

A Gromov-Hausdorff Framework with Diffusion Geometry for Topologically-Robust Non-rigid Shape Matching

Alexander M. Bronstein · Michael M. Bronstein ·
Ron Kimmel · Mona Mahmoudi · Guillermo Sapiro

Received: 21 September 2008 / Accepted: 21 September 2009 / Published online: 8 October 2009
© Springer Science+Business Media, LLC 2009

Abstract In this paper, the problem of non-rigid shape recognition is studied from the perspective of metric geometry. In particular, we explore the applicability of diffusion distances within the Gromov-Hausdorff framework. While the traditionally used geodesic distance exploits the shortest path between points on the surface, the diffusion distance averages all paths connecting the points. The diffusion distance constitutes an intrinsic metric which is robust, in particular, to topological changes. Such changes in the form of shortcuts, holes, and missing data may be a result of natural non-rigid deformations as well as acquisition and representation noise due to inaccurate surface construction. The presentation of the proposed framework is complemented with examples demonstrating that in addition to the relatively low

complexity involved in the computation of the diffusion distances between surface points, its recognition and matching performances favorably compare to the classical geodesic distances in the presence of topological changes between the non-rigid shapes.

Keywords Non-rigid shape matching · Diffusion geometry · Gromov-Hausdorff distance · Topology · Missing data · Partiality

1 Introduction

Non-rigid shapes are ubiquitous in the world we live in, from microscopic bacteria to tissues and parts of our body. Since pattern recognition applications need to deal with objects encountered in everyday life, non-rigid shape analysis has become important in many modern applications, such as object retrieval and recognition, surface matching, navigation, and target detection and recognition. One of the cornerstone problems in the analysis of non-rigid shapes is the problem of *shape similarity*: given two objects, we need to tell how similar or dissimilar they are. This can be quantitatively expressed as a *distance* between two shapes. The main difficulty in such a comparison stems from the immense number of degrees of freedom present in the problem as a result of possible deformations that the non-rigid shapes can undergo.

A common way in which the problem of non-rigid shape similarity has been approached in the pattern recognition literature is to try to find a representation of shapes which is *invariant* to a given class of deformations (Olver 1999). Using such a representation, it is then possible to compare shapes regardless of their deformations, when these deformations are limited to the given class. Riemannian geometry is of help in finding such invariant representations (Do-

This paper is dedicated to Prof. Gromov in the occasion of receiving the Abel Prize in May 2009.

This work is partially supported by NSF, ONR, NGA, ARO, DARPA, NIH, and by the Israel Science Foundation (ISF grant No. 623/08).

A.M. Bronstein · M.M. Bronstein · R. Kimmel
Department of Computer Science, Technion–Israel Institute of
Technology, Haifa, Israel

A.M. Bronstein
e-mail: bron@cs.technion.ac.il

M.M. Bronstein
e-mail: mbron@cs.technion.ac.il

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

M. Mahmoudi · G. Sapiro (✉)
Department of Electrical and Computer Engineering, University
of Minnesota, Minneapolis, USA
e-mail: guille@umn.edu

M. Mahmoudi
e-mail: mahmo022@umn.edu

Carmo 1992). It is well-known, for example, that the *intrinsic* properties of a shape remain invariant under inelastic deformations, i.e., deformations that do not stretch or tear the object. Many recent papers, e.g., (Bronstein et al. 2005, 2006a, 2006b; Elad and Kimmel 2003; Gal et al. 2007; Hamza and Krim 2005; Lévy 2006; Mahmoudi and Sapiro 2009; Mémoli and Sapiro 2005b; Mémoli et al. 2006; Reuter et al. 2006; Ruggeri and Saupe 2008; Rustamov 2007; Schwartz et al. 1989; Tannenbaum et al. 2000), exploit this fact in order to construct deformation-invariant shape distances and surface matching techniques.

Elad and Kimmel (2003) introduced a method for the recognition of 3D shapes based on Euclidean embedding, extending previous efforts by Schwartz et al. (1989) (see also Tannenbaum et al. 2000). The key idea of the method is to consider a shape as a metric space, whose metric structure is defined by the *geodesic distances* between pairs of points on the shape.¹ Geodesic distances, being an intrinsic property of the shape, are invariant to any inelastic deformation the shape can undergo, or using metric geometry terminology, we can say that such deformations are *isometric* or *metric-preserving* with respect to the geodesic distances. Two non-rigid shapes are compared by first having their respective geodesic metric structures mapped into a low-dimensional Euclidean space using *multidimensional scaling* (MDS) (Cox and Cox 1994), and then rigidly matching the resulting images (called *canonical forms*). MDS allows to “undo” the non-rigid deformations of the shapes, leading to a bending-invariant shape comparison framework based on the pairwise geodesic distances. This method has been used in three-dimensional face recognition (Bronstein et al. 2005), analysis of articulated two-dimensional shapes and images (Ling and Jacobs 2005a, 2005b), texture mapping and object morphing (Bronstein et al. 2007a; Grossman et al. 2002; Zigelman et al. 2002), and shape segmentation (Katz et al. 2005).

Due to the fact that the canonical forms method uses an intermediate metric space to compare two shapes, inaccuracies are introduced, as it is theoretically impossible to embed a generic metric structure into a finite-dimensional Euclidean space without distorting it. It was shown empirically in Bronstein et al. (2007b), Elad and Kimmel (2002, 2003), Walter and Ritter (2002) that using spaces with non-Euclidean (non-flat) geometry makes it possible to obtain more accurate representations, but can not avoid the embedding error completely.

Mémoli and Sapiro (2005b), proposed a metric framework for non-rigid shape comparison based on the *Gromov-Hausdorff distance*. This distance was introduced by Mikhail

Gromov (1981), as a way to compute similarity between metric spaces. Using the Gromov-Hausdorff formalism, the comparison of two shapes can be posed as direct comparison of pairwise distances on the shapes (basically, in the discrete case, the comparison up to permutations of the corresponding pairwise distance matrices or their corresponding submatrices). Since no fixed intermediate space is forced, the representation error inherent to canonical forms can be avoided. The Gromov-Hausdorff distance computation is an NP-hard problem, and together with a number of theoretical results, Mémoli and Sapiro (2005b) proposed a practical approximation scheme with explicit probabilistic bounds connecting the approximation to the actual Gromov-Hausdorff distance and the number of available sample points.

According to an alternative but mathematically equivalent definition (Burago et al. 2001), the Gromov-Hausdorff distance computation can be posed as measuring the distortion of embedding one metric space into another. Bronstein et al. (2006a, 2006b), observing the connection between this formulation of the Gromov-Hausdorff distance and MDS, proposed an efficient computation of such embedding based on a continuous optimization problem similar in spirit to MDS, referred to as *generalized MDS* (GMDS). This method follows the line of thought of embedding into non-Euclidean spaces where instead of embedding each shape into Euclidean, hyperbolic, or spherical spaces, as classically done in MDS, the embedding is done from one shape into the other. We exploit the GMDS computational framework for the examples in this paper. Alternative computational techniques based on the relaxation of the Gromov-Hausdorff distance have been recently proposed (Mémoli 2007).

All the aforementioned contributions on the Gromov-Hausdorff and MDS-based shape analysis frameworks considered geodesic distances as the invariant used to intrinsically compare non-rigid shapes. This is motivated by a number of fundamental reasons. First, many natural object deformations can be approximated as inelastic ones. Thus, methods based on geodesic distances allow good shape recognition accuracy. Second, there exists a plethora of efficient numerical methods for the computation of geodesic distances for diverse shape representations (Kimmel and Sethian 1998; Mémoli and Sapiro 2001, 2005a; Spira and Kimmel 2004; Weber et al. 2008; Yatziv et al. 2006).

At the same time, the notable drawback of the geodesic distances is their sensitivity to topological transformations. By modifying the connectivity of the shape, one can significantly alter the paths between points, and in particular the shortest one, which in turn, can result in significant changes of the geodesic distances. Inconsistent topology or “topological noise” are common phenomena in shapes acquired using 3D scanners or obtained as a result of point cloud triangulation (Wood et al. 2004). Thus, in order to be able to

¹Recall that the geodesic distance between two points is the length of the shortest path, traveling on the surface of the shape, that connects the points.

deal with real-life data, it is important for shape similarity methods to be topology-invariant or at least topologically robust. Of course, topological changes can result from natural non-rigid deformations as well, e.g., bending of an open hand until the finger tips touch. A practically useful shape recognition system should be able to deal with these transformations.

It should be noted, however, that the metric model of shapes allows to represent a shape as a generic metric space with any metric. Also, neither the canonical forms framework nor the Gromov-Hausdorff distance are necessarily limited to geodesic distances. Using another intrinsic distance, insensitive or robust to topological changes, instead of the geodesic distance, could potentially make these methods cope with topological noise and topological deformations in general. Bronstein et al. (2007c, 2009) showed that using Euclidean distances (which are robust to topology changes but not invariant or even robust to non-rigid deformations), and geodesic distances (which are invariant to non-rigid deformations but not to topology changes) together in a framework resembling both of the GMDS and iterative closest point (ICP) algorithms (Besl and McKay 1992; Chen and Medioni 1991), allows to obtain a shape similarity method more robust to topological changes than the one obtained with the geodesic distances alone. Also, it was shown in Bronstein et al. (2007c, 2009) that the Gromov-Hausdorff type distance between shapes modeled as metric spaces with Euclidean geometry allows to obtain an alternative formulation, and then use GMDS as a computation method for the ICP method, a classical approach for rigid shape comparison (for more detailed analysis, see Mémoli 2008).

Motivated by the generality of the metric model of shapes, and the need to cope with topological transformations in non-rigid shape analysis, in this paper we propose to use *diffusion geometry* in the Gromov-Hausdorff framework. *Diffusion distances*, introduced by Coifman and Lafon (2006) and Lafon (2004), are related to the probability of traveling on the surface from one point to another in a fixed number of random steps (random walk). There are a number of reasons that lead us to use these distances instead of geodesic ones when we deal with topological noise and changes. First, the diffusion distance is an average length of paths connecting two points on the shape, while the geodesic distance is the length of the shortest path. This naturally makes diffusion distances less sensitive to topological changes (Mahmoudi and Sapiro 2009; Ovsjanikov et al. 2009; Rustamov 2007; Sun et al. 2009). Secondly, both diffusion and geodesic distance are intrinsic, thus invariant to inelastic deformations. Finally, diffusion distance can be efficiently computed from the eigenvalues of a discrete approximation to the Laplace-Beltrami operator (or simply the eigenvalues of a weighted connectivity matrix).

The interested reader is referred to Jain and Zhang (2007), Mateus et al. (2008), Reuter et al. (2006), Rusta-

mov (2007) for some recent works on 3D shape recognition based on spectral methods, both for geodesic matrices and the Laplace-Beltrami operator, which is closely related to the diffusion distance (Hein et al. 2007). For example, Reuter et al. (2006), Ruggeri et al. (2009) consider the spectra (eigenvalues) of the Laplace-Beltrami operator and use this as bending invariant signature for shape recognition. Rustamov (2007) uses both eigenvalues and eigenvectors of the Laplace-Beltrami operator, in order to construct a shape representation and stresses its relevance for topological robustness. The diffusion distance was also exploited in Mahmoudi and Sapiro (2009) for 3D point cloud recognition using the framework of distance distributions by Osada et al. (2002) and to construct intrinsic descriptors for feature-based shape recognition (Ovsjanikov et al. 2009; Sun et al. 2009). Bronstein and Bronstein (2009) showed the relation between the approaches of Rustamov (2007) and Mahmoudi and Sapiro (2009).

Combining the Gromov-Hausdorff framework with diffusion distances provides a fundamental framework leading to a non-rigid shape comparison and matching approach which is robust to topological transformations such as holes and point-wise connectivity changes. The Gromov-Hausdorff distance is an invariant metric on the space of equivalence classes of shapes under shape isometries. It allows to compare shapes as metric spaces directly, without resorting to approximations such as low-dimensional embedding into a Euclidean space. Methods based on such pre-defined embeddings are theoretically sub-optimal (Gromov 1981) (see also Mémoli 2008), the Gromov-Hausdorff distance can not be achieved by projecting onto finite dimensional Euclidean spaces, and typically result in worse performance for our applications, as experimentally demonstrated in Bronstein et al. (2006a, 2006b) and in this paper. The Gromov-Hausdorff distance is also supported by a well-developed theory, relating it to metric geometry of shapes. The computation of the Gromov-Hausdorff distance not only permits to match isometric shapes, but also to quantify their distance from isometry, that is, their ϵ -isometry, that measures how much the intrinsic distances are preserved (Burago et al. 2001). Finally, an intrinsic correspondence between non-rigid shapes is a natural byproduct of the Gromov-Hausdorff framework.

The remainder of this paper is organized as follows. Following a description of the contributions of this paper, next, in Section 2 we present the metric approach to shape matching and recognition, and the Gromov-Hausdorff framework. Section 3 describes the basic diffusion distance theory. Section 4 summarizes the relations between different methods. Section 5 is devoted to the numerical computation of the proposed shape distance based on the GMDS algorithm, and Section 6 presents experimental results which include comparisons with previous approaches. Finally, Section 7 concludes the paper.

1.1 Contributions

In this paper, we address the challenge of invariant shape similarity and correspondence under a wide class of shape transformations, such as bending, connectivity change, and missing information. To achieve this, we put together two fundamental mathematical frameworks. The first is the metric approach for shape representation, modeling shapes as metric spaces and using Gromov's results on distances between metric spaces to measure shape similarity directly from their metric structures, without resorting to suboptimal approximate representations. The second mathematical framework is the spectral analysis of surfaces, connecting Laplace–Beltrami operators of surfaces to random walks and giving rise to the diffusion geometry.

Each one of these two fundamental theories have been recently shown to be very successful for different aspects of shape analysis, when acting separately. We here show that putting them together permits to address new problems such as shape matching under these challenging conditions. The generality of the metric approach allows us to use different metrics, and in particular, diffusion metrics derived from the Laplace–Beltrami operator. We show that the class of invariance (robustness) obtained by using the diffusion metric is wider than in the case of the geodesic metric, and the resulting shape similarity and correspondence approaches are less sensitive to topological changes of the shapes and to missing information.

Finally, previous experimental results (Mahmoudi and Sapiro 2009) demonstrated that the Gromov–Hausdorff framework is accuracy-wise superior to distance distribution methods in the geodesic case. Here, we provide experimental evidence of the same behavior in the case of diffusion distances.

2 Metric Approach for Shape Matching and Recognition

In this section, we present some basic concepts that constitute the core of the metric approach for shape recognition, staying mostly at the intuitive level. For a rigorous and insightful treatment of the topic, the reader is referred to Burago et al. (2001), Gromov (1981).

2.1 Basic Notions in Metric Geometry

We model a non-rigid shape as a *metric space* (X, d_X) , where X is a two-dimensional smooth compact connected and complete Riemannian surface (possibly with boundary) embedded into \mathbb{R}^3 , and $d_X : X \times X \rightarrow \mathbb{R}$ is a *metric* measuring distances between pairs of points on X . The key idea of the metric approach is to compare shapes as metric spaces.

Two shapes (X, d_X) and (Y, d_Y) are *similar* if the metrics between pairs of corresponding points on X and Y coincide, i.e., there exists a bijective map $\varphi : X \rightarrow Y$ such that $d_Y \circ (\varphi \times \varphi) = d_X$. Such a φ is called an *isometry* and X and Y in this case are said to be *isometric*. Isometry implies that in terms of intrinsic metric geometry, the two shapes are indistinguishable and thus are equivalent.

The notion of isometry can be relaxed in order to define similarity of shapes. We will refer to a set $C \subset X \times Y$ of pairs such that for every $x \in X$ there exists at least one $y \in Y$ such that $(x, y) \in C$, and similarly for every $y \in Y$ there exists an $x \in X$ such that $(x, y) \in C$, as a *correspondence* between X and Y . Note that a correspondence C is not necessarily a function. We can define the *distortion* of the correspondence as the discrepancy between the corresponding metrics,

$$\text{dis}(C) := \sup_{(x,y), (x',y') \in C} |d_X(x, x') - d_Y(y, y')|.$$

We say that the shapes X and Y are ϵ -*isometric* if there exists a correspondence C with $\text{dis}(C) \leq \epsilon$. Such a C is called an ϵ -*isometry*. ϵ -*isometry* can be regarded as a criterion of shape similarity. For small values of ϵ , the shapes are similar, and for large values of ϵ , the shapes are dissimilar in a metric sense.

2.2 Distance Distributions

Osada et al. (2002) proposed to describe shapes using *shape distributions*. In particular, the D_2 distribution in the authors' terminology is the histogram of pairwise Euclidean distances between the points of the shape. Such a description is invariant to Euclidean transformations. Two shapes can be compared by comparing their corresponding distributions. The work in Boutin and Kemper (2004) provides further theoretical basis for this approach.

In Hamza and Krim (2005), Mahmoudi and Sapiro (2009), it was proposed to use intrinsic distances instead of the Euclidean ones. This way, a deformation-invariant descriptor is obtained; shape comparison can be done same as in Osada et al. (2002). The authors in Mahmoudi and Sapiro (2009) showed two types of distances: geodesic and diffusion. Rustamov (2007) used shape distribution applied to Euclidean distances on eigenmaps of the Laplace–Beltrami operator, which is equivalent to distribution of commute time (see detailed discussion in Section 4).

2.3 Canonical Forms

One of the first attempts in the computer vision community to regard non-rigid shapes as metric spaces and formulate shape similarity in terms of metric geometry was proposed by Elad and Kimmel (2003). The metric geometry of (X, d_X) (d_X being the geodesic distance in Elad and

Kimmel 2003, and the one derived from the spectral framework later in Rustamov 2007), is represented in some fixed metric space $(\mathbb{Z}, d_{\mathbb{Z}})$ (typically, a low-dimensional Euclidean space), in such a way that the metric $d_{\mathbb{Z}}$ is as close as possible to the metric d_X by means of a minimum-distortion embedding $\varphi : X \rightarrow \mathbb{Z}$. The low-dimensional representation of (X, d_X) was dubbed in Elad and Kimmel (2003) a *canonical form*.

The canonical form can be computed by solving the *multidimensional scaling* problem (Cox and Cox 1994),

$$\min_{\varphi: X \rightarrow \mathbb{Z}} \max_{x, x' \in X} \|d_X(x, x') - d_{\mathbb{Z}}(\varphi(x), \varphi(x'))\|,$$

where $\varphi(X)$ gives the canonical form. In general, the canonical form is not unique, and is defined up to an isometry in \mathbb{Z} . Also, in most cases the representation is not exact, since in general there is no isometry between (X, d_X) and the pre-defined space $(\mathbb{Z}, d_{\mathbb{Z}})$.

Given two shapes (X, d_X) and (Y, d_Y) , their canonical forms $\varphi(X)$ and $\psi(Y)$ are computed. $\varphi(X)$ and $\psi(Y)$ can be compared as subsets of the metric space $(\mathbb{Z}, d_{\mathbb{Z}})$ using the *Hausdorff distance*,

$$d_{\mathbb{H}}^{\mathbb{Z}}(\varphi(X), \psi(Y)) = \max \left\{ \max_{y \in \psi(Y)} \min_{x \in \varphi(X)} d_{\mathbb{Z}}(x, y), \max_{x \in \varphi(X)} \min_{y \in \psi(Y)} d_{\mathbb{Z}}(x, y) \right\}.$$

Since the canonical forms are defined up to an isometry in $(\mathbb{Z}, d_{\mathbb{Z}})$, in order to undo the isometry ambiguity, rigid shape comparison algorithms such as ICP (Besl and McKay 1992; Chen and Medioni 1991), can be employed. In essence, ICP algorithms minimize the Hausdorff distance between the canonical forms on all the isometries in $(\mathbb{Z}, d_{\mathbb{Z}})$,

$$d_{\text{ICP}}(\varphi(X), \psi(Y)) = \min_{i \in \text{Iso}(\mathbb{Z})} d_{\mathbb{H}}^{\mathbb{Z}}(\varphi(X), i \circ \psi(Y)).$$

2.4 Spectral Canonical Forms

The method of Elad and Kimmel is intimately related to *spectral embeddings* such as those proposed by Belkin and Niyogi (2003), Rustamov (2007), Coifman and Lafon (2006) and Lafon (2004). The main idea of these approaches is to create an Euclidean representation of the shape by taking a few eigenvectors of a Laplacian (Laplace-Beltrami operator) defined on the shape. The metric induced by the Laplace-Beltrami operator is the diffusion metric described in details in Sect. 3.2.

There are two major differences between Elad-Kimmel canonical forms obtained using MDS and canonical forms obtained by spectral embedding. First, unlike the Elad-Kimmel canonical forms always defined up to an isometry in the embedding space, the spectral canonical form are unique if the Laplace-Beltrami operator has unique eigenvalues (no

multiplicity). In case of eigenvalues with multiplicity greater than one, the degrees of freedom in the spectral canonical form can come from permutations of the Laplace-Beltrami eigenfunctions (or more generally, also reflections in subspaces spanned by the multiple eigenfunctions) (Mateus et al. 2008; Ovsjanikov et al. 2008).

Secondly, Lafon et al. and Rustamov showed that if all the eigenvectors are used (implying an infinite-dimensional Euclidean representation), the embedding is isometric, that is, the resulting canonical form is a theoretically *exact representation* of the diffusion geometry of the shape. In practice, however, a finite-dimensional approximation is used by taking the first eigenfunctions of the Laplace-Beltrami operator (Belkin and Niyogi 2003; Rustamov 2007). The implications of such an approximations are discussed in Sect. 4.

Finally, note that the canonical forms consider a pre-defined space to map onto, the Euclidean space, and in the case of the spectral canonical form, the map itself is pre-selected as well. The distance presented next optimizes over these selections as well.

2.5 Gromov-Hausdorff Distance

An elegant framework to represent similarity of metric spaces as a distance was proposed by Gromov (Burago et al. 2001; Gromov 1981) and introduced into the non-rigid shape recognition area in Mémoli and Sapiro (2005b).

The source of inaccuracy of canonical forms was the fixed metric space $(\mathbb{Z}, d_{\mathbb{Z}})$ (and fixed embedding map). It is generally impossible to select a common metric space in which the geometry of any shape can be accurately represented. However, $(\mathbb{Z}, d_{\mathbb{Z}})$ and the map can be chosen in an optimal way for given two surfaces, resulting in the following shape distance:

$$d_{\text{GH}}(X, Y) := \inf_{\substack{\varphi: X \rightarrow \mathbb{Z} \\ \psi: Y \rightarrow \mathbb{Z} \\ \mathbb{Z}}} d_{\mathbb{H}}^{\mathbb{Z}}(\varphi(X), \psi(Y)), \quad (1)$$

where φ and ψ are isometric embeddings, and $d_{\mathbb{H}}^{\mathbb{Z}}$ is the Hausdorff distance in $(\mathbb{Z}, d_{\mathbb{Z}})$. d_{GH} is called the *Gromov-Hausdorff distance*.

For compact surfaces and our shape recognition framework, the Gromov-Hausdorff distance can also be expressed in terms of the distortion obtained by embedding one surface into another,

$$d_{\text{GH}}(X, Y) := \frac{1}{2} \inf_C \text{dis}(C), \quad (2)$$

where the infimum is taken over all correspondence C , and $\text{dis}(C)$ is the distortion defined above. The two expressions (1) and (2) are equivalent (Burago et al. 2001).

The Gromov-Hausdorff distance is a metric on the quotient space of metric spaces under the isometry relation, and

thus, in the context of the metric space model for shape recognition, is a good candidate for a shape distance (Mémoli and Sapiro 2005b). Being a metric particularly implies that $d_{GH}(X, Y) = 0$ if and only if X and Y are isometric. More generally, if $d_{GH}(X, Y) \leq \epsilon$, then X and Y are 2ϵ -isometric and conversely, if X and Y are ϵ -isometric, then $d_{GH}(X, Y) \leq 2\epsilon$ (Burago et al. 2001). The latter property relates the Gromov-Hausdorff distance to the notion of ϵ -isometry and makes it a good criterion of similarity between shapes.

2.6 Choice of a Metric

The metric approach we have described and the Gromov-Hausdorff distance do not specify any particular choice of the metric d_X . In general, d_X is independent of X and can be defined quite arbitrarily. There are, however, two natural choices of d_X . The first choice is the *geodesic metric*, measuring the length of the shortest intrinsic path between a pair of points, constructing the *intrinsic geometry* of X . The second choice is the extrinsic *Euclidean metric*, measuring the length of a line in \mathbb{R}^3 connecting two points on X that relates to the *extrinsic geometry* in which X is embedded, i.e., \mathbb{R}^3 (see Mémoli 2008).

Extrinsic geometry is invariant to rigid transformations of the shape (rotation, translation, and reflection), which preserve Euclidean distances. However, nonrigid deformations may change the extrinsic geometry (see example in Fig. 1). As a result, the Euclidean metric is not suitable for the comparison of shapes with significant bending or other type of non-rigid deformations. The intrinsic geometry on the other hand is invariant to inelastic shape deformations which, do not stretch or tear the shape. As a particular case, it is also invariant to rigid transformations. Therefore, the geodesic metric is a good choice for comparing non-rigid shapes, as has been confirmed by numerous results as those mentioned in the introduction.

Another important type of transformations a shape can undergo are those changing the shape *topology*. Omitting formal definitions, the topology of X can be thought of as a collection of neighborhoods of every point on X . This defines the *connectivity* of the shape—which points can be reached by a small step from a neighbor point. Topology can change as a result of non-rigid deformations, such as a person bending, fingers touching, or a paper folding. Topological dissimilarities between shapes can also be the result of noise, typically arising in acquisition of shapes by a 3D scanner. Connectivity changes do not change the geometry (location of the points) of the shape, therefore, the Euclidean metric is not sensitive to such transformations. Yet, geodesic distances can drastically change as a result of minor connectivity changes, since such changes alter the shortest paths between two points on the shape.




Fig. 1 Illustration of the difference between intrinsic and extrinsic similarity for comparison of non-rigid shapes with different topology. Bending the fingers makes the left and the middle shapes (open palm and bent hand, without fingers yet touching, respectively) extrinsically dissimilar while being intrinsically similar. Additionally gluing/touching the fingers (bottom right) makes the shapes intrinsically dissimilar, as the connectivity change modifies the intrinsic geometry

It is therefore important to use a metric which is both intrinsic (and thus invariant to non-rigid deformations) and also invariant (or at least robust) to topological changes. In this paper, we propose to use in the Gromov-Hausdorff framework a different metric that while being intrinsic and invariant to bends, is also robust to topological changes, thereby simultaneously enjoying important properties of both extrinsic and intrinsic geometry. This distance is explained next.

3 Diffusion Geometry

In Coifman and Lafon (2006), Lafon (2004) (see also Belkin and Niyogi 2003 for a related early effort), Lafon et al. introduced diffusion maps and diffusion distances as a method for data parametrization, embedding, and dimensionality reduction. Informally, the diffusion distance between two points on a shape is related to heat diffusion on a shape (hence the name), or equivalently, the probability of arriving from one point to another by a random walk in fixed time. Since the diffusion distance operates on paths defined on the shape, it is an intrinsic property, and as a result, is bending-invariant. Moreover, since the diffusion distances have an effect of averaging over all possible paths connecting two points, while the geodesic distance is the length of just the minimal one, the former is more robust than the geodesic distance in cases where topological changes are present (Sun et al. 2009; Ovsjanikov et al. 2009).

As an illustration, imagine the hand shape and two points on the tips of the index finger and the thumb (see Fig. 1). If the two fingertips are not touching, then both geodesic and diffusion distances between the two points are large, as all paths connecting the two points travel throughout the whole hand. Yet, if the hand is bent in such a way that the fingertips touch each other, the minimal geodesic will re-route itself through the shortcut across the fingertips instead of going through the hand, leading to a significant change in the geodesic distance. For the diffusion distance, this new path added as a result of the topology change is averaged with the other paths, which reduces the effect of such a change. The lesser sensitivity to topological changes attributed to the averaging property of the diffusion distance, may come at the expense of a potential reduction in discriminative power, as it usually happens in a tradeoff between invariance and discriminativity. However, many applications require robustness to topological changes, where this approach is advantageous. Besides the above properties, the diffusion distance is a metric, and thus, a valid candidate for the definition of a metric space used in the Gromov-Hausdorff shape model.

3.1 Heat Kernels

Formally, in order to define the diffusion distance, we resort to the *heat equation*,

$$\left(\Delta_X + \frac{\partial}{\partial t}\right)u = 0, \tag{3}$$

governing the distribution of heat u on the surface. Here, Δ_X denotes the *Laplace-Beltrami operator*, a generalization of the Laplacian to non-Euclidean domains (here, we define the Laplacian as a positive semi-definite operator, and thus it comes with a negative sign). The fundamental solution $h_t(x, z)$ of the heat equation (3), also called the *heat kernel*, is the solution of the heat equation with a point heat source at $x \in X$. The heat kernel $h_t(x, z)$ provides the heat value at time t at point $z \in X$. Interpreting the heat kernel probabilistically, $h_t(x, z)$ can be associated with the probability of getting to point z by means of a random walk on X of length t starting at point x .²

For compact manifolds, the Laplace-Beltrami operator has discrete eigendecomposition of the form

$$\Delta_X \phi_i = \lambda_i \phi_i, \tag{4}$$

where $\lambda_0 \geq \lambda_1 \geq \dots$ are eigenvalues and ϕ_0, ϕ_1, \dots are eigenfunctions. Since Δ_X is a positive semi-definite operator, $\lambda_0 = 0, \lambda_1, \lambda_2, \dots \geq 0$ and $\phi_0 = const$.

²More precisely, given a Brownian motion W_t on X and a Borel set $C \subset X$, the integral $\int_C h_t(x, z) dz$ is the probability $P(W_t \in C)$ of the Brownian motion ending up in C after time t (Coifman and Lafon 2006; Lafon 2004; Ovsjanikov et al. 2009; Sun et al. 2009).

The heat kernel can be presented as (Jones et al. 2008)

$$h_t(x, z) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x) \phi_i(z) \tag{5}$$

in the Laplace-Beltrami eigenbasis. In practical computation, the heat kernel is approximated by taking the sum over the first $k + 1$ largest eigenvalues,

$$\tilde{h}_t(x, z) = \sum_{i=0}^k e^{-\lambda_i t} \phi_i(x) \phi_i(z). \tag{6}$$

Since the eigenvalues typically decay fast, a small k can be sufficient (see Sect. 5.1 for more details).

3.2 Diffusion Distance

The *diffusion distance* is defined as a cross-talk between two heat kernels,

$$\begin{aligned} d_{X,t}^2(x, y) &= \|h_t(x, \cdot) - h_t(y, \cdot)\|_{L_2(X)}^2 \\ &= \int_X |h_t(x, z) - h_t(y, z)|^2 dz. \end{aligned} \tag{7}$$

It is inversely related to the connectivity of points x and y by paths of length t (i.e., if there are many such paths connecting x and y , the distance $d_{X,t}(x, y)$ is small) (Coifman and Lafon 2006; Lafon 2004; Ovsjanikov et al. 2009; Sun et al. 2009). As will be shown later, $d_{X,t}$ is a metric and therefore can be used to define a valid metric space in our framework.

Using the eigenfunctions of Δ_X , the diffusion distance can also be expressed as (Coifman and Lafon 2006)

$$d_{X,t}^2(x, y) = \sum_{i=1}^{\infty} e^{-2\lambda_i t} (\phi_i(x) - \phi_i(y))^2, \tag{8}$$

which allows its discretization as will be shown in Sect. 5.1.

3.3 Commute Time

A notion related to the diffusion distance is the *commute time*, given by

$$\delta^2 X(x, y) = 2 \int_0^{\infty} d_{X,t}^2(x, y) dt \tag{9}$$

$$= \sum_{i=1}^{\infty} \frac{1}{\lambda_i} (\phi_i(x) - \phi_i(y))^2. \tag{10}$$

While diffusion distance reflects the connectivity of two points by paths of length t , the commute time is the sum of the diffusion length of all possible paths connecting a pair of points (Qiu and Hancock 2007; Bronstein and Bronstein 2009). Commute time is also a metric.

3.4 Eigenmaps

Maps of the form $\Phi(x) = (\alpha(\lambda_0)\phi_0(x), \alpha(\lambda_1)\phi_1(x), \dots)$ where α is some function of the eigenvalues, are referred to as *Laplacian eigenmaps* (Belkin and Niyogi 2003). Informally, $\Phi(x)$ can be considered as an infinite-dimensional canonical form of X ; for this reason, we also use the term *spectral canonical form*.

In particular, for $\alpha(\lambda_i) = e^{-\lambda_i t}$, the map

$$\Phi_{X,t}(x) = (e^{-\lambda_0 t} \phi_0(x), e^{-\lambda_1 t} \phi_1(x), \dots)$$

given by the heat kernel, is called a *diffusion map* (Coifman and Lafon 2006). For $\alpha(\lambda_i) = \lambda_i^{-1/2}$, $i \geq 1$, the map $\Phi_X(x) = (\lambda_1^{-1/2} \phi_0(x), \lambda_2^{-1/2} \phi_1(x), \dots)$ was referred to as *global point signature* (GPS) embedding by Rustamov (2007).

Unlike canonical forms computed by MDS, an eigenmap is uniquely defined (i.e., there are no degrees of freedom related to the isometry in the embedding space) if the Laplace-Beltrami operator has no eigenvalues of multiplicity greater than one. Otherwise, the ambiguity in the definition of the eigenmap is up to switching between the eigenfunction corresponding to the eigenvalues with multiplicity and changes in their signs. For example, if $\lambda_i = \lambda_{i+1} = \dots = \lambda_{i+l-1}$ is an eigenvalue of multiplicity l , then the components $i, \dots, i + l - 1$ of the GPS embedding $\Phi_X(x)$ are of the form $c_i \lambda_i^{-1/2} \phi_{\pi_i}(x), \dots, c_{i+l-1} \lambda_i^{-1/2} \phi_{\pi(i+l-1)}(x)$, where π is an arbitrary permutation on the indices and $c_i, \dots, c_{i+l-1} \in \{-1, 1\}$. Such ambiguities arise in cases of symmetric shapes (Ovsjanikov et al. 2008). In general, two eigenmaps differ by a permutation of coordinates, i.e., $\Pi\Phi(x)$, where Π is a permutation of indices $\{0, 1, \dots\}$.

In practice, a finite-dimensional eigenmap is constructed by taking the first $k + 1$ components of the eigenmap, $\tilde{\Phi}(x) = (\alpha(\lambda_0)\phi_0(x), \dots, \alpha(\lambda_k)\phi_k(x))$. In simple cases (small k and no ambiguity due to eigenvalue multiplicity), the comparison of such eigenmaps can be performed using standard rigid similarity methods such as ICP. In case of $k \gg 1$ and shape symmetries resulting in multiple eigenvalues, such a comparison is a challenging problem (Ovsjanikov et al. 2008).

As a possible remedy, Rustamov (2007) proposed applying Osada’s method (Osada et al. 2002) to eigenmaps, representing shapes as histograms of pairwise Euclidean distances measured on the GPS embedding $\tilde{\Phi}_X(x)$. This way, explicit matching of the canonical forms with all the related problems is avoided. In particular, the Euclidean distances are insensitive to problems related to eigenfunction permutations, since $\|\Phi(x) - \Phi(y)\|_2 = \|\Pi\Phi(x) - \Pi\Phi(y)\|_2$ for any permutation Π . Moreover, the comparison of distance distributions is computationally efficient. The price paid in this approach is the reduction of *discriminativity*, as there

can be many shapes with equal distance distributions but different geometry (see Boutin and Kemper 2004).

4 Relations Between Methods

4.1 Distance Distributions

The observation

$$\|\Phi_{X,t}(x) - \Phi_{X,t}(y)\|_2 = d_{X,t}(x, y), \tag{11}$$

and

$$\|\Phi_X(x) - \Phi_X(y)\|_2 = \delta_X(x, y) \tag{12}$$

provides an alternative expression of the diffusion distance and commute time (see details and proof in Bronstein and Bronstein 2009). In practical computations, we use the approximation (6), yielding an approximate diffusion distance $\tilde{d}_{X,t}(x, y) = \|\tilde{\Phi}_{X,t}(x) - \tilde{\Phi}_{X,t}(y)\|_2$ and approximate commute time $\tilde{\delta}_X(x, y) = \|\tilde{\Phi}_X(x) - \tilde{\Phi}_X(y)\|_2$.

Equations (11) and (12) have a few implications. First, by virtue of (11) and (12), both the diffusion distance and the commute time are metrics. Secondly, the distribution of diffusion distances $\tilde{d}_{X,t}$ approximated with a finite number of eigenvalues and eigenfunctions on the shape is equivalent to the distribution of the Euclidean distances on the diffusion map $\tilde{\Phi}_{X,t}(X)$. Similarly, the distribution of approximate commute times $\tilde{\delta}_X$ is equivalent to the distribution of the Euclidean distances on the GPS embedding $\tilde{\Phi}_X(X)$. In other words, the Euclidean distances measured on the eigenmaps coincide with distances arising from diffusion geometry measured on the shape itself.

Using the relation between diffusion and commute time distance, we conclude that the methods of Rustamov (2007) and Mahmoudi and Sapiro (2009) are equivalent (up to integration over all t) and arrive to a similar shape distance in different ways. Computationally, diffusion distances are usually more advantageous than commute time, in the sense that they require less eigenfunctions and eigenvalues (Bronstein and Bronstein 2009).

4.2 Gromov-Hausdorff Distance and ICP

Equations (11) and (12) also imply that the metric spaces $(X, \tilde{d}_{X,t})$ and $(\tilde{\Phi}_{X,t}(X), \|\cdot\|_{\mathbb{R}^{k+1}})$ are isometric. Given two shapes X and Y , the Gromov-Hausdorff distance between X and Y with the approximated diffusion distances $\tilde{d}_{X,t}$ and $\tilde{d}_{Y,t}$ is equal to the Gromov-Hausdorff distance in the space of spectral canonical forms with the Euclidean metric,

$$\begin{aligned} d_{GH}((X, \tilde{d}_{X,t}), (Y, \tilde{d}_{Y,t})) \\ = d_{GH}((\tilde{\Phi}_{X,t}(X), \|\cdot\|_{\mathbb{R}^{k+1}}), (\tilde{\Phi}_{Y,t}(Y), \|\cdot\|_{\mathbb{R}^{k+1}})). \end{aligned}$$

Mémoli (2008) showed that the Gromov-Hausdorff distance in the Euclidean space is equivalent to the ICP distance,

$$c \cdot d_{\text{ICP}} \leq d_{\text{GH}} \leq d_{\text{ICP}}$$

in the sense of equivalence of metrics ($c > 0$ is a constant). This result implies that, from the computational point of view, the distance resulting from the comparison of spectral canonical forms using ICP is metrically equivalent to the method proposed in this paper once the diffusion distance has been approximated with a finite number of eigenfunctions as detailed above. Note, however, that metric equivalence should not be confused with equality: for example, L_1 and L_2 metric are equivalent but not equal.

Particularly important for shape comparison, the Gromov-Hausdorff distance (unlike the ICP distance) is explicitly related to the notion of ϵ -isometry, our criterion of shape similarity. Moreover, ICP-based comparison of shapes is impractical for large dimensions ($k \gg 1$) and is further complicated by the problem of eigenfunction permutation in case of symmetries. Our method based on direct approximation of the Gromov-Hausdorff distance does not suffer from these problems. This means that while we can approximate the diffusion distance with $k \gg 1$, and then compute the Gromov-Hausdorff distance $d_{\text{GH}}((X, \tilde{d}_{X,t}), (Y, \tilde{d}_{Y,t}))$, we cannot in practice use ICP to compute $d_{\text{GH}}((\hat{\Phi}_{X,t}(X), \|\cdot\|_{\mathbb{R}^{k+1}}), (\hat{\Phi}_{X,t}(Y), \|\cdot\|_{\mathbb{R}^{k+1}}))$ when $k \gg 1$.

Finally, shape comparison based on the Gromov-Hausdorff framework is known to outperform distance distribution-based methods from the point of view of accuracy in the case of geodesic metric d_X (Bronstein et al. 2006a, 2006b; Mahmoudi and Sapiro 2009) (being, however, computationally more expensive). We expect the same behavior in the case of diffusion geometry, i.e., the Gromov-Hausdorff framework to be superior to methods based on distributions of diffusion distances (Rustamov and Mahmoudi & Sapiro).

4.3 Choice of the Scale

The advantage of diffusion geometry is the presence of time parameter t , allowing to compare shapes at different scales. At the same time, it introduces the problem of scale selection. Similarly to Mahmoudi and Sapiro (2009), we used a fixed scale selected empirically to maximize the algorithm performance in our experiments. This straightforward approach does not make use of the multi-scale nature of diffusion geometry.

In particular, some of the shape transformations (e.g. topological changes and deformations) are less prominent at certain scales (Bronstein and Bronstein 2009). If such transformations are present at the same time, the choice of a single scale is impossible. In Bronstein and Bronstein (2009),

it was shown that using all the scales for shape comparison results in better performance than a single scale.

The problem of scale selection was also addressed by Mémoli (2009) as a follow-up of the present paper. Mémoli proposed the *spectral Gromov-Wasserstein distance*, in which the scale is part of the optimization problem and sup is taken over all $t > 0$. This way, the most discriminative scale that tells two shapes apart is chosen (Mémoli 2009).

5 Numerical Computation

In order to make the problems of intrinsic shape similarity and correspondence computationally tractable, our first step consists of the discretization of the shapes X and Y and their corresponding metrics d_X and d_Y . We first assume the shapes to be represented as discrete approximations \hat{X} and \hat{Y} of the underlying continuous smooth surfaces X and Y , respectively. In our implementation, triangular meshes were used for this purpose, and points on \hat{X} and \hat{Y} were represented in barycentric coordinates.

5.1 Discrete Geodesic and Diffusion Distances

The first step is the computation of the discretized metrics on \hat{X} and \hat{Y} , denoted here as $\hat{d}_{\hat{X}}$ and $\hat{d}_{\hat{Y}}$, respectively. The discretized metrics $\hat{d}_{\hat{X}}$ and $\hat{d}_{\hat{Y}}$ can be represented as symmetric matrices of distances.

The geodesic metric is computed using the fast marching algorithm on triangulated surfaces (Kimmel and Sethian 1998).³ The main idea of fast marching is to simulate wave-front propagation on a triangular mesh, associating the time of arrival of the front with the distance traveled, assuming constant propagation speed. This allows to measure the distance from a point to all the rest of the points on the shape (computation that can be performed in linear time (Yatziv et al. 2006)). Repeating the process for all the points used as a source, the pairwise geodesic distances are obtained.

The diffusion metric is computed using formula (8). For this purpose, we first compute the discrete approximation of the Laplace-Beltrami operator on the mesh, which has the following generic form

$$(\Delta_{\hat{X}} f)_i = \frac{1}{a_i} \sum_j w_{ij} (f_i - f_j), \quad (13)$$

where $f : \hat{X} \rightarrow \mathbb{R}$ is a scalar function defined on the mesh \hat{X} , represented as a vector of function values at the vertices of the mesh, w_{ij} are weights, and a_i are normalization coefficients. In matrix notation, (13) can be written as

$$\Delta_{\hat{X}} f = A^{-1} L f, \quad (14)$$

³Basic versions of the fast marching algorithm is available online at <http://tosca.cs.technion.ac.il/book/resources.html>.

where $A = \text{diag}(a_i)$ and $L = \text{diag}(\sum_{l \neq i} w_{il}) - (w_{ij})$.

Different discretizations (Bobenko and Springborn 2007; Floater and Hormann 2005; Zhang 2004) of the Laplace-Beltrami lead to different choice of A and W . In this paper, we used the popular *cotangent weight* scheme (Meyer et al. 2003), in which $w_{ij} = \cot \alpha_{ij} + \cot \beta_{ij}$ (α_{ij} and β_{ij} are the two angles opposite to the edge between vertices i and j in the two triangles sharing the edge) for j in the 1-ring neighborhood of vertex i and zero otherwise, and a_i is proportional to the sum of the areas of the triangles sharing the vertex x_i . It can be shown (Wardetzky et al. 2008) that this discretization preserves many important properties of the continuous Laplace-Beltrami operator, such as positive semi-definiteness, symmetry, and locality, and in addition it is numerically consistent, i.e., converges to the continuous Laplace-Beltrami operator.⁴ The later property is especially important since it gives consistent approximation of the eigenfunctions across different shape triangulations.

By solving the *generalized eigendecomposition* problem (Lévy 2006)

$$A\phi = \lambda L\phi, \tag{15}$$

k smallest eigenvalues $\lambda_0, \dots, \lambda_k$ and corresponding eigenfunctions $\phi_0, \dots, \phi_k : \hat{X} \rightarrow \mathbb{R}$ of the discretized Laplace-Beltrami operator are computed. The discrete diffusion distance between points x_i, x_j is approximated as

$$d_{X,t} x_i, x_j \approx \sum_{l=1}^k e^{-2\lambda_l t} (\phi_{l;i} - \phi_{l;j})^2. \tag{16}$$

The boundary conditions used are Dirichlet. In some cases, e.g. when dealing with shapes with boundaries or missing parts, Neumann boundary conditions may be advantageous. In such cases, finite elements methods (FEM) used by Reuter et al. (2009) could be a more accurate discretization. We should also note that though we used triangular meshes, both the geodesic and diffusion distances can be efficiently computed for other surface representations as well, including point cloud data (Belkin et al. 2009). Therefore, the framework introduced here is not limited to meshes.

5.2 Discrete Gromov-Hausdorff Distance

Following the description in Sect. 2.5, our next goal is to construct a discrete approximation of a correspondence between the discrete shapes and evaluate its distortion. In the GMDS framework, correspondence is found by embedding points of \hat{X} (typically, subset of the vertices of \hat{X}) into \hat{Y} and vice versa (note that images of the points under these

⁴Here, convergence is understood in the following sense: the solution of a discrete PDE with the discrete Laplace-Beltrami operator converges to the solution of the continuous counterpart under some conditions on the mesh refinement.

embeddings do not necessarily fall on the vertices of \hat{Y} or \hat{X} , respectively). The distortion is computed by measuring the difference between the metric between the points in the original shape and the metric between their corresponding images under the embedding in the other shape.

Let us fix two sufficiently dense finite samplings $P = \{p_1, \dots, p_m\}$ and $Q = \{q_1, \dots, q_n\}$ of \hat{X} and \hat{Y} , respectively.⁵ A discrete correspondence between the shapes is defined as $C = (P \times Q') \cup (Q \times P')$, where $P' = \{p'_1, \dots, p'_n\}$ and $Q' = \{q'_1, \dots, q'_m\}$ are some (different) sets of samples on \hat{X} and \hat{Y} corresponding to Q and P , respectively. One can think of C as the union of the graphs of two discrete functions $\varphi : P \rightarrow \hat{Y}$ and $\psi : Q \rightarrow \hat{X}$, parametrizing the class of all discrete correspondences. P' and Q' are the optimization variables in the GMDS problem.

Given two sets P and P' on \hat{X} , we can construct an $m \times n$ distance matrix $D(P, P')$, whose elements are the distances $\hat{d}_{\hat{X}}(p_i, p'_j)$ (either geodesic or diffusion, depending on the context). In these terms, the distortion of the correspondence C can be written as⁶

$$\text{dis}(C) = \left\| \begin{pmatrix} D(P, P) & D(P, P') \\ D(P, P')^T & D(P', P') \end{pmatrix} - \begin{pmatrix} D(Q', Q') & D(Q', Q) \\ D(Q', Q)^T & D(Q, Q) \end{pmatrix} \right\|,$$

where $\|\cdot\|$ is some norm on the space of $(m+n) \times (m+n)$ matrices. The selection of the infinity norm $\|D\|_\infty = \max_{i,j} |d_{ij}|$ is consistent with the Gromov-Hausdorff distance, however, in practice more robust norms like the Frobenius norm $\|D\|_F^2 = \text{trace}(DD^T)$ are often preferable (see Bronstein et al. 2006a; Mémoli 2007; Mémoli and Sapiro 2005b for discussions on the regularization of the infinity norm in the Gromov-Hausdorff framework by other l_p norms).

The discretization of $\text{dis}(C)$ leads directly to a discretized approximation of the Gromov-Hausdorff distance between shapes, which can be expressed as

$$\hat{d}_{\text{GH}}(\hat{X}, \hat{Y}) := \frac{1}{2} \min_{P', Q'} \text{dis}(C).$$

Note that only P' and Q' participate as continuous minimization variables, while P and Q are constants (given

⁵We use the *farthest point sampling* (Hochbaum and Shmoys 1985) strategy to produce r -separated r -coverings of the shape, where the parameter r controls the radius of the sampling (see Mémoli and Sapiro 2005b and Bronstein et al. 2006a for details on the relationships between the Gromov-Hausdorff distance between these discrete coverings and the underlying continuous spaces). Such farthest sampling strategy can also be easily applied both to meshes and point cloud data.

⁶For a detailed explanation why this expression is equivalent to the Gromov-Hausdorff distance, see for example Burago et al. (2001), Mémoli (2008).

samples on the respective shapes). The above minimization problem is solved using GMDS—a numerical procedure resembling in its spirit the standard multidimensional scaling.

We use barycentric coordinates to represent points on \hat{X} and \hat{Y} . In these coordinates, a point p_i lying in a triangle t_i on \hat{X} is represented as a convex combination of the triangle vertices (corresponding to the indices t_i^1 , t_i^2 , and t_i^3) with the weights $u_i = (u_i^1, u_i^2, u_i^3)^T$. We will denote by $T = (t_1, \dots, t_m)^T$ the vector of triangle indices and by $U = (u_1, \dots, u_m)$ the $3 \times m$ matrix of coordinates corresponding to the sampling P . Similarly, the samplings P' , Q , and Q' are represented as (T', U') , (S, V) and (S', V') . For the sake of notation simplicity, we are going to use these interchangeably.

It was shown in Bronstein et al. (2008) that a first-order approximation of a geodesic distance between p'_i and p'_j on \hat{X} can be expressed as the quadratic form

$$D_{ij}(P', P') \approx u_i^T \begin{pmatrix} D_{t_i^1, t_j^1}(P, P) & D_{t_i^1, t_j^2}(P, P) & D_{t_i^1, t_j^3}(P, P) \\ D_{t_i^2, t_j^1}(P, P) & D_{t_i^2, t_j^2}(P, P) & D_{t_i^2, t_j^3}(P, P) \\ D_{t_i^3, t_j^1}(P, P) & D_{t_i^3, t_j^2}(P, P) & D_{t_i^3, t_j^3}(P, P) \end{pmatrix} u_j.$$

Other distance terms are expressed similarly. Using tensor notation, we can write

$$\text{dis}(C) \approx \|(U, U') \mathcal{D}_{\hat{X}}(T, T')(U, U') - (V, V') \mathcal{D}_{\hat{Y}}(S, S')(V, V')\|_{\mathbb{F}}^2,$$

where $\mathcal{D}_{\hat{X}}(T, T')$ is a rank four tensor whose ij -th elements are defined as the 3×3 distance matrices above, and $\mathcal{D}_{\hat{Y}}(S, S')$ is defined in a similar way.

The resulting objective function $\text{dis}(C)$ is a fourth-order polynomial with respect to the continuous coordinates U', V' , also depending on the discrete index variables T' and S' . However, when all indices and all coordinate vectors except one, say, u'_i , are fixed, the function becomes convex and quadratic with respect to u'_i . A closed-form minimizer of $\text{dis}(u'_i)$ is found under the constraints $u'_i \geq 0$ and $u'_i^1 + u'_i^2 + u'_i^3 = 1$, guaranteeing that the point p'_i remains within the triangle t'_i . The GMDS minimization algorithm proceeds iteratively by selecting u'_i or v'_i corresponding to the largest gradient of the objective function, updating it according to the closed-form minimizer, and updating the corresponding triangle index to a neighboring one in case the solution is found on the boundary of the triangle. The reader is referred to Bronstein et al. (2008) for further implementation details.

6 Experimental Results

In this section, we present experimental results showing the advantages of using the diffusion distance instead of the geo-

desic one for shape comparison with topological changes under the metric model. We perform experiments both with shape similarity and correspondence (matching).

6.1 Data

The experiments were performed on meshes taken from the TOSCA dataset (Bronstein et al. 2009).⁷ Seven classes of objects were used (see Figs. 2 and 3): centaur, horse, cat, dog, two male shapes (Michael and David), and one female shape (Victoria). In each class, the shape underwent different types of transformations (Fig. 4). The transformations included null (no transformation), isometry (near-isometric bending), topology (connectivity change obtained by welding some of the shape vertices), topology + isometry, triangulation (different triangulations of the same shape), and partiality (missing information, obtained by making holes in the shape). These transformations simulate typical artifacts and deformations encountered in real life in the shapes reconstructed from 3D scanners data. Multiple instances of each transformation were created for each shape class. Totally, the dataset consisted of 245 different shape instances (centaur: 22, horse: 31, Michael: 55, David: 29, Victoria: 39, dog: 32, cat: 37). Shapes in the null transformations were sampled at approximately 1500 points; shapes in triangulation transformation had between 250 to 2000 points.


6.2 Similarity

In the first set of experiments, we compared five different approaches for shape similarity: the proposed Gromov-Hausdorff framework with diffusion distances, Gromov-Hausdorff framework with geodesic distances (Bronstein et al. 2006b; Mémoli and Sapiro 2005b), canonical forms with geodesic distances (Elad and Kimmel 2003), canonical forms with diffusion distances (diffusion maps (Coifman et al. 2005)), and, for reference, Rustamov's GPS embedding (Rustamov 2007). These approaches were used to compute the similarity of shapes, with the goal to distinguish between different classes of shapes while being insensitive to intra-class transformations.

An L_2 version of the Gromov-Hausdorff distance was computed using the GMDS algorithm with a matrix of geodesic or diffusion distances measured on the shape, as described in Sect. 5. A multiresolution scheme was used for GMDS with $m = n = 64$ points at the finest scale. Initialization was at 8 points at the lowest scale, using the branch-and-bound algorithm described in Raviv et al. (2007). The typical convergence time for GMDS was about 10–30 sec.

⁷Dataset is available online at http://tosca.cs.technion.ac.il/book/resources_data.html.

Fig. 2 The base shapes used in the experiments: centaur, horse, dog, cat, woman, and two males



Discrete geodesic distances were computed using fast marching (Kimmel and Sethian 1998), as described in Sect. 5.1. As detailed above, diffusion distances with time parameter $t = 2500$ were computed using formula (8) by performing eigendecomposition of the Laplace-Beltrami operator on the mesh. $k = 200$ eigenvalues were used. These values were experimentally found to produce good results.

Canonical forms with geodesic distances were computed by applying least squares MDS to the matrix of geodesic distances measured on the shape, embedding them into \mathbb{R}^3 (Elad and Kimmel 2003). From our experiments, this dimensionality was sufficient and increasing the dimensionality of the embedding space did not lead to significant improvement. For canonical forms with diffusion distances, the embedding was computed by taking the first three non-trivial eigenvectors of the Laplace-Beltrami operator of the shape to obtain a representation in \mathbb{R}^3 (Belkin and Niyogi 2003; Coifman et al. 2005).⁸ In both cases, the canonical forms in \mathbb{R}^3 were compared using ICP (while in the infinite dimensional case this is not needed for the diffusion embedding, since the embedding is unique and invariant, it helps in the finite and discrete scenario which provides only an approximation of the embedding).


GPS embedding was computed using $k = 200$ eigenfunctions of the Laplace-Beltrami operator. Commute time distributions were approximated using histograms with 500 bins.

The similarity quality was first quantitatively measured by plotting the *receiver operating characteristic* (ROC) curves for each approach, representing a tradeoff between the *false acceptance rate* (FAR) and the *false rejection rate* (FRR). Each ROC curve was computed as follows: the confusion matrix (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 FRR 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).

In the first experiment, we used a subset of the whole dataset, excluding topology and partiality (missing data) transformations. The partial dataset consisted of 101 different shape instances (centaur: 10, horse: 12, Michael: 24, David: 11, Victoria: 16, dog: 13, cat: 15). All the instances within each shape class were approximately isometries. Figures 5 and 6 show the confusion matrices (null transformation to all) obtained using the Gromov-Hausdorff and canonical shape frameworks with geodesic and diffu-

⁸We used three eigenvectors since ICP is impractical in higher dimensions.

Fig. 3 A few representative shapes used in the experiments



sion distances. Figure 7 shows the ROC curve (FAR vs FRR), representing the accuracy of shape recognition for all the five approaches. The EER is shown in Table 1 (first row).⁹

In the second experiment, we used the entire dataset, including topology and partiality transformations. Figures 8 and 9 show the confusion matrices obtained using different approaches. Figure 10 shows the ROC curve (FAR vs FRR), representing the accuracy of shape recognition for all the approaches. The corresponding EER is shown in Table 1 (second row).

We can make a few observations at this point. First, the Gromov-Hausdorff framework is more accurate than the canonical forms approach and than methods based on

⁹A slightly better performance of the Gromov-Hausdorff framework with diffusion distances over geodesic ones in case of Isometry transformations may be explained by the fact that the transformations are not truly isometric and involve some elastic deformations.

Table 1 EER (in %) on the two datasets obtained by the five different approaches. Experiment I included shapes with Null, Isometry, and Triangulation transformations. Experiment II included shapes with Null, Isometry, Triangulation, Topology, Isometry + Topology, and Partiality transformations. Best results are in bold

	Gromov-Hausdorff		Canonical forms		GPS
	Geodesic	Diffusion	Geodesic	Diffusion	
Experiment I	4.95	2.22	15.01	17.82	25.74
Experiment II	15.49	2.02	16.32	16.44	16.97

distance distributions. This corresponds to results previously reported in the literature (Mémoli and Sapiro 2005b; Bronstein et al. 2006b), and reflects the fact that canonical forms can be regarded as an approximation to a distance between metric spaces. The inferior performance of GPS embedding is attributed to the discriminative power loss when passing from shape geometry to distributions.

Fig. 4 Different types of shape transformations used in the experiments: null, isometry, topology, and partiality




Fig. 5 Confusion matrices for shape similarity using the Gromov-Hausdorff framework with geodesic (*top*) and diffusion (*bottom*) distances on the data of the first experiment. Shown are distances from null transformations of each shape class (*rows*) to all instances in the dataset (*columns*). Brighter color represents larger distances (smaller similarity)





Fig. 6 Confusion matrices for shape similarity using the canonical forms framework with geodesic (*top*) and diffusion (*bottom*) distances on the data of the first experiment. Shown are distances from null transformations of each shape class (*rows*) to all instances in the dataset (*columns*). Brighter color represents larger distances (smaller similarity)



Second, observe the behavior of the Gromov-Hausdorff framework with diffusion and geodesic distances in the first and second experiments. Both approaches (GMDS with diffusion and geodesic distances) behave well in the first experiment where the dataset contains only nearly-isometric transformations (EER of 2.22% and 4.95%, respectively). However, the picture changes notably in the second experiment when non-isometric transformations (topology and partiality) are added: the EER in the geodesic case increases more than three times to 17.55%, while in the diffusion case it remains approximately the same (2.28%). This is a clear evidence that the proposed approach is less sensitive to topological noise and missing parts.

Third, the proposed approach is insensitive to different sampling and triangulations of the shapes. Finally, one can observe from the confusion matrices that horse-like shapes (centaur and horse) and human-like shapes (two males and female) are similar, though still distinguishable from each other.

As an additional shape recognition quality criterion, for each shape class we looked at the inter/intra cluster ratio,

defined as the ratio between the average inter-cluster distance (average of all the distances from the null shape of the class to all the instances of shapes from other classes) and the average intra-cluster distance (average of all the distances between the transformations of a shape of the same class). Small intra-cluster distance means different deformations of the shape are considered similar; large inter-cluster distance means different shape classes are considered dissimilar. Higher ratio of the two implies better ability to discriminate between different shape classes and lower sensitivity to transformations.

Table 2 shows the inter/intra cluster distance ratio for GMDS with geodesic and diffusion distances on the dataset of the second experiment. GMDS with diffusion distances performs significantly better in the case of topological transformations and missing data (average ratio of 2.76 for geodesic versus 5.61 for diffusion).

6.3 Correspondence

In the second class of experiments, we compared the performance of the Gromov-Hausdorff framework with dif-

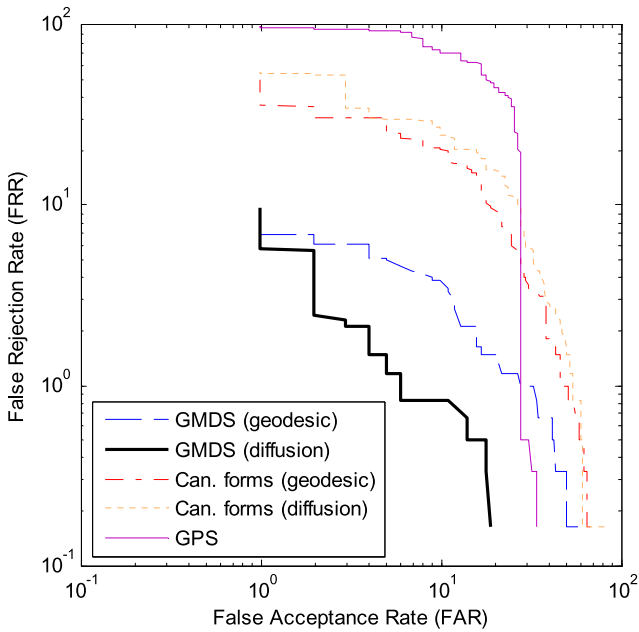


Fig. 7 ROC curves for the first experiment

fusion and geodesic distances on the problem of shape matching and correspondence. TOSCA shapes with known groundtruth correspondence were used. The correspondence computed by both methods was compared to the groundtruth.

Table 2 Inter-cluster and intra-cluster average distance and the intra/inter-cluster average distance ratio obtained by GMDS with geodesic and diffusion distances on the second (full) dataset. Distances were normalized by maximum distance. Larger inter/intra cluster ratio means better recognition

	Geodesic			Diffusion		
	Intra	Inter	Inter/Intra	Intra	Inter	Inter/Intra
Centaur	0.067	0.397	5.94	0.074	0.516	6.94
Horse	0.064	0.652	10.24	0.041	0.679	16.65
Michael	0.118	0.327	2.76	0.088	0.413	4.68
David	0.069	0.255	3.70	0.063	0.490	7.83
Victoria	0.138	0.288	2.08	0.059	0.402	6.82
Dog	0.065	0.362	5.61	0.068	0.676	9.96
Cat	0.070	0.302	4.31	0.060	0.513	8.55

Fig. 8 Confusion matrices for shape similarity using the Gromov-Hausdorff framework with geodesic (*top*) and diffusion (*bottom*) distances on the data of the second experiment. Shown are distances from null transformations of each shape class (*rows*) to all instances in the dataset (*columns*). Brighter color represents larger distances (smaller similarity)

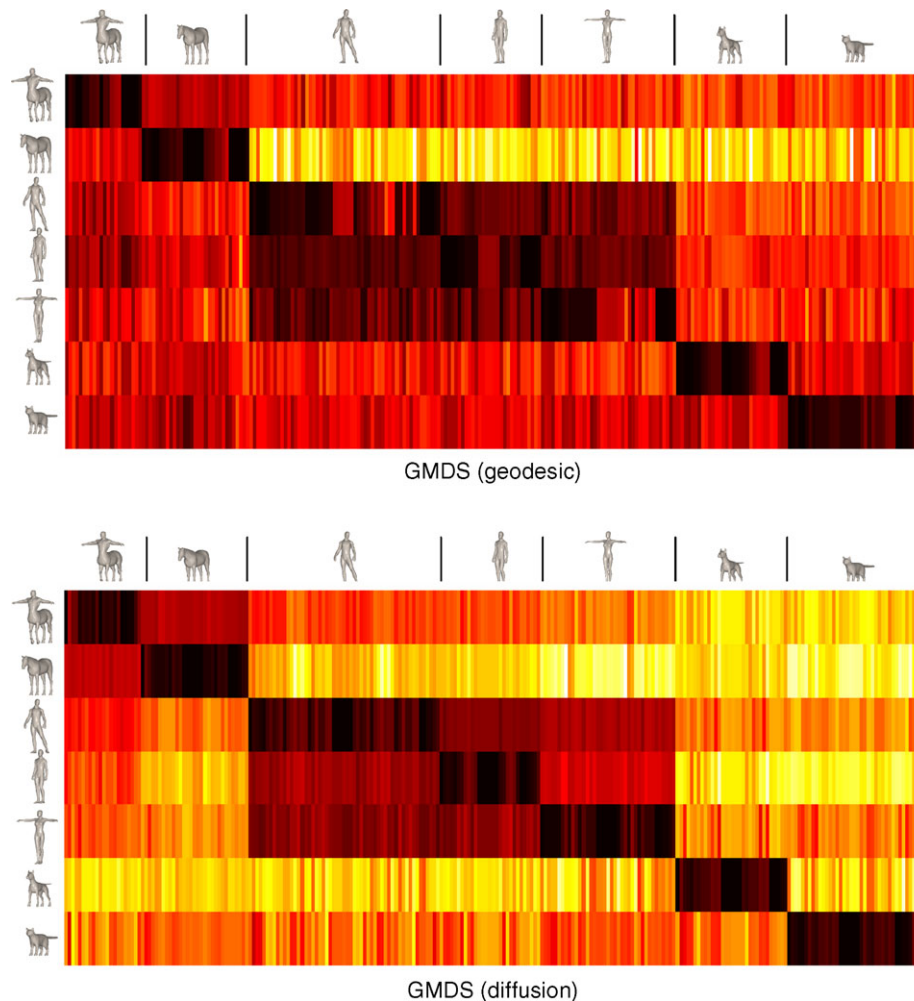
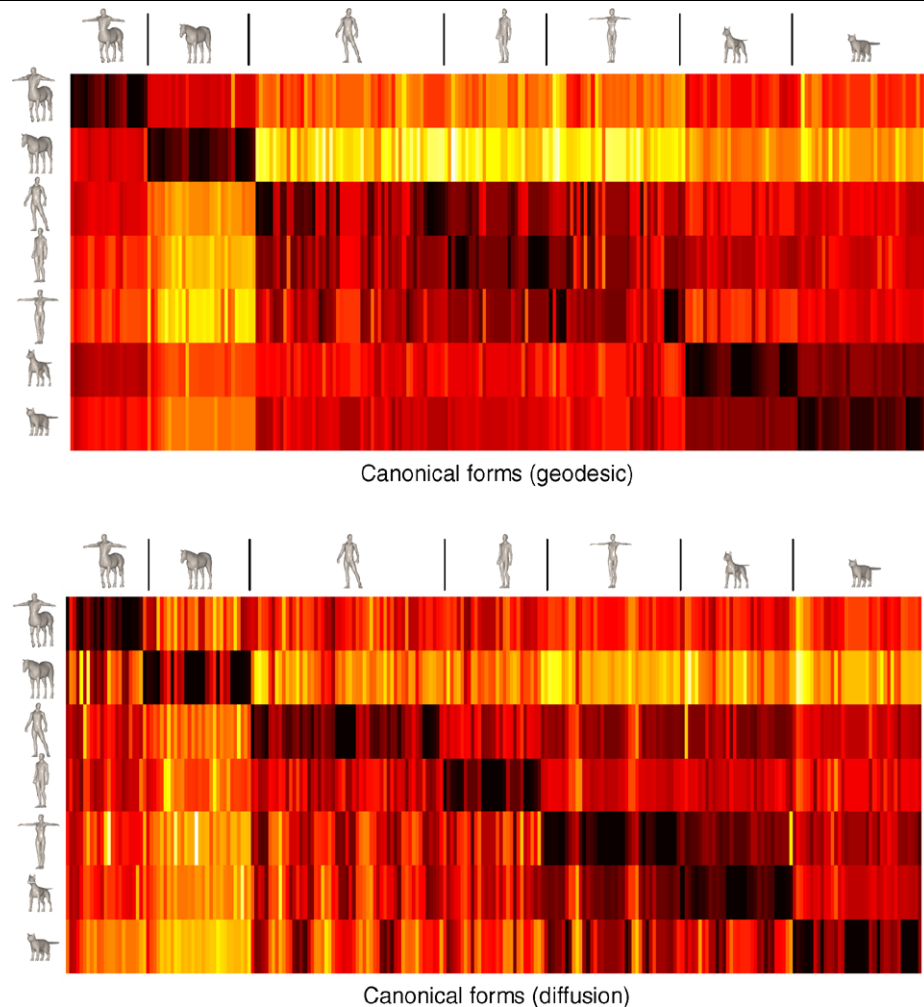


Fig. 9 Confusion matrices for shape similarity using the canonical forms framework with geodesic (*top*) and diffusion (*bottom*) distances on the data of the second experiment. Shown are distances from null transformations of each shape class (*rows*) to all instances in the dataset (*columns*). Brighter color represents larger distances (smaller similarity)



To quantify the correspondence quality, let $C = P \times Q$ be the computed correspondence and let $C_0 = P_0 \times Q_0$ be the groundtruth correspondence. Let us denote $P'_0 = \{p \in P : (p, q) \in C_0 \text{ and } q \in Q\}$ (the true corresponding points to Q in X) and $Q'_0 = \{q \in Q : (p, q) \in C_0 \text{ and } p \in P\}$ (the true corresponding points to P in Y). We define the correspondence quality as

$$d(C, C_0) = \frac{1}{2} (\|D_X(P_0, P'_0)\|_\infty + \|D_Y(Q_0, Q'_0)\|_\infty),$$

where D_X and D_Y are the corresponding geodesic distance matrices. In simple words, this criterion finds the maximum geodesic distance between the set of computed corresponding points and the true corresponding points. Geodesic distances are used in $d(C, C_0)$.

Tables 3 and 4 show the correspondence quality between different transformations of the male and dog objects, obtained by GMDS with geodesic and diffusion distances, respectively. When the transformation is an approximate isometry (first row in both tables), both methods produce correspondence of approximately equal quality (24.03 and

Table 3 Correspondence quality for the male object

	Geodesic	Diffusion
Isometry	24.03	20.88
Topology	29.79	6.13

Table 4 Correspondence quality for the dog object

	Geodesic	Diffusion
Isometry	30.11	30.51
Topology	21.26	7.96

20.88 for male, 30.11 and 30.51 for dog). The correspondence error is non-zero because the transformation is not truly isometric. However, the picture is different in the case of a topological transformation (obtained by locally changing the connectivity only without deforming the shape). The correspondence error of GMDS with geodesic distances is almost five times larger than that of GMDS with diffusion

distances in the male shape (29.79 versus 6.13); correspondence error of GMDS with geodesic distances is over twice

larger than that of GMDS with diffusion distances in the dog shape (21.26 versus 7.96). This is yet another indication that the proposed approach is less sensitive to topological changes.

Figures 11 and 12 provide a visualization of these phenomena. Figure 11 shows a correspondence between different near-isometric deformations of the dog object. The shapes are nearly isometric and both methods produce good correspondence. Figure 12 shows the correspondence between male shapes with different connectivity. The correspondence produced by GMDS with geodesic distances is bad (note for example the magnified part showing the corresponding points on the feet), while the one obtained with diffusion distances is good.

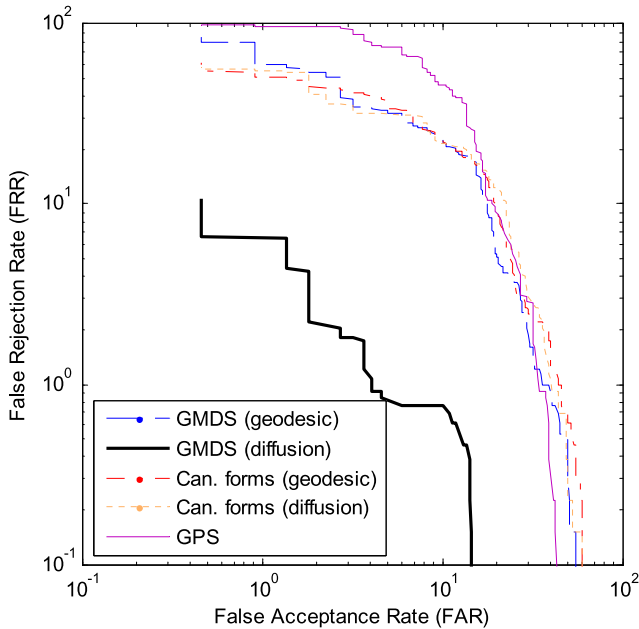


Fig. 10 ROC curves for the second experiment

7 Conclusions

In this paper, we addressed the problems of shape similarity and correspondence in the presence of topological changes. We used the metric approach, modeling shapes as metric spaces and posing the problem of shape similarity as the similarity between metric spaces. We showed that the Gromov-Hausdorff distance, previously applied to geodesic

Fig. 11 Correspondence between two isometric instances of a dog shape. *Left to right:* ground truth correspondence, correspondence using GMDS with geodesic distances, and correspondence using GMDS with diffusion distances. Voronoi regions around the corresponding points are shown in different colors


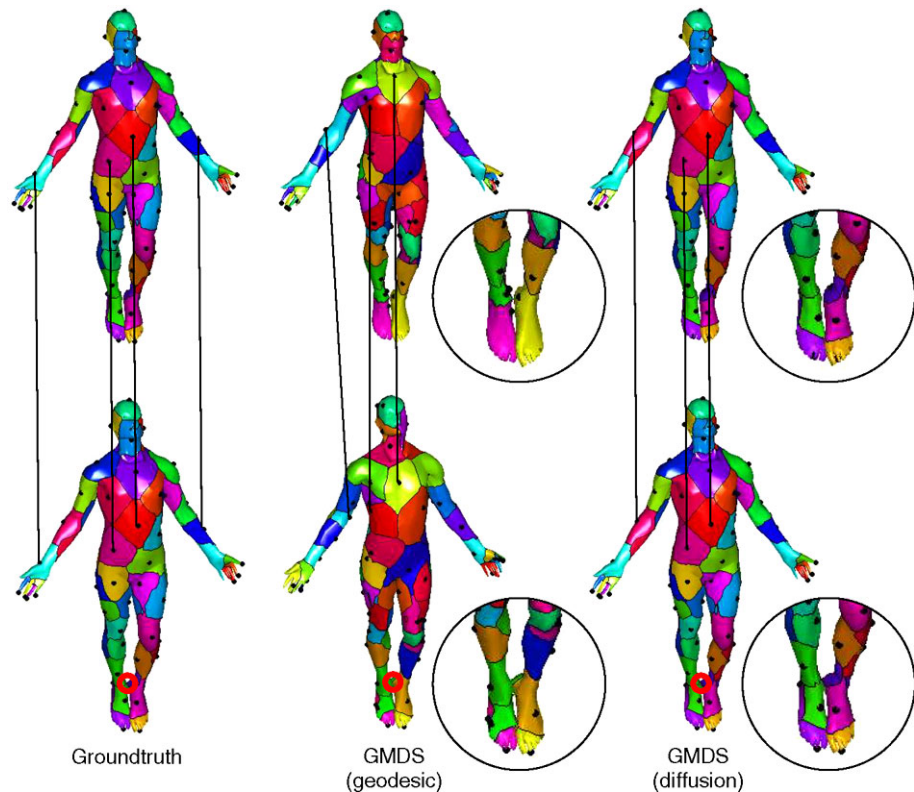


Fig. 12 Correspondence between two instances of a male shape with different connectivity. *Left to right:* ground truth correspondence, correspondence using GMDS with geodesic distances, and correspondence using GMDS with diffusion distances. The different topology is obtained by welding the vertices at a point marked by the *red circle*. Voronoi regions around the corresponding points are shown in different colors



metrics for bending-invariant shape recognition, can be applied to shapes endowed with diffusion geometry, leading to a topologically robust approach for non-rigid shape comparison and matching.

In particular, we showed how replacing the geodesic distance between pair of surface points by the diffusion distance, leads to recognition improvements for data with topological variations such as holes and connectivity changes. This robustness to holes is a first step toward the recognition of partial shapes, since the missing portion can be considered as a “hole.”

In addition to the practical consequences brought by the proposed framework, as with the works on canonical forms, these results suggest moving beyond the classical use of geodesics for intrinsic non-rigid shape matching in the Gromov-Hausdorff metric framework. Thereby, the use of other intrinsic distances, as well as other kernels in the diffusion framework, deserves further study. The combination of such distances may lead to further performance improvements as well. The combination of the Gromov-Hausdorff and diffusion framework with topological features as those described in Singh et al. (2007) is of great practical interest as well. The close relationship of the diffusion distance with graph methods raises the issue of considering results from the graph theory community (Chung 1997), including graph matching algorithms (Farin et al. 2003), and their relationship with the framework here detailed.

At a theoretical level, many existing questions are emphasized and new ones are posed. The geometry of the shape space defined by the Gromov-Hausdorff metric with diffusion distances is of great theoretical and practical significance. In addition, the study of the classes of shape transformations under which the diffusion geometry is invariant (“diffusion isometries”) and their relation to “geodesic isometries” should give an important insight on the cases in which each of the methods is preferable.

Acknowledgements GS thanks Dr. Facundo Memoli, Stanford University, with whom the travel into the Gromov-Hausdorff and metric shape comparison world started. We thank the anonymous reviewers, they helped us improve the presentation of the paper and motivated us to provide additional details and examples.

References

- Belkin, M., & Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 13, 1373–1396.
- Belkin, M., Sun, J., & Wang, Y. (2009). Constructing Laplace operator from point clouds in R^d . In *Proc. SIAM symp. discrete algorithms* (pp. 1031–1040).
- Besl, P. J. & McKay, N. D. (1992). A method for registration of 3D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 14, 239–256.
- Bobenko, A. I., & Springborn, B. A. (2007). A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete and Computational Geometry*, 38(4), 740–756.

- Boutin, M., & Kemper, G. (2004). On reconstructing n -point configurations from the distribution of distances or areas. *Advances in Applied Mathematics*, 32(4), 709–735.
- Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2005). Three-dimensional face recognition. *International Journal of Computer Vision (IJCV)*, 64(1), 5–30.
- Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006a). Efficient computation of isometry-invariant distances between surfaces. *SIAM Journal on Scientific Computing*, 28(5), 1812–1836.
- Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006b). Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Science (PNAS)*, 103(5), 1168–1172.
- 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). Expression-invariant representation of faces. *IEEE Transactions on Image Processing*, 16(1), 188–197.
- Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2007c). Rock, paper, and scissors: extrinsic vs. intrinsic similarity of non-rigid shapes. In *Proc. int. conf. computer vision (ICCV)*.
- Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2008). *Numerical geometry of nonrigid shapes*. Berlin: Springer.
- Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2009). Topology-invariant similarity of nonrigid shapes. *International Journal of Computer Vision (IJCV)*, 81(3), 281–301.
- Bronstein, M. M., & Bronstein, A. M. (2009). *On a relation between shape recognition algorithms based on distributions of distances* (Technical Report CIS-2009-14). Technion.
- Burago, D., Burago, Y., & Ivanov, S. (2001). *A course in metric geometry. Graduate studies in mathematics*, Vol. 33. Providence: AMS.
- Chen, Y., & Medioni, G. (1991). Object modeling by registration of multiple range images. In *Proc. conf. robotics and automation*.
- Chung, F. R. K. (1997). *Spectral graph theory*. Providence: AMS.
- Coifman, R. R., & Lafon, S. (2006). Diffusion maps. *Applied and Computational Harmonic Analysis*, 21, 5–30.
- Coifman, R. R., Lafon, S., Lee, A. B., Maggioni, M., Nadler, B., Warner, F., & Zucker, S. W. (2005). Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Science (PNAS)*, 102(21), 7426–7431.
- Cox, T. F., & Cox, M. A. A. (1994). *Multidimensional scaling*. London: Chapman and Hall.
- Do Carmo, M. P. (1992). *Riemannian geometry*. Boston: Birkhäuser.
- Elad, A., & Kimmel, R. (2002). Spherical flattening of the cortex surface. In *Geometric methods in bio-medical image processing* (Vol. 2191, pp. 77–89). Berlin: Springer.
- Elad, A., & Kimmel, R. (2003). On bending invariant signatures for surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 25(10), 1285–1295.
- Farin, D., de With, P. H. N., Effelsberg, W., & De, P. H. N. (2003). Recognition of user-defined video object models using weighted graph homomorphism. In *SPIE IVCP*.
- Floater, M. S., & Hormann, K. (2005). Surface parameterization: a tutorial and survey. *Advances in Multiresolution for Geometric modelling 1*.
- Gal, R., Shamir, A., & Cohen-Or, D. (2007). Pose-oblivious shape signature. *IEEE Transactions on Visualization and Computer Graphics*, 13(2), 261–271.
- Gromov, M. (1981). Structures métriques pour les variétés riemanniennes. *Textes mathématiques*, no. 1.
- Grossman, R., Kiryati, N., & Kimmel, R. (2002). Computational surface flattening: a voxel-based approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 24(4), 433–441.
- Hamza, A. B., & Krim, H. (2005). Probabilistic shape descriptor for triangulated surfaces. In *Proc. ICIP* (pp. 1041–1044).
- Hein, M., Audibert, J.-Y., & von Luxburg, U. (2007). Graph Laplacians and their convergence on random neighborhood graphs. *Journal of Machine Learning Research*, 8, 1325–1368.
- Hochbaum, D. S., & Shmoys, D. B. (1985). A best possible heuristic for the k -center problem. *Mathematics of Operations Research*, 10, 180–184.
- Jain, V., & Zhang, H. (2007). A spectral approach to shape-based retrieval of articulated 3D models. *Computer-Aided Design*, 39, 398–407.
- Jones, P. W., Maggioni, M., & Schul, R. (2008). Manifold parametrizations by eigenfunctions of the Laplacian and heat kernels. *Proceedings of the National Academy of Sciences*, 105(6), 1803.
- Katz, S., Leifman, G., & Tal, A. (2005). Mesh segmentation using feature point and core extraction. *The Visual Computer*, 21(8), 649–658.
- Kimmel, R., & Sethian, J. A. (1998). Computing geodesic paths on manifolds. *Proceedings of the National Academy of Science (PNAS)*, 95(15), 8431–8435.
- Lafon, S. (2004). *Diffusion maps and geometric harmonics*. Ph.D. dissertation, Yale University.
- Lévy, B. (2006). Laplace-Beltrami eigenfunctions towards an algorithm that “understands” geometry. In *Int’l conf. shape modeling and applications*.
- Ling, H., & Jacobs, D. (2005a). Deformation invariant image matching. In *Proc. ICCV*.
- Ling, H., & Jacobs, D. (2005b). Using the inner-distance for classification of articulated shapes. In *Proc. CVPR*.
- Mahmoudi, M., & Sapiro, G. (2009). Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models*, 71(1), 22–31.
- Mateus, D., Horaud, R. P., Knossow, D., Cuzzolin, F., & Boyer, E. (2008). Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In *Proc. IEEE conference on computer vision and pattern recognition (CVPR)*.
- Mémoli, F. (2007). On the use of Gromov-Hausdorff distances for shape comparison. In *Point based graphics*, Prague.
- Mémoli, F. (2008). Gromov-Hausdorff distances in Euclidean spaces. In *Proc. non-rigid shapes and deformable image alignment (NORDIA)*.
- Mémoli, F. (2009). Spectral Gromov-Wasserstein distances for shape matching. In *Proc. non-rigid shape analysis and deformable image alignment*.
- Mémoli, F., & Sapiro, G. (2001). Fast computation of weighted distance functions and geodesics on implicit hyper-surfaces. *Journal of Computational Physics*, 173(2), 730–764.
- Mémoli, F., & Sapiro, G. (2005a). Distance functions and geodesics on submanifolds of R^d and point clouds. *SIAM Journal of Applied Mathematics*, 65, 1227–1260.
- Mémoli, F., & Sapiro, G. (2005b). A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5, 313–346.
- Mémoli, F., Sapiro, G., & Thompson, P. (2006). Geometric surface and brain warping via geodesic minimizing Lipschitz extensions. In *MFCA-2006 international workshop on mathematical foundations of computational anatomy*, Copenhagen.
- Meyer, M., Desbrun, M., Schroder, P., & Barr, A. H. (2003). Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics, III*, 35–57.
- Olver, P. (1999). *Classical invariant theory*. Cambridge: Cambridge University Press.
- Osada, R., Funkhouser, T., Chazelle, B., & Dobkin, D. (2002). Shape distributions. *ACM Transactions on Graphics (TOG)*, 21(4), 807–832.
- Ovsjanikov, M., Sun, J., & Guibas, L. (2008). Global intrinsic symmetries of shapes. *Computer Graphics Forum*, 27, 1341–1348.

- Ovsjanikov, M., Bronstein, A. M., Bronstein, M. M., & Guibas, L. (2009). Shape Google: a computer vision approach for invariant retrieval of non-rigid shapes. In *Proc. non-rigid shape analysis and deformable image alignment*.
- Qiu, H., & Hancock, E. R. (2007). Clustering and embedding using commute times. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11), 1873–1890.
- Raviv, D., Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2007). Symmetries of non-rigid shapes. In *Proc. workshop on non-rigid registration and tracking through learning (NRTL)*.
- Reuter, M., Wolter, F.-E., & Peinecke, N. (2006). Laplace-Beltrami spectra as “shape-DNA” of surfaces and solids. *Computer Aided Design*, 38, 342–366.
- Reuter, M., Biasotti, S., Giorgi, D., Patanè, G., & Spagnuolo, M. (2009). Discrete Laplace-Beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33, 381–390.
- Ruggeri, M. R., & Saube, D. (2008). Isometry-invariant matching of point set surfaces. In *Proc. eurographics 2008 workshop on 3D object retrieval*.
- Ruggeri, M. R., Patanè, G., Spagnuolo, M., & Saube, D. (2009). Spectral-driven isometry-invariant matching of 3D shapes. *International Journal of Computer Vision*.
- Rustamov, R. M. (2007). Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. symp. geometry processing (SGP)* (pp. 225–233).
- Schwartz, E. L., Shaw, A., & Wolfson, E. (1989). A numerical solution to the generalized mapmaker’s problem: flattening nonconvex polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 11, 1005–1008.
- Singh, G., Mémoli, F., & Carlsson, G. (2007). Topological methods for the analysis of high dimensional data sets and 3d object recognition. In *Point based graphics*, Prague.
- Spira, A., & Kimmel, R. (2004). An efficient solution to the eikonal equation on parametric manifolds. *Interfaces and Free Boundaries*, 6(3), 315–327.
- Sun, J., Ovsjanikov, M., & Guibas, L. (2009). A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. symp. geometry processing (SGP)*.
- Tennenbaum, J. B., de Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500), 2319–2323.
- Walter, J., & Ritter, H. (2002). On interactive visualization of high-dimensional data using the hyperbolic plane. In *Proc. int’l conf. knowledge discovery and data mining (KDD)* (pp. 123–131).
- Wardetzky, M., Mathur, S., Kälberer, F., & Grinspun, E. (2008). Discrete Laplace operators: no free lunch. In *Conf. computer graphics and interactive techniques*.
- Weber, O., Devir, Y. S., Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2008). Parallel algorithms for approximation of distance maps on parametric surfaces. *ACM Transactions on Graphics*, 27(4).
- Wood, Z., Hoppe, H., Desbrun, M., & Schröder, P. (2004). Removing excess topology from isosurfaces. *ACM Transactions on Graphics*, 23.
- Yatziv, L., Bartesaghi, A., & Sapiro, G. (2006). $O(N)$ implementation of the fast marching algorithm. *Journal of Computational Physics*, 212(2), 393–399.
- Zhang, H. (2004). Discrete combinatorial Laplacian operators for digital geometry processing. In *SIAM conference on geometric design* (pp. 575–592).
- Zigelman, G., Kimmel, R., & Kiryati, N. (2002). Texture mapping using surface flattening via multi-dimensional scaling. *IEEE Transactions Visualization and Computer Graphics (TVCG)*, 9(2), 198–207.