mathbf{Z} = \mathbf{Z}'$ and go to Step 1.

The update of $\mathbf{Z}$ in Step 3 can be safeguarded by evaluating the true cost function (with $\mathcal{I}(\mathbf{Z}')$ and $\mathbf{Y}^*(\mathbf{Z}')$ instead of $\mathcal{I}(\mathbf{Z})$ and $\mathbf{Y}^*(\mathbf{Z})$). In our implementation, no safeguard was used, and the minimization in Step 3 was done using conjugate gradients.

The initialization of the algorithm can be done in several ways, the simplest of which is $\mathbf{Z} = \mathbf{X}$. This choice works well when the extrinsic dissimilarity between X and Y is not too large; for large $d_{\mathrm{ext}}(X, Y)$, the algorithm will suffer from poor convergence similarly to most ICP methods. Another choice is to initialize $\mathbf{Z}$ by the corresponding points on $\hat{Y}$ resulting from the solution of the GMDS problem. This choice is suitable for objects having sufficiently similar intrinsic geometries, making the intrinsic correspondence computed by GMDS meaningful. A more robust initialization scheme can be constructed based on branch-and-bound global optimization proposed in [30] for the initialization of ICP algorithms, and adopted in [61, 9] for the initialization of GMDS. This family of approaches consists of computing local descriptors (in our case, reflecting intrinsic geometric properties) for a set of prominent feature points on both shapes, followed by finding the best correspondence between the two sets of features. The branch-and-bound technique is used for fast pruning of the search space, guaranteeing global optimality of the found solution on one hand, while maintaining reasonably low complexity on the other.

As both in ICP and GMDS, a multi-resolution scheme can improve significantly the convergence speed of our minimization algorithm. In a multi-resolution approach, an initial solution is found by first computing d_{joint} between coarse versions of X and Y, and subsequently interpolating it to higher resolution levels. This allows to practically eliminate many spurious local minima.

5 Results

In order to assess the proposed approach, three experiments were performed. In the first experiment, we show the computation of the set-valued Pareto distance between different nonrigid shapes. In the second experiment, we evaluate the discriminative power of the joint similarity criterion and compare its performance to extrinsic and intrinsic similarity criteria. In the third experiment, we compare joint similarity with state-of-the-art shape matching methods.

The shapes in our experiments were taken from the nonrigid objects database available online at `tosca.cs.technion.ac.il`. The joint similarity computation was implemented in MATLAB. The Dijkstra algorithm was written in C. No code optimization was performed.

5.1 Pareto distance

In the first experiment, we compared three different objects: a man, a woman and a gorilla. As a model, we used the shape of aiming man (shown in red in Figure 8), sampled at 1000 points. Three probes were compared to the model: a near-isometric deformation of the man shape (a man with open hands), a woman and a gorilla. Each probe was represented by 100 points. Subsampling was performed using the *farthest point sampling* method [38, 37, 31]. The Pareto frontier was obtained by computing the joint distance for thirteen different values of λ. The computation of each point on the Pareto frontier took several minutes.

Figure 7 shows the Pareto distances between the model shape and the probe objects. The Pareto frontiers clearly indicate that the man shape is more similar to its deformed version than to a woman and even less to the gorilla.

Figure 8 shows the matching obtained as a result of computation of the joint distance. Small values of λ give larger weight to intrinsic similarity, which results in the probe (depicted in blue) remaining almost rigid. Increasing λ, more weight is given to the extrinsic similarity, which results in the probe bent to better fit the model. Thus, "traversing" the solutions

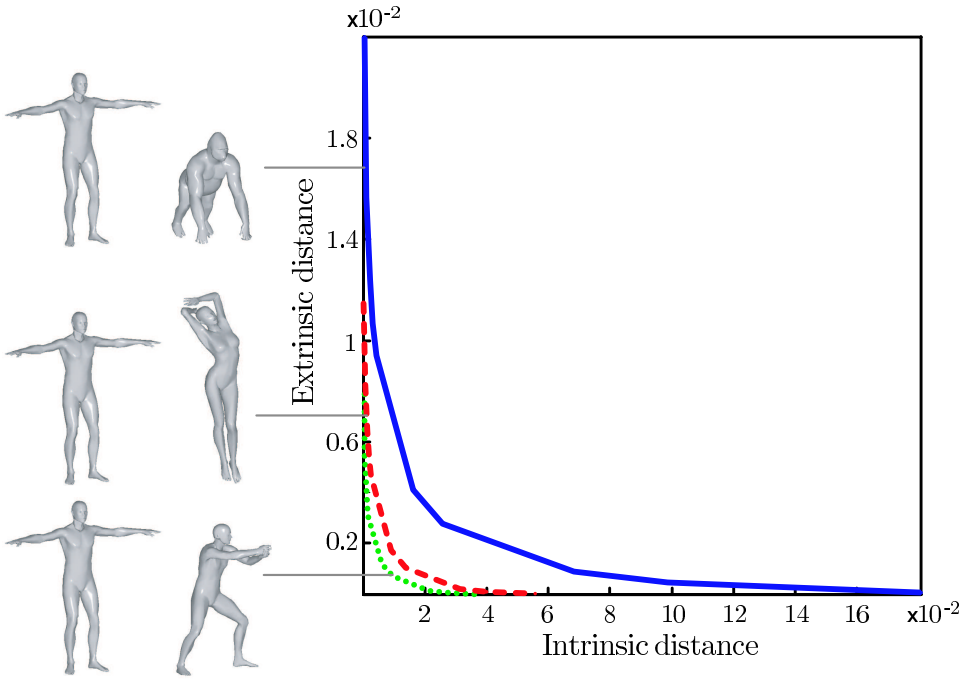


Figure 7: Pareto distances between a man shape and its deformed version (green curve), man and woman (red curve) man and gorilla (blue curve).

on the Pareto frontier, one can obtain a continuous morphing between the probe and the model. Note that unlike many morphing methods used in computer graphics which assume compatible meshing of the source and the target shapes, here the two shapes have arbitrary number of samples, and arbitrary triangulations.

5.2 Comparison of d_{H}, d_{GH} and d_{joint}

In the second experiment, we compared different distances on a data set of six shape classes. Each object appeared in a variety of instances, obtained by near-isometric deformations. In addition, for each of the deformations, a version with different topology was created by welding the shapes at a set of points (marked by red circles in Figure 9). In total, the data set contained the following shapes: 5 cats, 9 dogs, 13 gorillas, 7 lions, 13 males and 13

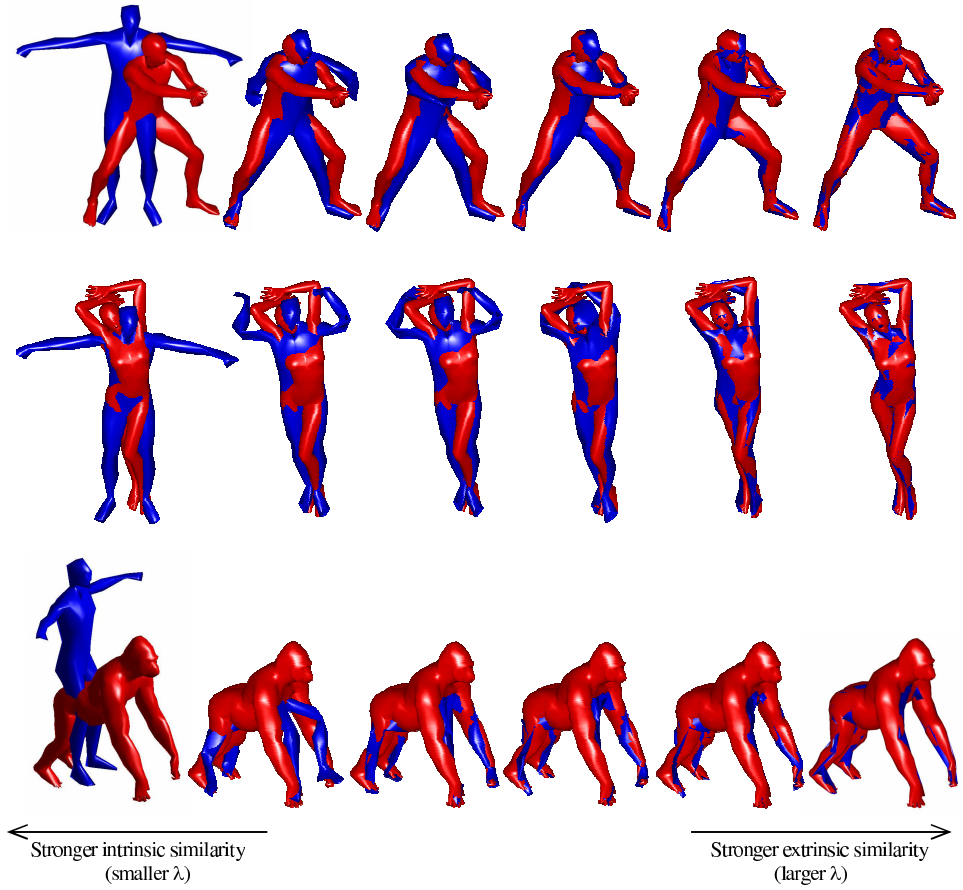


Figure 8: Matching a man shape to different shapes (another version of the man, a woman and a gorilla) produced by joint similarity using different values of λ. Gradually increasing the value of λ results in a morphing effect.

females (total of 60 shapes; 27 of them with welding).

We compared three shape distances: the Gromov-Hausdorff distance d_{GH} (representing intrinsic dissimilarity), extrinsic dissimilarity computed using an ICP algorithm, and the proposed joint distance d_{joint}. For the approximation of the Gromov-Hausdorff distance, we used GMDS. The distance was computed using 100 points on the embedded shape and 1000 points on the shape into which embedding was performed. For the ICP method, meshes with 1000 sampled were used. For the computation of the joint distance, the model and the probe were represented by 1000 and 100 points, respectively. The typical computation time of the extrinsic, intrinsic and joint similarity between a pair of shapes was a few seconds, 30 - 60 seconds, and a few minutes, respectively.

The recognition accuracy was assessed both qualitatively and quantitatively. The first assessment consisted of presenting the shapes as points in the Euclidean space, with the Euclidean distance representing the distances between the shapes (Figures 10–13). Such plots are straightforwardly obtained using MDS and allow to visually represent the similarity relations between the shapes.

The second, quantitative assessment consisted of computing the *receiver operating characteristic* (ROC) curves for each similarity criterion, representing a tradeoff between the *false acceptance rate* (FAR) and the *false rejection rate* (FRR). Each ROC curve was computed as follows: the matrix of distances between different shapes was thresholded by a value ranging from zero to the maximum distance value. Shapes with distances falling below the threshold were regarded similar (i.e., instances of the same object); those with distances above the threshold were regarded dissimilar (different objects). The FAR was computed as the percentage of dissimilar shapes wrongfully identified as similar. The FRR was computed as the percentage of similar shapes wrongfully identified as dissimilar. For small values of the threshold, the FAR is small (two shapes must have a very small distance in order to be considered similar), while the FAR is large. For large values of the threshold, the FAR is large and the FRR is small. Ideally, both should be as small as possible, meaning that the recognition is accurate. A

single number capturing the recognition error was computed as the point at which the values of FAR and FRR coincide (referred to as *equal error rate* or EER).

Figure 10 visualizes the Gromov-Hausdorff distance between the shapes in the absence of topological noise (using a subset of the database of 33 shapes without welding). The Gromov-Hausdorff distance appears insensitive to deformations, which is seen as tight clusters in the figure, and allows to distinguish between different objects almost perfectly (with EER of 1.14%, see Figure 14). This idealistic picture changes dramatically when shapes with topological noise are added. Figure 11 shows that shapes with welded points (represented as hollow circles in the figure) are significantly different from the original ones. This is confirmed by a significant drop in the recognition rate (EER of 7.68%).

Figure 12 visualizes the ICP distance between the shapes. While the extrinsic distance is insensitive to topological noise (which is clearly seen from the hollow circles, representing shapes with welding, coinciding with points, representing shapes without welding), it is sensitive to nonrigid deformations. Overall, the recognition rate is poor (EER of 10.34%).

Finally, Figure 13 visualizes the joint distance. The joint similarity criterion combines the advantages of the two distances, and is insensitive to both isometries and topological changes. It significantly outperforms the extrinsic and intrinsic distances used separately, achieving an EER of 1.61%.

5.3 Comparison to other methods

In the third experiment, we compared the joint similarity criterion with two state-of-the-art methods: shape DNA [62, 50, 65] and D2 shape distribution [59]. The shape DNA method is based on comparison of the Laplace-Beltrami spectrum of the shapes. The discrete Laplace-Beltrami operator was computed on shapes subsampled at 500 points. We used the approximation described in [65]. Twenty largest eigenvalues were used as shape descriptors. The descriptors were compared using Euclidean norm. In the D2 shape distribution method, the shape is described by a histogram of Euclidean distances between uniform samples. We used uniform subsam-

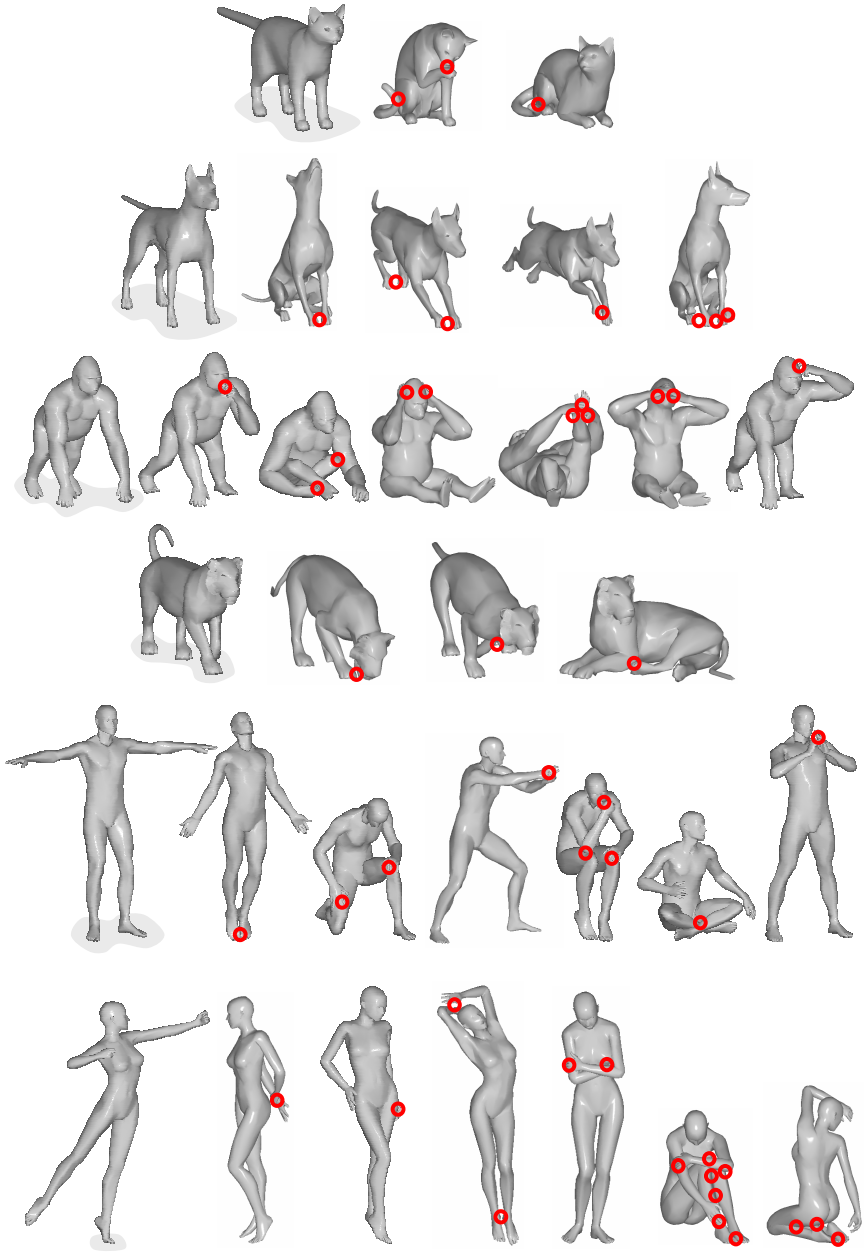


Figure 9: The set of objects used in the second experiment. Topological noise was modeled by welding the meshes at points indicated by red circles.



Figure 10: Visualization of the intrinsic similarity (Gromov-Hausdorff distance) between nonrigid shapes. No topological noise is present.

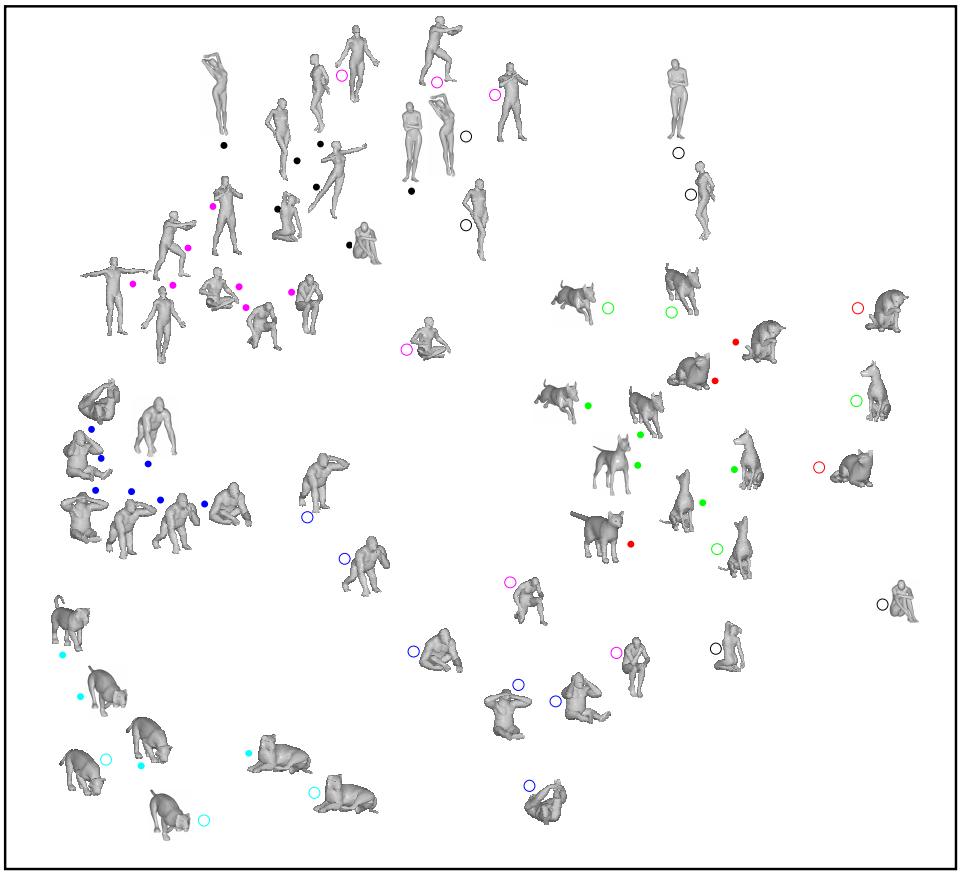


Figure 11: Visualization of the intrinsic similarity (Gromov-Hausdorff distance) between nonrigid shapes. Shapes with different topology created by welding are marked by hollow circles.

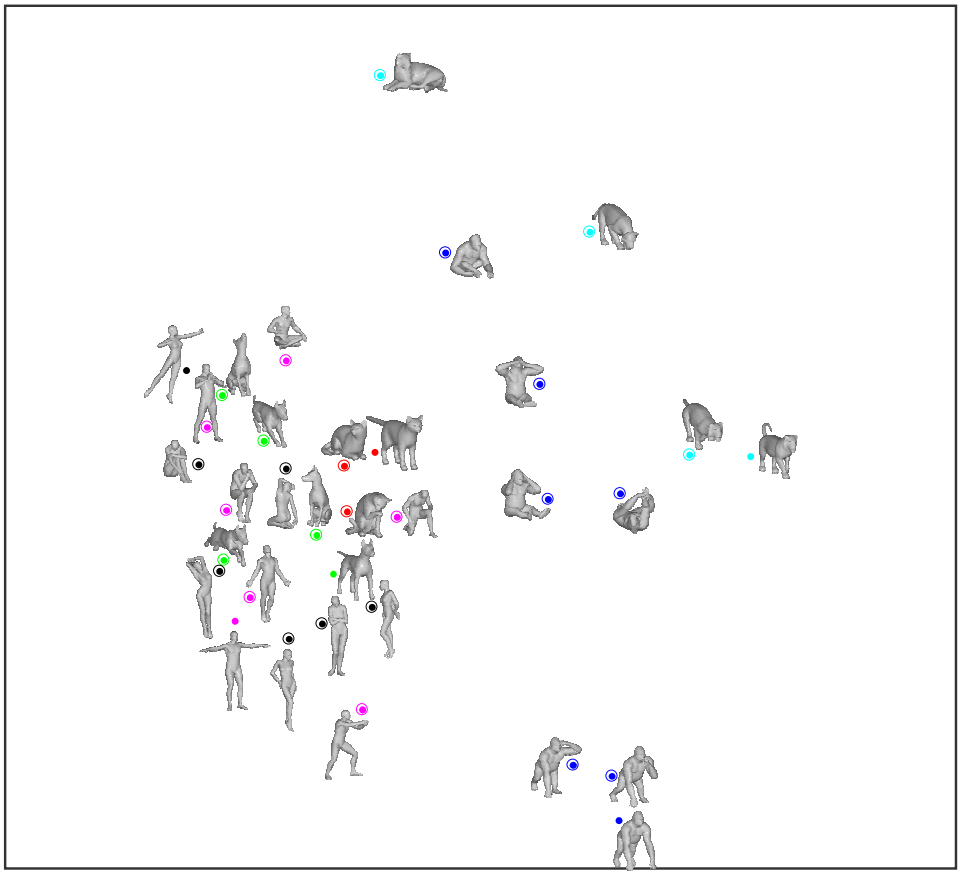


Figure 12: Visualization of the extrinsic similarity (computed using ICP) between nonrigid shapes. Shapes with different topology created by welding are marked by hollow circles.

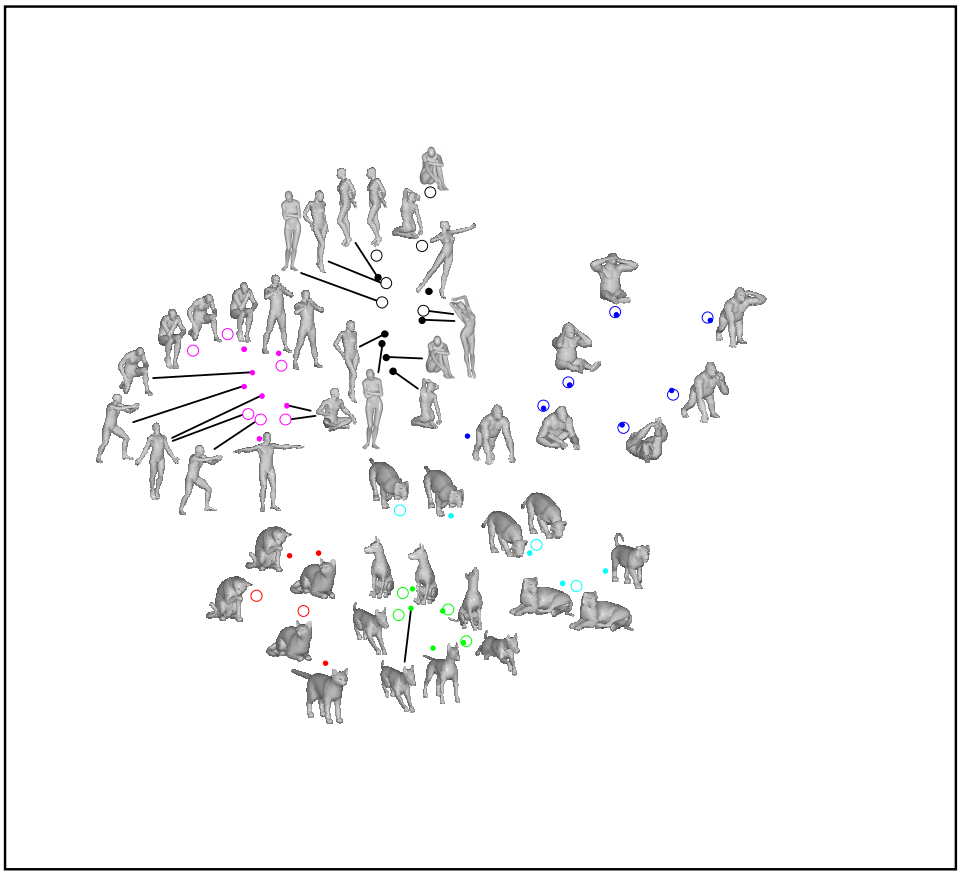


Figure 13: Visualization of the joint intrinsic-extrinsic similarity computed using the proposed method. Shapes with different topology created by welding are marked by hollow circles.

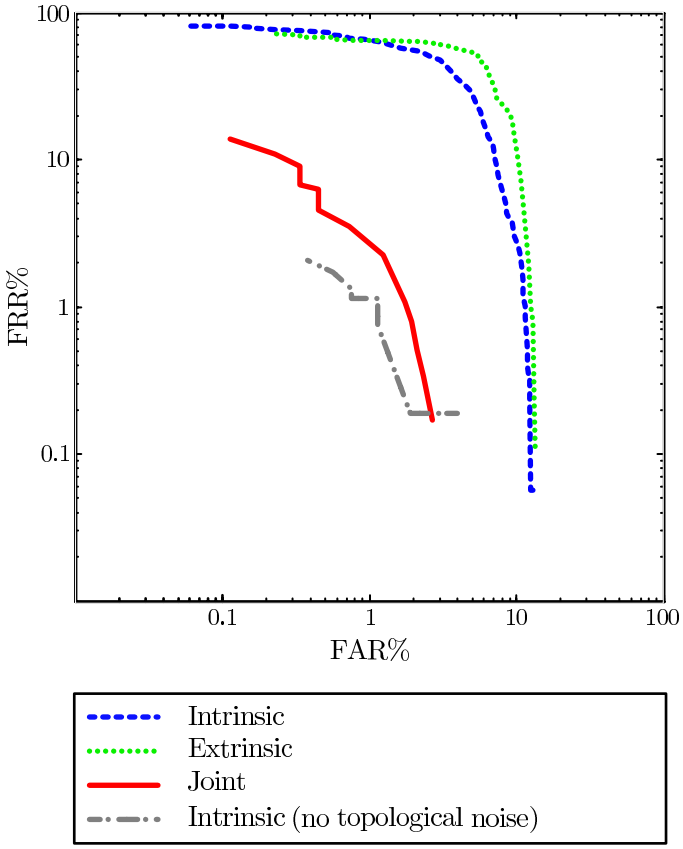


Figure 14: ROC curves describing the recognition power of intrinsic, extrinsic and joint similarities in the second experiment.

pling based on farthest point sampling. We used histograms with 125 bins. Comparison of histograms was performed using the *Earth Moving distance* (EMD).

Figure 15 shows the performance of these method on the same database used in the second experiment. The EER of the shape DNA and the shape distribution methods is 8.02% and 11.8%, respectively. One can see that the proposed joint similarity significantly outperforms the the shape descriptor methods achieving the EER of 1.61%.

The disadvantage of our method is the computational complexity and comparison efficiency: each comparison requires solving a complicated optimization problem (in current implementation, each comparison takes a few minutes). For comparison, the computation of the Laplace-Beltrami spectrum took 4.6 sec in our implementation and the comparison of two spectra using the Euclidean norm was negligible. The computation of the D2 shape distribution took 4.2 sec and the computation of the EMD between two distribution was below one second.

6 Conclusions

We presented a new approach for the computation of nonrigid shape similarity as a tradeoff between extrinsic and intrinsic similarity criteria. Our approach can be illustratively presented as deforming one shape in order to make it the most similar to another from an extrinsic point of view, while trying to preserve as much as possible its intrinsic geometry. The joint intrinsic and extrinsic similarity appears to be advantageous over traditional purely extrinsic or intrinsic similarity criteria. While extrinsic similarity is sensitive to strong nonrigid deformations and intrinsic similarity is sensitive to topology changes, our joint similarity criterion allows to gracefully handle both geometric and topological deformations. Experimental results prove that it can be used in situations where intrinsic and extrinsic similarities fail.

The numerical framework presented in this paper extends beyond shape similarity problems. As a byproduct of our similarity computation, we ob-

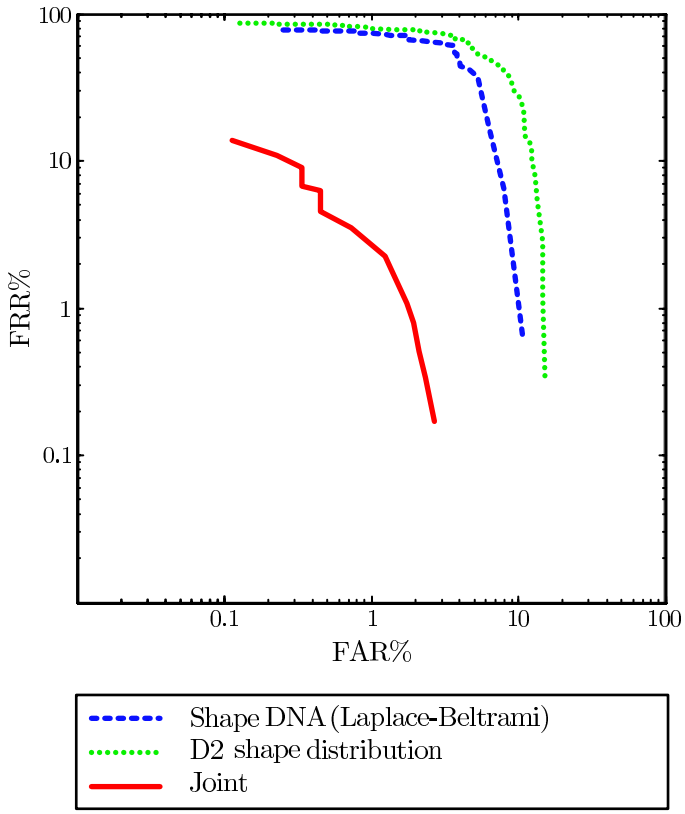


Figure 15: ROC curves describing the recognition power of joint similarity, shape DNA and D2 shape distribution methods in the third experiment.

tain nonrigid alignment or correspondence of two shapes. This potentially allows to employ the proposed framework for morphing problems in computer graphics in a way similar to [27, 45]. Additional potential applications are inverse problems arising in shape reconstruction [2, 67]. Our intrinsic and extrinsic distances can be used as priors for regularization of such problems.

References

[1] B. Amberg, S. Romdhani, and T. Vetter, *Optimal step nonrigid ICP algorithms for surface registration*, Proc. CVPR, 2007.

[2] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis, *SCAPE: shape completion and animation of people*, Proc. SIGGRAPH **24** (2005), no. 3, 408–416.

[3] M. Belkin and P. Niyogi, *Laplacian eigenmaps for dimensionality reduction and data representation*, Neural Computation **13** (2003), 1373–1396.

[4] M. Berger, *A panoramic view of Riemannian geometry*, Springer, Berlin, 2002.

[5] P. J. Besl and N. D. McKay, *A method for registration of 3D shapes*, IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) **14** (1992), 239–256.

[6] I. Borg and P. Groenen, *Modern multidimensional scaling - theory and applications*, Springer, 1997.

[7] A. M. Bronstein and M. M. Bronstein, *Not only size matters: regularized partial matching of nonrigid shapes*, Proc. Workshop on Non-rigid Shape Analysis and Deformable Image Alignment (NORDIA), 2008.

[8] A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel, *Partial similarity of objects, or how to compare a centaur to a horse*, Int'l J. Computer Vision (IJCV), to appear.

[9] ———, *Analysis of two-dimensional non-rigid shapes*, Int'l J. Computer Vision (IJCV) **78** (2008), no. 1, 67–88.

[10] A. M. Bronstein, M. M. Bronstein, Y. S. Devir, R. Kimmel, and O. Weber, *Parallel algorithms for approximation of distance maps on parametric surfaces*, Proc. ACM Trans. Graphics (2008), submitted.

[11] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, *Three-dimensional face recognition*, Int'l J. Computer Vision (IJCV) **64** (2005), no. 1, 5–30.

[12] ———, *Efficient computation of isometry-invariant distances between surfaces*, SIAM J. Scientific Computing **28** (2006), no. 5, 1812–1836.

[13] ———, *Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching*, Proc. National Academy of Science (PNAS) **103** (2006), no. 5, 1168–1172.

[14] ———, *Robust expression-invariant face recognition from partially missing data*, Proc. European Conf. Computer Vision (ECCV), 2006, pp. 396–408.

[15] ———, *Calculus of non-rigid surfaces for geometry and texture manipulation*, IEEE Trans. Visualization and Computer Graphics **13** (2007), no. 5, 902–913.

[16] ———, *Expression-invariant representation of faces*, IEEE Trans. Image Processing **16** (2007), no. 1, 188–197.

[17] ———, *Rock, paper, and scissors: extrinsic vs. intrinsic similarity of non-rigid shapes*, Proc. Int. Conf. Computer Vision (ICCV), 2007.

[18] ———, *Numerical geometry of nonrigid shapes*, Springer, 2008.

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

[20] A. M. Bruckstein, N. Katzir, M. Lindenbaum, and M. Porat, *Similarity-invariant signatures for partially occluded planar shapes*, IJCV **7** (1992), no. 3, 271–285.

[21] D. Burago, Y. Burago, and S. Ivanov, *A course in metric geometry*, Graduate studies in mathematics, vol. 33, AMS, 2001.

[22] Y. Chen and G. Medioni, *Object modeling by registration of multiple range images*, Proc. Conf. Robotics and Automation, 1991.

[23] H. Chui and A. Rangarajan, *A new point matching algorithm for non-rigid registration*, Computer Vision and Image Understanding **89** (2003), no. 2-3, 114–141.

[24] F. R. K. Chung, *Spectral graph theory*, AMS, 1997.

[25] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, *Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps*, Proc. National Academy of Sciences (PNAS) **102** (2005), no. 21, 7426–7431.

[26] D. Donoho and C. Grimes, *Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data*, Proc. National Academy of Science (PNAS), vol. 100, 2003, pp. 5591–5596.

[27] I. Eckstein, J. P. Pons, Y. Tong, C. C. J. Kuo, and M. Desbrun, *Generalized surface flows for mesh processing*, Proc. Symposium on Geometry Processing, 2007.

[28] A. Elad and R. Kimmel, *Geometric methods in bio-medical image processing*, vol. 2191, ch. Spherical flattening of the cortex surface, pp. 77–89, Springer, 2002.

[29] ———, *On bending invariant signatures for surfaces*, IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) **25** (2003), no. 10, 1285–1295.

[30] N. Gelfand, N. J. Mitra, L. Guibas, and H. Pottmann, *Robust global registration*, Proc. Symp. Geometry Processing (SGP) (2005).

[31] T. F. Gonzalez, *Clustering to minimize the maximum intercluster distance*, Theoretical Computer Science **38** (1985), no. 2, 293–306.

[32] C. Gordon, D. L. Webb, and S. Wolpert, *One cannot hear the shape of the drum*, Bulletin AMS **27** (1992), no. 1, 134–138.

[33] H. Groemer, *Geometric applications of Fourier series and spherical harmonics*, Cambridge University Press, New York, 1996.

[34] M. Gromov, *Structures métriques pour les variétés riemanniennes*, Textes Mathématiques, no. 1, 1981.

[35] R. Grossman, N. Kiryati, and R. Kimmel, *Computational surface flattening: a voxel-based approach*, IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) **24** (2002), no. 4, 433–441.

[36] D. Hahnel, S. Thrun, and W. Burgard, *An extension of the ICP algorithm for modeling nonrigid objects with mobile robots*, Proc. IJCAI (2003).

[37] D. S. Hochbaum and D. B. Shmoys, *A best possible approximation algorithm for the k-center problem*, Mathematics of Operations Research **10** (1985), no. 2.

[38] ———, *A best possible heuristic for the k-center problem*, Mathematics of Operations Research **10** (1985), no. 2, 180–184.

[39] B. K. P. Horn, *Closed-form solution of absolute orientation using unit quaternions*, J. Optical Society of America (JOSA) A **4** (1987), 629–642.

[40] D. Jacobs, D. Weinshall, and Y. Gdalyahu, *Class representation and image retrieval with non-metric distances*, IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) **22** (2000), 583–600.

[41] ———, *Class representation and image retrieval with non-metric distances*, IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) **22** (2000), 583–600.

[42] M. Kac, *Can one hear the shape of a drum?*, American Mathematical Monthly **73** (1966), 1–23.

[43] S. Katz, G. Leifman, and A. Tal, *Mesh segmentation using feature point and core extraction*, The Visual Computer **21** (2005), no. 8, 649–658.

[44] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, *Rotation invariant spherical harmonic representation of 3D shape descriptors*, Proc. Symposium on Geometry Processing, 2003, pp. 156–164.

[45] M. Kilian, N. J. Mitra, and H. Pottmann, *Geometric modeling in shape space*, Proc. SIGGRAPH, vol. 26, 2007.

[46] R. Kimmel and J. A. Sethian, *Computing geodesic paths on manifolds*, Proc. National Academy of Sciences (PNAS) **95** (1998), no. 15, 8431–8435.

[47] L. J. Latecki and R. Lakaemper, *Shape similarity measure based on correspondence of visual parts*, IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) **22** (2000), 1185–1190.

[48] L. J. Latecki, R. Lakaemper, and D. Wolter, *Optimal Partial Shape Similarity*, Image and Vision Computing **23** (2005), 227–236.

[49] S. Leopoldseder, H. Pottmann, and H. Zhao, *The d2-tree: A hierarchical representation of the squared distance function*, Tech. report, Institute of Geometry, Vienna University of Technology, 2003.

[50] B. Lévy, *Laplace-Beltrami eigenfunctions towards an algorithm that "understands" geometry*, Int'l Conf. Shape Modeling and Applications, 2006.

[51] H. Ling and D. Jacobs, *Deformation invariant image matching*, Proc. ICCV, 2005.

[52] ———, *Using the inner-distance for classification of articulated shapes*, Proc. CVPR, 2005.

[53] N. Linial, E. London, and Y. Rabinovich, *The geometry of graphs and some its algorithmic applications*, Combinatorica **15** (1995), 333–344.

[54] N. Litke, M. Droske, M. Rumpf, and P. Schroder, *An image processing approach to surface matching*.

[55] F. Mémoli and G. Sapiro, *A theoretical and computational framework for isometry invariant recognition of point cloud data*, Foundations of Computational Mathematics **5** (2005), 313–346.

[56] N. J. Mitra, N. Gelfand, H. Pottmann, and L. Guibas, *Registration of point cloud data from a geometric optimization perspective*, Proc. Eurographics Symposium on Geometry Processing, 2004, pp. 23–32.

[57] B. Mohar, *Graph theory, combinatorics, and applications*, vol. 2, ch. The Laplacian spectrum of graphs, pp. 871–898, Wiley, 1991.

[58] M. Novotni and R. Klein, *3D Zernike descriptors for content based shape retrieval*, 2003, pp. 216–225.

[59] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, *Shape distributions*, ACM Transactions on Graphics (TOG) **21** (2002), no. 4, 807–832.

[60] E. Paquet, M. Rioux, A. Murching, T. Naveen, and A. Tabatabai, *Description of shape information for 2-D and 3-D objects*, Signal Processing: Image Communication **16** (2000), no. 1-2, 103–122.

[61] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel, *Symmetries of non-rigid shapes*, Proc. Workshop on Non-rigid Registration and Tracking through Learning (NRTL), 2007.

[62] M. Reuter, F.-E. Wolter, and N. Peinecke, *LaplaceBeltrami spectra as "shape-DNA" of surfaces and solids*, Computer Aided Design **38** (2006), 342–366.

[63] G. Rosman, M. M. Bronstein, A. M. Bronstein, and R. Kimmel, *Topologically constrained isometric embedding*, Workshop on Human Motion Understanding, Modeling, Capture and Animation (D. Metaxas, R. Klette, and B. Rosenhahn, eds.), Lecture Notes in Computer Science, Springer, 2007, to appear.

[64] S. T. Roweis and L. K. Saul, *Nonlinear dimensionality reduction by locally linear embedding*, Science **290** (2000), no. 5500, 2323–2326.

[65] R. M. Rustamov, *Laplace-Beltrami eigenfunctions for deformation invariant shape representation*, Proc. Symp. Geometry Processing (SGP), 2007, pp. 225–233.

[66] M. E. Salukwadze, *Vector-valued optimization problems in control theory*, Academic Press, 1979.

[67] M. Salzmann, J. Pilet, S. Ilic, and P. Fua, *Surface deformation models for nonrigid 3D shape recovery*, Trans. PAMI **29** (2007), no. 8.

[68] E. L. Schwartz, A. Shaw, and E. Wolfson, *A numerical solution to the generalized mapmaker's problem: flattening nonconvex polyhedral surfaces*, IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI) **11** (1989), 1005–1008.

[69] H. Y. Shum, M. Hebert, and K. Ikeuchi, *On 3D shape similarity*, School of Computer Science, Carnegie Mellon University, 1995.

[70] A. Tal, M. Elad, and S. Ar, *Content based retrieval of VRML objects - an iterative and interactive approach*, Proc. Eurographics Workshop on Multimedia, 2001.

[71] J. W. H. Tangelder and R. C. Veltkamp, *A survey of content based 3D shape retrieval methods*, Proc. Shape Modeling Applications, 2004, pp. 145–156.

[72] M. R. Teague, *Image analysis via the general theory of moments*, J. Optical Society of America (JOSA) **70** (1979), 920–930.

[73] J. B. Tennenbaum, V. de Silva, and J. C. Langford, *A global geometric framework for nonlinear dimensionality reduction*, Science **290** (2000), no. 5500, 2319–2323.

[74] D. V. Vranic, D. Saupe, and J. Richter, *Tools for 3D-object retrieval: Karhunen-Loeve transform and spherical harmonics*, Proc. IEEE Fourth Workshop on Multimedia Signal Processing, 2001, pp. 293–298.

[75] J. Walter and H. Ritter, *On interactive visualization of high-dimensional data using the hyperbolic plane*, Proc. Int'l Conf. Knowledge Discovery and Data Mining (KDD), 2002, pp. 123–131.

[76] M. Yu, I. Atmosukarto, W. K. Leow, Z. Huang, and R. Xu, *3D model retrieval with morphing-based geometric and topological feature maps*, Proc. CVPR, vol. 2, 2003.

[77] C. Zhang and T. Chen, *Efficient feature extraction for 2D/3D objects in meshrepresentation*, Proc. IEEE ICIP, vol. 3, 2001.

[78] ———, *Indexing and retrieval of 3D models aided by active learning*, ACM Multimedia (2001), 615–616.

[79] Z. Zhang and H. Zha, *Principal manifolds and nonlinear dimension reduction via local tangent space alignment*, Technical Report CSE-02-019, Department of Computer Science and Engineering, Pennsylvania State University, 2002.

[80] Z. Y. Zhang, *Iterative point matching for registration of free-form curves and surfaces*, Int'l J. Computer Vision (IJCV) **13** (1994), 119–152.

[81] G. Zigelman, R. Kimmel, and N. Kiryati, *Texture mapping using surface flattening via multi-dimensional scaling*, IEEE Trans. Visualization and Computer Graphics (TVCG) **9** (2002), no. 2, 198–207.