Quantifying 3D shape similarity using maps: Recent trends, applications and perspectives

S. Biasotti[1], A. Cerri[1], A. Bronstein[2], M. Bronstein[3]

[1] Istituto di Matematica Applicata e Tecnologie Informatiche, Consiglio Nazionale delle Ricerche, Italy
[2] School of Electrical Engineering, Tel Aviv University, Israel
[3] Institute of Computational Science, University of Lugano (USI), Switzerland

Abstract

Shape similarity is an acute issue in Computer Vision and Computer Graphics that involves many aspects of human perception of the real world, including judged and perceived similarity concepts, deterministic and probabilistic decisions and their formalization. 3D models carry multiple information with them (e.g., geometry, topology, texture, time evolution, appearance), which can be thought as the filter that drives the recognition process. Assessing and quantifying the similarity between 3D shapes is necessary to explore large dataset of shapes, and tune the analysis framework to the user's needs. Many efforts have been done in this sense, including several attempts to formalize suitable notions of similarity and distance among 3D objects and their shapes.
In the last years, 3D shape analysis knew a rapidly growing interest in a number of challenging issues, ranging from deformable shape similarity to partial matching and view-point selection. In this panorama, we focus on methods which quantify shape similarity (between two objects and sets of models) and compare these shapes in terms of their properties (i.e., global and local, geometric, differential and topological) conveyed by (sets of) maps. After presenting in detail the theoretical foundations underlying these methods, we review their usage in a number of 3D shape application domains, ranging from matching and retrieval to annotation and segmentation. Particular emphasis will be given to analyse the suitability of the different methods for specific classes of shapes (e.g. rigid or isometric shapes), as well as the flexibility of the various methods at the different stages of the shape comparison process. Finally, the most promising directions for future research developments are discussed.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.5]: Computational Geometry and Object Modeling—Computer Graphics [I.3.6]: Methodology and Techniques—

1. Introduction

The technological advances in data acquisition and storage have been inducing an exponential growth of the volume of available data, also deeply modifying the approach we access to them. Indeed, such data are often stored using different formats and have to be analysed, interpreted and catalogued with significant computational efforts and experts' commitment. When dealing with large sets of data, storage is only one of the aspects to be solved: developing an *automatic* way to define *similarity* distances so that the database is indexed and efficiently queried is also a necessity.

For many decades, psychologists have studied how humans perceive a shape and how this perception affects the everyone' decisions and understandings. The literature concerning the psychological background of similarity assesses that the association between the perception and the conceptual model is in the mind of the observer [SJ99, Tve77, AP88, Ash92, Koe90]. Indeed, similarity involves many aspects of human perception of the real world, including judged and perceived resemblance, deterministic and probabilistic perceptions and decisions, and so on [SB11]. In summary, the formalization of the concept of shape similarity is a complex interaction process that depends on the observer, the visual content and the context.

From the computational point of view, the need of methods and algorithms able to quantify how much and where two or more shapes differ calls for a formal definition of the notion of shape (dis)similarity. Capturing the information contained in shape data typically takes the form of comput-

DOI: 10.2312/egst.20141039

ing meaningful properties, turning them into invariants, and defining opportune distances in the shape description space.

1.1. STAR focus and contribution

In this paper we focus on similarity assessment, by specifically targeting the 3D shape digital world. By shape similarity we mean *quantifying* through some measure how a shape resembles to another.

Assessing shape similarity is definitively not an easy task: first of all, there is neither a single *best* shape characterization nor a single *best* similarity measure, and shape comparison largely depends on the type of shapes to be analysed and on the properties that are considered relevant in the comparison process. An intuition of this is given by Figure 1, showing some models from the SHREC'08 classification benchmark [GM08]. Three categorization levels are proposed, reflecting as many different ways to conceive shape similarity. Indeed, models are classified with respect to functional (semantic), structural and geometric criteria.

In the rapidly growing field of the evaluation of 3D shape similarity, a number of strategies have been proposed, spanning from the direct definition of metrics between two objects to the approximation of a transformation between shapes and the evaluation of its distortion. While at the beginning the main efforts were mainly devoted to the transposition of well-known metrics into application domains [VH01, TV04, BKS*05, FKMS05, YLZ07, DP06, BKSS07, TV08], during last years the focus is moving to new techniques and more complex frameworks that allow a larger flexibility in the definition of similarity [SOG09, BBK*10, OBCS*12, DF04, DF07]. A number of interesting solutions comes from advances in pure and applied mathematics [BCF*08,CZ09,CFF*13,EH10,SOG09], as well as from the re-reading of classical mathematical theories and their adaptation to the discrete setting [BBK*10,OBCS*12,RBB*11].

In this scenario, we aim at providing a reasoned overview of the most recent advances in Computer Graphics, driven by the following guidelines:

- Methods that extract the shape structure through *functions* or *distances*. On the one hand, real- and vector-valued functions may be used to measure specific shape properties, such as the distance from a point or the Gaussian curvature. On the other hand, distances defined on the model representations provide insights on the corresponding shape distributions, such as geodesic and diffusion distances;
- Among the above methods, special emphasis is given to those techniques that quantify similarity in terms of *maps* between spaces. Many of these approaches fall in well-established mathematical frameworks, thus taking advantage of theoretical results on stability, robustness and invariance to shape transformations;
- Finally, we restrict our attention on methods published from 2008 on.

We will analyse the properties of the different approaches, possibly collocate them within the appropriate theoretical frameworks and discuss the extent of their applicability. In particular, we will analyse the properties of the methods with respect to the specific shape invariants they consider (e.g. rigid and non-rigid transformations), as well as the type of output they provide (e.g. full or partial similarity, sparse or dense correspondence),

1.2. Comparison to other surveys

If analysed through the lens of similarity assessment, shape retrieval, correspondence, alignment, symmetry detection, etc. are all different aspects of a more general problem: for instance, the knowledge of the punctual correspondences between two shapes might drive the definition of a distance in terms of distortion, and viceversa.

In general, shape retrieval methods target the evaluation of the distance of a query model from the objects in a database: most of existing surveys on 3D content based retrieval [VH01, TV04, BKS*05, FKMS05, YLZ07, DP06, BKSS07, TV08] mainly focus on the classification and discussion of methods, which target the conversion of statistical and geometric shape analysis into feature vectors or histograms. A common aspect of the methods reviewed in these papers is that, generally, they do not refer to any mathematical framework for similarity quantification, and there are no formal proofs of the stability/robustness of the distances adopted. Differently from these surveys, we focus on recent methods (since 2008) that follow a mathematical framework for the definition of the similarity distance.

More recent surveys have targeted specific aspects of similarity, such as shape registration [TCL*13], shape correspondence [vKZHCO11], symmetry detection [MPWC13], partial matching [LBZ*13], non-metric distances [SB11] and non-rigid shape retrieval [LGB*13]. For instance, Lian et al. [LGB*13] extend the analogous SHREC'2011 track [LGB*11]; using the same evaluation criteria the comparison is done in terms of performance of the retrieval over a common benchmark. In this review, we focus on approaches that explicitly use functions and/or distances to model the shape properties and invariants. We also consider deformations that are non-isometric; due to the closeness of the topics, there are some overlaps on the methods reviewed in previous works but our focus is different and mainly targeted on the mathematical aspects of similarity measurement.

The focus of recent surveys and tutorial works delivered by the authors in related fields, see for instance [BDF*08, BBK08, BFGS12, BCB12], is mainly on shape analysis and description, while the current review focuses on the problem of similarity quantification. This paper extends the EG tutorial [BFF*07] and illustrates three theoretical frameworks

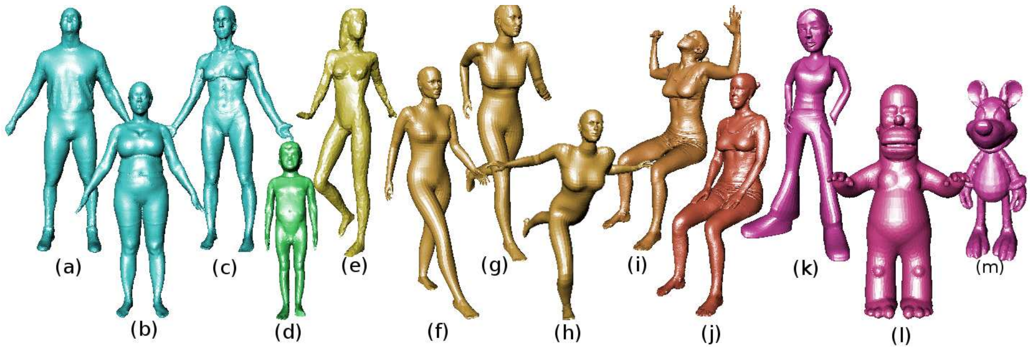

Figure 1: *Representative models of the humanoid class, SHREC'08 classification benchmark [GM08]. Models (a-d) have same pose but (d) differs by scale, (e) is a human model in a different pose and (f-h) are isometric deformations of the same template, (i-j) are two scans of the same model with a significant change of topology while (k-m) represent three virtual characters.*

that have been recently introduced in Computer Graphics [MS05, DF04, OBCS*12].

1.3. Organization

To drive the reader through the bunch of approaches and frameworks revised here, we firstly introduce the basic notions of mathematical concepts such as space, manifold, metric, shape transformation, diffusion geometry and algebraic topology, see Section 2. Depending on his/her background, the reader may skip this section or some of its parts.

Then, Section 3 is about the problem of similarity assessment and its mathematical modelling. We will discuss the properties of a number of similarity measures and introduce three formal frameworks for similarity assessment. Also, we describe some representative methods involved in the algorithmic evaluation of similarity.

Section 4 presents a taxonomy of the methods highlighting the emerging shape structure, the distances concretely used for similarity evaluation, and the invariance properties captured in the process. The proposed taxonomy also takes into account the type of input for each method, as well as the typology of output. The aim is to give a multi-faceted classification that might help the reader to compare methods not only on the basis on their algorithmic aspects, but also drive him/her in the choice of the method that better fulfils his/her requirements.

A more detailed analysis on the specific application domain for which methods have been proposed in then carried out in Section 5. Also, in Section 6 we review the recent development of shape benchmarks, which are becoming a crucial aspect for the quantitatively evaluate the performance of the methods.

Finally, Section 7 is devoted to the discussion of the potential of the methods proposed, also including perspectives, open issues, and future developments.

We believe that organizing the comparison of the various methods in this way may facilitate their analysis, possibly suggesting interesting research directions for the development of new approaches. In our opinion, the generality and flexibility of these methods may be of interest for a wide research community involved in visualization and topological modelling, beyond people working in shape matching and retrieval.

2. Mathematical Background

In this section we summarize the theoretical concepts which are necessary to model the shape similarity problem as presented in the rest of the paper.

2.1. Topological spaces and maps

A *topological space* is a set of points, along with a set of neighbourhoods for each point, that satisfy some axioms relating points and neighbourhoods. The definition of a topological space can be considered the most general notion of a mathematical space, and allows for the introduction of concepts such as continuity, connectedness, and closeness. Being so general, topological spaces are a central notion in mathematics. As for the digital 3D world, they are key ingredient to model the shape of the objects under study, as well as to reason about concepts like robustness and stability of shape analysis methods.

Topological spaces. A topological space (X,τ) is a set X on which a *topology* τ has been defined, that is, a collection of subsets of X called *open sets* and satisfying the following axioms:

- Both X and the empty set are open sets;
- Intersecting a finite number of open sets gives an open set;
- Any union of open sets is still an open set.

A *Hausdorff space* is a topological space in which distinct points admit disjoint neighbourhoods.

In what follows, we will refer to a topological space (X,τ) by simply mentioning the set X, omitting any reference to τ.

Maps. A *map* f between topological spaces is said to be *continuous* if the inverse image of every open set is an open set. A *homeomorphism* is a continuous bijection whose inverse is also continuous. Two topological spaces X, Y are said to be *homeomorphic* if there exists a homeomorphism $f: X \to Y$. From the viewpoint of topology, homeomorphic spaces are essentially identical. Properties of topological space which are preserved up to homeomorphisms are said to be *topological invariants*.

An important property of maps, which will be useful in the sequel, is *smoothness*. Roughly, a continuous map f is smooth if it has continuous partial derivative of all orders. Note, however, that this definition depends on the notion of partial derivative, which is usually well-defined only if the domain of f is an open set. Therefore, for an arbitrary subset $X \subseteq \mathbb{R}^n$ we need to adapt the above definition, stating that a continuous function $f: X \to \mathbb{R}^m$ is smooth if it can be locally extended to a smooth map on open sets; that is, around each point $x \in X$ we can find an open set $U \subseteq \mathbb{R}^n$ and a function $F: U \to \mathbb{R}^m$ such that F equals f on $X \cap U$, and whose partial derivative of all orders are continuous.

For $X \subseteq \mathbb{R}^n$, $Y \subseteq \mathbb{R}^m$, a smooth map $f: X \to Y$ is a *diffeomorphism* if it is bijective ant its inverse is smooth as well. In this case, X and Y are said to be *diffeomorphic*.

The notion of map plays a central role in shape analysis and comparison. On the one hand, they can by used to model spatial relations between two (or more) shapes represented by suitable topological spaces. On the other hand, real- or vector-valued maps can be adopted to encode measurements which are relevant to characterize the shapes under study. Throughout the paper, we will talk about *functions* rather than maps whenever referring to real- or vector-valued maps, in accordance with a quite common habit. Note, however, that the two concepts are completely equivalent from the mathematical viewpoint.

2.2. Metric spaces and transformations

Metric spaces can be seen as specifications of topological spaces. Their definition relies on the concept of *metric* (or *distance*), which describes a way to quantify the relative closeness between different entities, such as points, spaces or physical objects.

Metric spaces. A metric space (X,d) is a set X equipped with a metric, that is, a function $d: X \times X \to \mathbb{R}$ satisfying the following properties for all $x,y,z \in X$:

- $d(x,y) \geq 0$ (non-negativity);
- $d(x,y) = 0$ iff $x = y$ (reflexivity);
- $d(x,y) = d(y,x)$ (symmetry);
- $d(x,y) + d(y,z) \geq d(x,z)$ (triangle inequality).

Every metric space is a topological space in a natural way, by considering as open sets the open balls induced by d.

The Euclidean 3D space is an example of a metric space, where the metric is given by the well known *Euclidean distance*, that is, the distance between two points is the length of the straight line that joins them. The *geodesic distance* generalizes the concept of "straight line" to an arbitrary metric space (X,d): for two points in X, their geodesic distance is the length, measured with respect to d, of the shortest path between them, which is in turn referred to as *a geodetic*. More formally, a geodetic is a curve $\gamma: [a,b] \to X$ which is locally a distance minimizer: every $t \in [a,b]$ has a neighbourhood $J \subseteq [a,b]$ such that, for any $t_1, t_2 \in J$, the equality $d(\gamma(t_1), \gamma(t_2)) = |t_1 - t_2|$ holds.

Transformations. By the term *transformation*, we refer here to *structure-preserving* maps between spaces. As we will see later, relevant transformations from the viewpoint of shape similarity are isometries, affine transformations and homeomorphisms.

Isometries are distance-preserving maps, taking elements of a metric space to another metric space such that the distance between the elements in the new metric space is equal to the distance between the elements in the original metric space. Formally, given two metric spaces (X,d_X), (Y,d_Y), a transformation $\phi: X \to Y$ is called an isometry if for any $x,y \in X$, $d_Y(\phi(x),\phi(y)) = d_X(x,y)$. Examples of isometries in the usual Euclidean space are *rigid motions*, that is, combinations of translations and rotations; shape properties that are invariant to rigid motions are also called *extrinsic* because they are related on how the surface is laid out in the Euclidean space.

Affine transformations, or simply *affinities*, preserve straight lines (i.e., all points lying on a line initially still lie on a line after transformation) and ratios of distances between points lying on a straight line (e.g., the midpoint of a line segment remains the midpoint after transformation). They do not necessarily preserve angles or lengths, but do have the property that sets of parallel lines will remain parallel to each other after being affinely transformed. In particular, a map $\phi: X \to Y$ is an affine transformation if and only if for every family $\{(a_i, \lambda_i)\}_{i \in I}$ of weighted points $a_i \in X$ such that $\sum_{i \in I} \lambda_i = 1$, we have $f(\sum_{i \in I} \lambda_i a_i) = \sum_{i \in I} \lambda_i f(a_i)$. Examples of affine transformations include translation, geometric contraction, expansion, homothety, reflection, rotation, scale and compositions of them.

A larger class of transformations, also including isometries and affinities, is that of homeomorphisms, which preserve topological properties of spaces such as compactness, connectedness and Hausdorffness (the property of being Hausdorff). However, from the shape comparison point of view, topological invariance is not in general a reasonable requirement, admitting, e.g., that a horse surface model is topologically equivalent to a sphere and to a human surface model. This fact opened the way to the development of theoretical frameworks to improve the topological analysis of spaces by taking into account the additional information provided by real functions defined on the spaces themselves, such as Morse theory [Mil63] and other related frameworks [FM99, ELZ02].

2.3. Manifolds

To ease the analysis of a shape and look at it as if we locally were in "our" traditional Euclidean space, it is necessary to consider the notion of manifold. A Hausdorff space X is a *n-dimensional manifold* if it is locally homeomorphic to $\mathbb{R}^n$; that is, each point $x \in X$ admits a neighbourhood $V \subseteq X$ homeomorphic to an open set of $\mathbb{R}^n$. Such local homomorphism, is called *a coordinate system on* V, and allows for identifying any point $v \in V$ with a n-tuple of $\mathbb{R}^n$. X is a *n-dimensional manifold with boundary* if every point has a neighbourhood homeomorphic to an open set of either $\mathbb{R}^n$ or the half-space $H^n = \{u = (u_1, \ldots, u_n) \in \mathbb{R}^n | u_n \geq 0\}$. The boundary of X, namely ∂X, consists of those points of X which only have neighbourhoods locally homeomorphic to H^n. Note that, according the above definitions, any manifold is also a manifold with (possibly empty) boundary, while the converse does not hold in general.

A manifold X is *smooth* if it is equipped with a notion of differentiability. We prefer here to skip the technicalities needed to formally define such a notion, referring the reader to [Hir97] for further details. We rather point out that, having a notion of differentiability at a hand, we can do differential calculus on X and talk about concepts like tangent vector, vector field and inner product. All of these are functional to introduce *Riemannian manifolds*.

Riemannian manifold. If X is a smooth manifold of dimension n, at each point $x \in X$ we can consider the *tangent space* $T_x(X)$, a vector space that intuitively contains all possible vectors passing tangentially through x, see Figure 2 for an intuition. If we glue together all tangent spaces $T_x(X)$, thus considering $\bigcup_{x \in X} T_x(X)$, we get the *tangent bundle* $T(X)$. A *vector field* on X is then a section of $T(M)$, that is, a smooth map from $F : X \to T(M)$ which assign each point $x \in X$ to a tangent vector $F(x) = v \in T_x(X)$. On each tangent space $T_x(X)$ we can define an inner product (i.e. a symmetric, positive definite bilinear form) $g_x : T_x(X) \times T_x(X) \to \mathbb{R}$. A *Riemannian metric* g is a collection of inner products $\{g_x\}_{x \in X}$ that smoothly vary point by point, in the sense that if F and G are vector fields on X, then $x \mapsto g_x(F(x), G(x))$ is a smooth map.

Note that, in practice, a Riemannian metric is a symmetric tensor that is positive definite. Indeed, once a local system

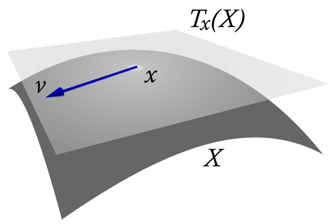

Figure 2: *Tangent plane* $T_x(X)$ *in* x*. The vector* $v \in T_x(X)$ *is a tangent vector.*

of coordinates is fixed for a point x, we can completely define each g_x by the inner products $g_{ij}(x) = g_x(v_i, v_j)$, with $\{v_1, v_2, \ldots, v_n\}$ a basis in $\mathbb{R}^n$. The collection $\{g_{ij}(x)\}$ is thus made of real symmetric and positive-definite $n \times n$ matrices, smoothly varying in x: It is called *a metric tensor* g_{ij}.

A *Riemannian manifold* is a n-dimensional differentiable manifold X equipped with a Riemannian metric g of metric tensor g_{ij}. Endowing a manifold with a Riemannian metric makes it possible to define various geometric notions on the manifold, such as angles, lengths of curves, curvature and gradients. The Riemannian metric on the surface does not depend on the particular embedding of the surface; properties that preserves this metric structure are called *intrinsic properties* of the surface.

2.4. Basics on diffusion geometry

In [CL06], Coifman and Lafon proposed the framework of diffusion geometry as a method for data parametrization, embedding, and dimensionality reduction. We summarize here some key ingredients of this framework, with particular reference to the *heat kernel signature* [SOG09], also known as the autodiffusion function [GBAL09], and the *diffusion distance*. Informally, diffusion geometry is related to the heat diffusion on the data (hence the name), which is turn closely connected with the notion of *Laplace operator*.

2.4.1. Laplace operator

The Laplace operator, briefly Laplacian, is a differential operator given by the divergence of the gradient of a real-valued function f defined on the Euclidean space $\mathbb{E}^n$:

$$\Delta f(p) := \mathrm{div}(\mathrm{grad}\, f(p)) = \nabla \cdot \nabla f(p) = \sum_i \frac{\partial^2 f}{\partial x_i^2}(p),$$

where grad and div are the gradient and divergence on the space and the point $p \in \mathbb{E}^n$ is represented by the Cartesian coordinates $p = (x_1, \ldots, x_n)$. Therefore, the Laplacian requires that the function f is at least twice-differentiable.

Intuitively, the Laplace operator generalizes the second order derivative to higher dimensions, and is a characteristic of the irregularity of a function, indeed $\Delta f(p)$ measures the difference between $f(p)$ and its average in a small neighbourhood of $p \in \mathbb{E}^n$.

The generalization of the Laplace operator to manifolds equipped with a Riemannian metric is called the *Laplace-Beltrami operator* of f and its computation requires complex calculations, that can be greatly simplified by the so-called *exterior calculus (EC)* [GDP*05].

The Laplace-Beltrami operator admits an eigendecomposition with non-negative eigenvalues λ_i and corresponding orthonormal eigenfunctions ϕ_i satisfying $\Delta\phi_i = -\lambda_i\phi_i$. Here orthonormality is meant in the sense of the inner product $\langle f, g \rangle = \int_X f \cdot g \, d\mu$, induced on a Riemannian manifold X by the associated Riemannian metric. Moreover, if we further assume that X is compact[†], we have that the spectrum is discrete, $0 \leq \lambda_1 \leq \lambda_2 \leq \dots$. In general, the eigenbasis of the Laplace-Beltrami operator is referred to as the harmonic basis of the manifold, and the functions ϕ_i as manifold harmonics [LV08, WZL*10]. The use of Laplacian eigenbasis has been shown to be fruitful in many computer graphics applications and several techniques in shape analysis, synthesis, and correspondence rely on the harmonic bases that allow for using classical tools from harmonic analysis on manifolds. For a detailed discussion on the main properties of the Laplace-Beltrami operator, we refer the reader to [Reu06, Ros97, WMKG07].

Several discrete Laplace-Beltrami operators exist [LV08], allowing for practical computation on a manifold discretization. For example, suppose to have a triangulation T with vertices $V := \{p_i, i = 1, \dots, n\}$. A function f on T is defined by linearly interpolating the values $f(p_i)$ of f at the vertices of T. This is done by choosing a base of piecewise-linear *hat-functions* φ_i, with value 1 at vertex p_i and 0 at all the other vertices. Then f is given as $f = \sum_{i=1}^{n} f(p_i)\varphi_i$. Discrete Laplace-Beltrami operators are usually represented as:

$$\Delta f(p_i) := \frac{1}{d_i} \sum_{j \in N(i)} w_{ij} \left[f(p_i) - f(p_j) \right],$$

where $N(i)$ denotes the index set of the 1-*ring* of the vertex p_i, i.e. the indices of all neighbors connected to p_i by an edge. The masses d_i are associated to p_i and the w_{ij} are the symmetric edge weights. If $V = \mathrm{diag}(v_1, \dots, v_n)$ is the diagonal matrix whose elements are $v_i = \sum_{j \in N(i)} w_{ij}$, $W = (w_{ij})$ and $D = \mathrm{diag}(d_1, \dots, d_n)$, then we can set $A := V - W$ and finally represent the discrete Laplacian-Beltrami operator on T as the $n \times n$ matrix given by $L := D^{-1}A$ (generally not symmetric).

Depending on the different choices of the edge weights and masses, discrete Laplacian operators are distinguished between *geometric operators* and *finite-element operators* [RBG*09]. A deep analysis of different discretizations of the Laplace-Beltrami operator in terms of the correctness of their eigenfunctions with respect to the continuous case is shown in [RBG*09]. Unless some special cases (see, for example, [BSW08, BS07, Sin06, HAvL05]), the discrete Laplace-Beltrami operator would not converge to the continuous one. In addition, when dealing with intrinsic shape properties, it should be independent or at least minimally dependent on the triangular mesh and thus the discrete approximation has to preserve the geometric properties of the Laplace-Beltrami operator. Unfortunately, Wardetzky et al. in [WMKG07] showed that for a general mesh, it is theoretically impossible to satisfy all properties of the Laplace-Beltrami operator at the same time, and thus the ideal discretization does not exist. This result also explains why there exists such a large diversity of discrete Laplacians, each having a subset of the properties that make it suitable for certain applications and unsuitable for others [BBK08].

2.4.2. Heat kernel and diffusion distance

Formally, the heat kernel signature and the diffusion distance can be expressed in terms of the *heat equation*. For a compact Riemannian manifold X, the diffusion process on X is described by the partial differential equation:

$$\left(\frac{\partial}{\partial t} + \Delta \right) u(t, x) = 0, \qquad (1)$$

where Δ denotes the Laplace-Beltrami operator associated with the Riemannian metric of X. The heat equation governs the distribution of heat u from a source point $x \in X$. The initial condition of the equation is some initial heat distribution $u(0, x)$ at time $t = 0$; if X has a boundary, appropriate boundary conditions must be added.

The *heat kernel* $h_t(x, y)$ is a fundamental solution of equation (1), with point heat source at x, and heat value at y after time t: it represents the amount of heat transferred from x to y in time t due to the diffusion process (Figure 3). By the eigendecomposition of Δ, the heat kernel can be written as

$$h_t(x, y) = \sum_{i \geq 0} \exp^{-\lambda_i t} \phi_i(x)\phi_i(y).$$

Since coefficients λ_i rapidly decay, the heat kernel is generally approximated by the truncated sum:

$$h_t(x, y) = \sum_{i=1}^{N} \exp^{-\lambda_i t} \psi_i(x)\psi_i(y).$$

The heat kernel has many nice properties, among which invariance to isometries; being related to the Riemannian metric of X, this means that the heat kernel is an intrinsic property of the manifold. Also, the heat kernel is multi-scale: for small vales of t, $h_t(x, \cdot)$ only reflects local properties of the manifold around the base point x, while for large values of t it captures the global structure of X from the point of view

† A compact manifold is a manifold that is compact as a topological space. A topological space X is compact if, from any union of open sets giving X, it is possible to extract a finite subfamily whose union is still X.

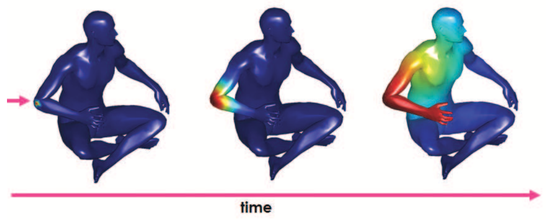

Figure 3: *The heat kernel represents the amount of heat transferred from a source point in time t.*

of x. Finally, the heat kernel is stable under small perturbations of the underlying manifold. All these properties make the heat kernel an optimal candidate for the definition of informative functions and distances to be used for shape description, such as the heat kernel signature (HKS) [SOG09] and the diffusion function. The HKS at a time t, denoted by HKS_t, is defined as

$$HKS_t(x) = h_t(x,x),$$

for any $x \in X$; the diffusion distance d_t between two points $x, y \in X$ at time t is given by

$$d_t^2(x,y) = h_t(x,x) + h_t(y,y) - 2h_t(x,y).$$

The computation of the spectrum of the discrete Laplacian is the main computational bottleneck for the evaluation of the heat kernel, and hence of HKS_t and d_t; in fact, it takes from $O(n)$ to $O(n^3)$ operations, according to the sparsity of the Laplacian matrix. Recently, a discrete and spectrum-free computation of the diffusion kernel on a 3D shape (either represented as a triangulation or a point cloud) has been proposed in [PS13], based on the computation of the full spectrum via the Chebyshev approximation [CMV69, ML03] of the weighted heat kernel matrix.

2.5. Basics on algebraic and differential topology

A fundamental issue in Shape Analysis is the study of basic models and methods for representing and generating. Since discretization strategies play a fundamental role in the way the results stated in a smooth context can be achieved in discrete ones, in this Section we briefly review some basic concepts that are at the bases of 3D shape representations [Req80, Man88].

2.5.1. Simplicial complexes

In order to construct topological spaces, one can take a collection of simple elements and glue them together in a structured way. Probably the most relevant example of this construction is given by simplicial complexes, whose building-blocks are called simplices. A detailed dissertation on simplicial complexes can be found in [Mun00].

A *k-simplex* Δ^k in $\mathbb{R}^n$, $0 \leq k \leq n$, is the convex hull of $k+1$ affinely independent points $A_0, A_1, \ldots, A_k$, called *vertices*. Figure 4 shows the simplest examples of simplices: Δ^0 is a point, Δ^1 an interval, Δ^2 a triangle (including its interior), Δ^3 a tetrahedron (including its interior), see Figure 4.

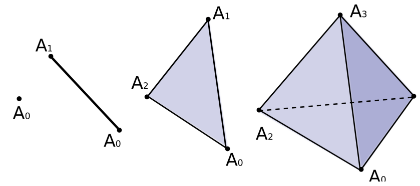

Figure 4: *Examples of 0-, 1-, 2- and 3-simplices.*

A *face* of a k-simplex Δ^k is a simplex whose set of vertices is a non-empty subset of the set of vertices of Δ^k. A *finite simplicial complex* is defined as a finite collection of simplices that meet only along a common face, together with their faces of any dimension. Triangulations are examples of simplicial complexes: the vertices, edges and faces correspond to 0-, 1- and 2-simplices, respectively. The *dimension* of a simplicial complex is the maximum dimension of its simplices.

2.5.2. Homology groups

The approach adopted by algebraic topology is the translation of topological problems into an algebraic language, in order to solve them more easily. A typical case is the construction of algebraic structures to describe topological properties, which is the core of homology theory, one of the main tools of algebraic topology.

The homology of a space is an algebraic object which reflects the topology of the space. The *homology* of a space X is denoted by $H_*(X)$, and is defined as a sequence of groups $\{H_q(X) : q = 0, 1, 2, \ldots\}$, where $H_q(X)$ is called the *q-th homology group* of X. The homology $H_*(X)$ is a topological invariant of X. The rank of $H_q(X)$, called the *q-th Betti number of* X and denoted by β_q, is roughly a measurement of the number of different holes in the space X. For three-dimensional data the Betti numbers β_0, β_1 and β_2 count the number of connected components, tunnels and voids, respectively.

In the literature there are various types of homologies [Spa66]. One of the most popular is (integer) *simplicial homology*, which relies on the concept of simplicial complex.

2.5.3. Basic concepts on Morse theory

Morse theory can be seen as the investigation of the relation between functions defined on a manifold and the shape of the manifold itself. The key feature in Morse theory is that information on the topology of the manifold is derived from the information about the critical points of real functions defined on the manifold. In particular, Morse theory

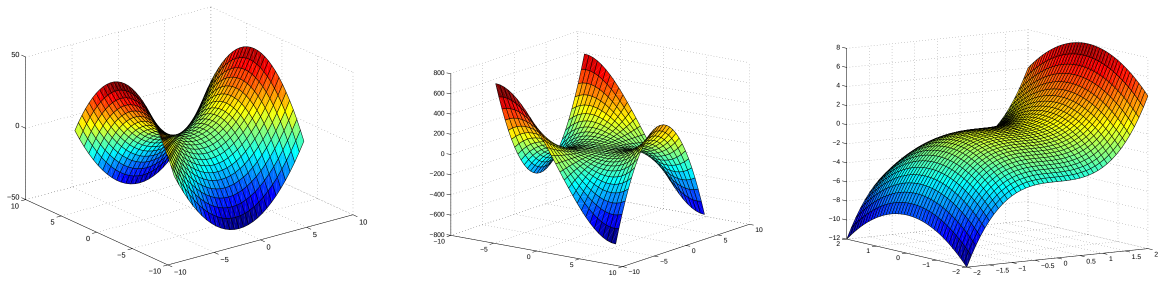

(a) (b) (c)

Figure 5: *(a) The graph of* $f(x,y) = x^2 - y^2$. *The point* $(0,0)$ *is a non-degenerate critical point. (b) and (c) The graphs of* $f(x,y) = x^3 - 3xy^2$ *(a "monkey saddle") and* $f(x,y) = x^3 - y^2$. *In both cases the point* $(0,0)$ *is a degenerate critical point.*

provides the mathematical background underlying several descriptors, such as Reeb graphs, size theory, persistent homology and Morse shape descriptors. A basic reference for Morse theory is [Mil63], while details about notions of geometry and topology can be found, for example, in [Hir97].

Let X be a smooth, compact n-dimensional manifold without boundary, and $f : X \to \mathbb{R}$ a smooth function defined on it. Then, a point x of X is a *critical point* of f if all the first order partial derivatives vanish at x, that is,

$$\frac{\partial f}{\partial x_1}(x) = 0, \ldots, \frac{\partial f}{\partial x_n}(x) = 0,$$

with respect to a local coordinate system $(x_1, \ldots, x_n)$ at x. A real number is a *critical value* of f if it is the image of a critical point. Points (values) which are not critical are said to be *regular*. A critical point p is *non-degenerate* if the determinant of the *Hessian* matrix of f at x,

$$H_f(p) = \left(\frac{\partial^2 f}{\partial x_i \partial x_j}(p) \right)$$

is not zero; otherwise the critical point is *degenerate*. Figure 5 shows some examples of non-degenerate and degenerate critical points. For a non-degenerate critical point p, the number of negative eigenvalues of the Hessian $H_f(p)$ of f at p is referred to as the *index* of p. Then, $f : M \to \mathbb{R}$ is a *Morse function* if all its critical points are non-degenerate.

An important property is that a Morse function defined on a compact manifold admits only finitely many critical points, each of which is isolated. This means that, for each critical point x, it is always possible to find a neighbourhood of x not containing other critical points. Moreover, Morse theory asserts that changes in the topology of a manifold endowed with a Morse function occur in the presence of critical points, and according to their index; these changes in the topology can be interpreted in terms of homology.

On the basis of these results, it is possible to choose regular values $t_0 < t_1 < \cdots < t_m$ bracketing the m critical values for f, and consider the *sublevel sets* $X_i = \{x \in X \,|\, f(x) \le t_i\}$. Moreover, if λ is the index of the i-th critical point, when sweeping from X_{i-1} to X_i there are two possibilities for how homology can change: either $\beta_\lambda(X_i) = \beta_\lambda(X_{i-1}) + 1$ or $\beta_{\lambda-1}(X_i) = \beta_{\lambda-1}(X_{i-1}) - 1$. The analogous approach to study the changes in the level sets $\{x \in X \,|\, f(x) = t\}$, $t \in \mathbb{R}$, is proposed in [Mil65].

3. Assessing similarity between spaces

Assessing the similarity between shapes can be posed as the problem of defining a suitable function $d : \mathcal{X} \times \mathcal{X} :\to \mathbb{R}$, taking a pair of input objects from a universe $\mathcal{X}$ to a real number that represents a similarity score for the two objects [SB11]. Such a function d is called a *pairwise similarity* function. Often the inverse concept is required, namely a *dissimilarity* function δ, where a higher dissimilarity score stands for a lower similarity score, and vice versa. Hence, a dissimilarity δ equivalent to a similarity d must fulfill $d(X,Y) \ge d(X,Z) \iff \delta(X,Y) \le \delta(X,Z), \forall\, X,Y,Z \in \mathcal{X}$.

The choice between similarity and dissimilarity function mainly depends on the application domain; however there exist many situations where the formula/algorithm defining the function is available in just one of the two forms, while its manual transformation into the inverse is not straightforward [SB11]. The application scenario is also strongly related to the properties that the chosen (dis)similarity functions is required to satisfy, such as being a metric or not. Being a metric means to fulfill all the postulates listed in Section 2.2. Assuming, e.g, that a dissimilarity function δ has been fixed, *reflexivity* permits zero dissimilarity just for identical objects while *non-negativity* guarantees that every two distinct objects are somehow positively dissimilar. In addition, the triangle inequality is a kind of transitivity property that is really useful for indexing a database [ZADB06]: if (X,Y) and (X,Z) are close with respect to δ (that is, small dissimilarity), also (X,Z) are.

A number of (dis)similarity functions exist in the literature, which do not fulfil one or more of the metric axioms. Such functions are generally referred to as *non-metrics*

[SB11], presenting more specific names according to the particular metric axiom they miss. In case reflexivity is not guaranteed, then we have a *pseudometric*; a *quasi-metric* if symmetry is not satisfied, a *semi-metric* if triangle inequality is missing. The paradigm here is that, being not constrained by metric postulates, non-metrics offers a larger freedom of problem modelling. Indeed, several psychological theories suggest the metric axioms could substantially limit the expressive power of (dis)similarity functions [SJ99, Tve77]. In particular, reflexivity and non-negativity have been refuted by claiming that different objects could be differently self-similar [Kru78, Tve77]. The triangle inequality is the most attacked property. Some theories point out that similarity does not have to be transitive [AP88, TG82], as shown by a well-known example: a man is similar to a centaur, the centaur is similar to a horse, but the man is completely dissimilar to the horse.

Beyond (a subset of) metric axioms, a notion of continuity is often required for a (dis)similarity function, such as robustness with respect to different discretizations of spaces and small perturbations in the input measurements. Last but not least, invariance to some classes (groups) of transformations may be required, thus allowing the similarity assessment to be independent, e.g, to orientation, scaling or rigid movements of the considered objects. Formally, a similarity function d (a dissimilarity function δ, respectively) is invariant under a chosen group of transformations G if for all transformations $g \in G$ and all $X,Y \in \mathcal{X}$, we have $d(g(X),Y)) = d(X,Y)$ (resp. $\delta(g(X),Y)) = \delta(X,Y)$).

In the last decade, new emphasis has been given to assess the dissimilarity between two shapes by modelling them as suitable spaces and to formally quantify similarity in terms of the distortion needed to deform one space into the other. Besides the classical approaches for similarity assessment, in what follows we summarize three different theoretical approaches that have been recently and successfully introduced in Computer Graphics. For a detailed review of other existing distances, and related application fields, we refer to the *Encyclopedia of Distances* [DD09].

3.1. A standard approach

The use of concise descriptions, which are usually referred to as *shape descriptors*, instead of the whole model representation is a common strategy in shape comparison. Therefore, the first challenge is to identify the shape properties that better characterize the object under study and are highly discriminative; in our settings, this translates in the selection of the functions used to detect the main shape features [BDF*08]. A good shape descriptor should be robust and endowed with adequate invariance properties. Indeed, robustness guarantees that small changes in the input data, such as noise or non-relevant details, do not result in substantial changes in the associated shape descriptors. Invariance properties are related to the application domain; for instance, in case shape alignment rotations and translations.

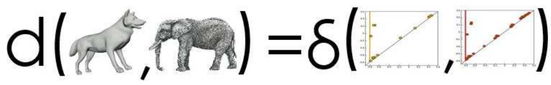

Figure 6: *The dissimilarity d between two objects is computed as the distance δ between their descriptors.*

Having a good shape descriptor at hand, the problem of assessing the similarity between two shapes can be recast into the comparison of the associated descriptors, according to a suitable (dis)similarity measure taking into account the remarks discussed above in this section, see Figure 6.

The use of shape descriptors is largely acknowledged in the literature and a variety of methods has been proposed [BKS*05, TV04, DP06, TV08, BDF*08, vKZHCO11, WZL*10, TCL*13]. During years, the situation has evolved from 3D descriptors heuristically introduced [BKS*05], motivated by techniques and practices inherited from vision (projection-based descriptions), geometry (statistics of surface curvature or geodesic distances), or signal processing (object samples in the frequency domains), to more sophisticated and mathematically sound frameworks leading to detect salient shape's feature yet showing robustness to noise and different group of transformations.

3.1.1. Examples

Among the variety of methods proposed in the literature, we have selected some representative ones that extract shape information in the form of functions, and use that information to derive shape descriptors.

The method proposed by Mademlis et al [MDTS09] adopts the potential of a Newtonian field defined in the space outside the shape. The 3D descriptor is the combination of independent histograms (36 in the paper) related to surface proximity, field intensity and curvature. Histograms that come from the Newtonian field are compared with the normalized distance $d(H_1,H_2) = \sum_{i=1}^{k} \frac{2(H_1(i)-H_2(i))^2}{H_1(i)+H_2(i)}$ where k is the number of bins of the histograms H_1,H_2, while the diffusion distance is used to compare the curvature-based histograms. The robustness of the method to small shape variations derives from the preprocessing step (voxels simplify small shape details) and from the stability of the volumetric function. In addition to rotation and translation invariance, scale-invariance is achieved through a pre-processing step in which all shapes are normalized and voxelized.

The approach proposed by Smeets et al. in [SHVS12] achieves intrinsic invariance through the computation of the geodesic distances between surface samples. These values are stored in a matrix (GDM) that is unique up to vertex permutations. A first shape description consists of two histograms: the first one stores all upper triangular values of the GDM, the second one contains an histogram of the average of the geodesic distances. A second descriptor is provided by the singular value decomposition (SVD) of the GDM. Several possible feature vector distances are proposed (e.g. χ^2 and L^p) but the authors mainly refer to the Jensen-Shannon divergence. The method is shown to be robust under near isometric deformations (articulated object and faces) but the use of geodesics make it sensitive to topological changes (e.g., open/close mouth, two fingers that touch themselves, etc.).

The scale invariant heat kernel signature (SI-HKS) was proposed by Bronstein et al. [BK10c] to overcome the scale dependence of the Heat kernel signature. Once the HKS is computed, the scale dependence is avoided locally normalizing the heat kernel. Roughly speaking, the local normalization is done in terms of scaling and shift in time: scale is obtained from the logarithm of h and its discrete derivative with respect to time, while the shift is seen as a different phase that it is discarded through a complex representation of the discrete Fourier transform. The SI-HKS at each point of the shape is approximated through soft quantization by the closest geometric words in a precomputed vocabulary of 48 elements. The resulting shape description is a 48-dimensional bag of features that is compared using the L^1 distance. The SI-HKS fully satisfies intrinsic invariance and scale independence. Moreover, the choice of the discretization scheme for the Laplace-Beltrami operator (e.g. point wise or mesh-based) makes this signature available for many different inputs (e.g., point clouds or meshes).

Besides the use of histograms, graph-based signatures are well suited when structure and shape parts are relevant for the application. As a representative example we describe the technique for finding corresponding parts in structurally different objects proposed by Shapira et al. [SSS*10]. First, a hierarchical shape partition is obtained using the shape diameter function (SDF) [SSCO08]. Second, each part is described by a local signature made of: i) a normalized histogram of the SDF values and the size of that part as a percentage of the whole model, ii) a set of shape distribution signatures from [OFCD02], iii) the curvature-based histogram proposed in [BCG08], also known as the conformal geometry signature. Then, the comparison is done with a bipartite graph matching approach that is based on a flow algorithm able to take into account both local geometrical features and the part hierarchies. The robustness of the method with respect to small deformations depends on the robustness of the partition technique. Being based on a hierarchical representation, this method tackles the problem of comparing parts from very different shapes, even with different topology; however, the use of the SDF makes it particularly suitable for articulated shapes.

Finally, we summarize the method proposed in [ZBVH09, ZBH12] for comparing and matching textured 3D shapes. Such a method builds upon a scale-space derived from different normalized Gaussian derivatives through the Difference-of-Gaussians (DoG) operator [Low04], and incorporates in a unique paradigm geometry and photometric information. The operator is computed on a scalar function defined on the manifold, which in the original paper is either the mean curvature, the Gaussian curvature or the photometric appearance of a vertex (the mean of the RGB channels). The computation of the scale-space does not alter the surface geometry (differently from the similar approach in [CCFM08]). A local descriptor, called MeshHOG, is obtained as a two-level histogram of the projections of the gradient vectors onto the three planes associated with the local coordinate systems of the maxima and minima of the scale space representation. The final descriptor is obtained by concatenating the histogram values for each of the 3 orthonormal planes. Last, in order to have invariance to the mesh sampling (i.e. the selection of the feature points), the concatenated histograms are normalized through the L^2 norm, that is also used to compare two meshHOGs. Depending on the choice of the function (mean curvature, Gaussian, etc.) isometry invariance is satisfied, while the scale-space description guarantees robustness to noise [ZBH12].

3.2. Gromov-Hausdorff distance

The *Hausdorff distance* is probably one of the most immediate ways to assess the dissimilarity for two subsets of a metric space. Informally, it is the maximum distance of a set to the nearest point in the other set [HKR93]. Therefore, two sets are close in the Hausdorff distance if every point of one set is close to some point of the other set.

The Hausdorff metric $d_H(X,Y)$ between two non-empty subsets X and Y of a metric space (Z,d) is defined as:

$$d_H(X,Y) = \max\left\{ \sup_{x\in X}\inf_{y\in Y} d(x,y), \sup_{y\in Y}\inf_{x\in X} d(x,y) \right\},$$

where sup represents the supremum and inf the infimum. An extension of this concept is provided by the L_p-Hausdorff distance [Bad92] of which the Hausdorff metric represents the particular case $p=\infty$. Due to its relatively easy evaluation, Hausdorff distance and modifications are often used to measure similarity between rigid surfaces in $\mathbb{R}^3$, which give rise to a family of algorithms known in the literature as ICP [Zha94]. Also, the Hausdorff distance is very popular for shape comparison, ranging from images and digital terrain surfaces to 3D objects [HKR93, NSCE02]. However, while being a good match to compare extrinsic geometries, the Hausdorff distance is not invariant to isometries.

In order to face this problem, an isometry-invariant extension of the Hausdorff distance was introduced by Gromov [Gro99]. The *Gromov-Hausdorff distance* casts the comparison (and therefore the quantification of the similarity) of two shapes as a problem of comparing pairwise distances on metric spaces used to model the shapes themselves. Equivalently, the computation of the Gromov-Hausdorff distance

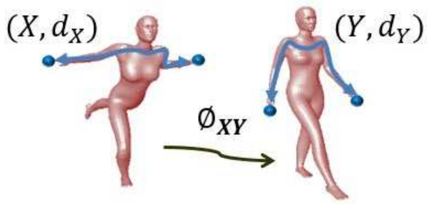

Figure 7: *Iconic representation of an isometric map ϕ_{XY} between two metric spaces (X,d_X) and (Y,d_Y).*

can be posed as the evaluation of how much the metric structure is preserved while mapping a space into the other.

The idea is to represent the comparison of two shapes as that between two metric spaces (X,d_X) and (Y,d_Y). For a map $\phi_{XY} : X \to Y$, we measure the *distortion* induced by ϕ_{XY} on the metric d_X as

$$\mathrm{dis}(\phi_{XY}) = \sup_{x,y \in X} |d_X(x,y) - d_Y(\phi_{XY}(x), \phi_{XY}(y))|. \quad (2)$$

Obviously, if $\mathrm{dis}(\phi_{XY}) = 0$ there is no distortion for d_X, and in fact we have that ϕ_{XY} is an isometry, see Figure 7. As for mappings $\phi_{XY} : Y \to X$, we can define $\mathrm{dis}(\phi_{YX})$ in the same way as in Eq. (2), by exchanging the roles of X and Y. Additionally, we consider the *joint distortion* $\mathrm{dis}(\phi_{XY}, \phi_{YX})$ given by

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

which roughly measures how far ϕ_{XY} and ϕ_{YX} are from being one the inverse of the other. The Gromov-Hausdorff d_{GH} distance between X and Y is then defined as:

$$d_{GH}(X,Y) = \inf_{\phi_{XY}, \phi_{YX}} \max\{\mathrm{dis}(\phi_{XY}), \mathrm{dis}(\phi_{YX}), \mathrm{dis}(\phi_{XY}, \phi_{YX})\}.$$

The combination of the metric approach with the Gromov-Hausdorff framework does not require any particular metric to be defined on spaces. Indeed, by choosing different metrics between points, we get different notions of distances between spaces [Gro99, M1́2, M1́1]. However, two possible choices appear quite natural here. The first one is to set d as the *geodesic metric*, thus defining the intrinsic geometry of X: In this case, d measures the length of the shortest path on X between two of its points. The second choice for d is the *Euclidean metric*, which relates to the extrinsic geometry of X: For two points in X, their distance is measured as the length of their connecting segment.

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. As a result, the Euclidean metric is not suitable for the comparison of nonrigid shapes. On the other hand, intrinsic geometry is invariant to inelastic shape deformations, which indeed are metric preserving. Therefore, the geodesic metric is a good choice for comparing non-rigid shapes, as has been confirmed by several contributions. However, other invariance classes can be relevant in applications, for example topological deformations or scaling. To this aim, more sophisticated choices are possible, such as the diffusion or the commute-time distance [WBBP12].

The Gromov-Hausdorff distance was first proposed for deformable shape analysis by Mémoli and Sapiro [MS05], together with an approximation scheme for discrte spaces. Given two samplings of (X,d_X) and (Y,d_Y) with the same number N of points, say (X_N,d_X) and (Y_N,d_Y) respectively, by restricting the attention to bijective mappings between X_N and Y_N, it is possible to approximate the Gromov-Hausdorff distance d_{GH} by a permutation distance:

$$d_{\widetilde{GH}}(X_N, Y_N) = \min_{\sigma \in \mathcal{P}(N)} \max_{0 \le i,j \le N} |d_X(x_i, x_j) - d_Y(x_{\sigma_i}, x_{\sigma_j})|,$$

with σ varying in the set $\mathcal{P}(N)$ of all permutation of N numbers, and $\sigma_i = \sigma(x_i)$. As shown in [MS05], $d_{\widetilde{GH}}$ approximates d_{GH} for randomized samplings.

3.3. Functional maps

The Gromov-Hausdorff framework allows for studying shape similarity through the comparison of pairwise distances defined on suitable spaces representing shapes. In practice, this problem is often approached by considering correspondence between points on the two shapes. For example, isometric matching techniques try to find correspondences between landmark points that preserve geodesic distances as well as possible. However, unless a small number of landmarks is considered, moving in this direction is in general computationally intractable, since the space of possible point correspondences is exponential in size.

Motivated by this problem, Ovsjanikov et al. [OBCS*12] proposed a novel representation of maps between shapes, based on looking for correspondences between real-valued functions defined on shapes, rather then between points on shapes. Formally, suppose that $T : M \to N$ is a bijective mapping between manifolds M and N representing two shapes, and $f : M \to \mathbb{R}$ is a real-valued function. The function on N corresponding to f is given by the relation $g = f \circ T^{-1}$. The induced transformation $T_F : \mathcal{F}(M,\mathbb{R}) \to \mathcal{F}(N,\mathbb{R})$, with $\mathcal{F}(M,\mathbb{R})$ and $\mathcal{F}(N,\mathbb{R})$ two generic spaces of the real-valued functions defined on M and N, respectively, is said to be the *functional map representation* of T.

Functional maps generalize the standard point-to-point map since every pointwise correspondence induces a mapping between function spaces, while the opposite is, in general, not true. Also, the knowledge of T can be recovered from the one of T_F by replacing f with suitable indicator functions [OBCS*12]. However, while T can be in principle very complicated, T_F is a linear map between function spaces. This has the main advantage that, after fixing a basis for the function space on each shape, a functional map represents the corresponding mapping as a change of basis matrix. More precisely, if $\mathcal{F}(M,\mathbb{R})$ is assumed to be equipped with a basis $\{\phi_i^M\}$, then any function $f \in \mathcal{F}(M,\mathbb{R})$ can be expressed as $\sum_i a_i \phi_i^M$, with a_i real coefficients. Moreover, if also $\mathcal{F}(N,\mathbb{R})$ is associated with a basis, say $\{\phi_j^N\}$, then T_F can be completely determined by these bases:

$$T_F(f) = T_F\left(\sum_i a_i \phi_i^M\right) = \sum_i a_i T_F\left(\phi_i^M\right) = \sum_j \sum_i a_i c_{ij} \phi_j^N,$$

with $T_F(\phi_i^M) = \sum_j c_{ij} \phi_N^j$ and c_{ij} is the generic element of the change of basis matrix. Remarkably, such a matrix is particularly simple to represent when the basis functions are orthonormal with respect to some inner product, as in the case of the eigenfunctions of the Laplace-Beltrami operator (see Section 2.4.1). Furthermore, being based on a linear algebraic formulation, the framework of functional maps opens the way to the practical usage of many common linear algebra tools, ranging from matrix multiplications to study map composition, matrix inversion to move from a bijection $T : M \to N$ to $T^{-1} : N \to M$, to principal component analysis (PCA) and singular value decomposition.

Figure 8 shows an example of a bijective map between two nearly isometric dog shapes, and the corresponding functional representation in the form of a matrix $C_{20\times 20}$. To get C, the first 20 Laplace Beltrami eigenfunctions were used as a the function basis for the function space in each shape. As shown by the picture, functional representations of nearly isometric maps are close to being sparse and diagonally dominant. In [ROA*13], functional maps are used to explore shape differences within a dataset of shapes. The main idea is to formalize the shape difference under a map by mimicking the Riemannian approach. Indeed, in that case local map distortions are expressed in terms of variations in the Riemannian metric. Inspired by this, the authors define the *conformal* and the *area preserving inner product* on the functions spaces of the considered shapes, and track the changes of tangent vectors before and after these vectors are mapped from one shape to the other. Such changes can be encoded in the form of linear operators (matrices) which are referred to as *conformal* and *area preserving shape differences*, respectively. With such a notion in a hand, a collection of shapes is then represented as a collection of shape differences from a base shape. In particular, matrices representing shape differences are vectorized and, after PCA has been applied, the associated coefficients are plotted along principal directions. In this new formulation, the database collection can be analyzed, e.g, to find the "average" shape, explore variability in shapes and defining shapes analogies.

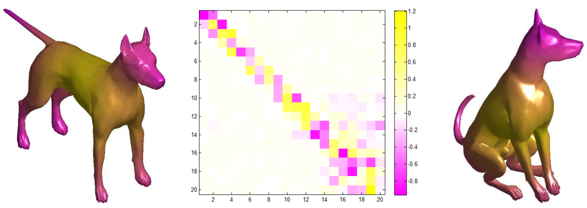

Figure 8: *Two shapes and the isometric mapping between them in the form its functional representation.*

Other possible applications of functional maps include isometric shape matching in the presence of symmetries [OMPG13] and attribute transfer [OBCS*12]. Also, in [SNB*12] *soft maps* are presented, providing a probabilistic relaxation of point-to-point shape correspondence. Similarly to functional maps, soft maps can be represented as probabilistic matrices, thus allowing for the usage of linear algebra tools in their analysis and manipulation.

3.4. The natural pseudo-distance

If we push further the idea of measuring the distortion of properties while transforming a shape into another, i.e. considering topological spaces instead of metric spaces, we get the concept behind the *natural pseudo-distance* [DF04, DF07, FM99]. The starting point is to model a shape as a pair (X, f), where X is a topological space equipped with a continuous real-valued function $f : X \to \mathbb{R}$ encoding a shape property of interest. To compare two pairs (X, f) and (Y, f), with X and Y homeomorphic, we can imagine to transform one space into the other trough a homomorphism $h : X \to Y$, and check how much the properties of the original shape have been preserved/distorted by h; this problem amounts to measure the difference between the functions f and $g \circ h$. In other words, shapes are considered similar if there exists a homeomorphism that preserves the properties conveyed by the functions.

Note that to represent a given shape it is possible to choose the topological space that best fits with the problem at hand. For example, we might want to fix $X = S$, with S a 2-dimensional manifold modelling the shape surface, but also the Cartesian product $S \times S$ in case the function f to be studied is a metric defined on S. Other possible choices could be the tangent space of S, or a projection of S onto a line, or the boundary of S, or the skeleton of S, and so on. Such a choice is driven by the set of properties that one wishes to capture.

More formally, the *natural pseudo-distance* between two pairs (X, f) and (Y, g) is defined by setting

$$d_{np}((X, f)\ ,\ (Y, f)) = \inf_{h\in H} \sup_{x\in X} |f(x) - g \circ h(x)|,$$

with h varying in the set H of homoeomorphisms from X to Y. If X and Y are not homeomorphic the pseudo-distance is set equal to ∞. Note however that the existence of a homeomorphism is not required for the shapes under study, but rather for the associated spaces X and Y. In this way, two objects are considered as having the same shape if and only if they share the same shape properties, i.e. the natural pseudo-distance between the associated size pairs vanishes.

The natural pseudo-distance offers a framework in which different shape properties can be plugged-in in the form of different real functions, so as to measure shape (dis)similarity up to different notions of invariance. Such a modular setting fostered the development of a topology-based approach to shape description and comparison based on the use of different classes of functions, describing both extrinsic and intrinsic properties of shapes. Some of them have been singled out as better suited than others to deal with specific problems, such as obtaining invariance under groups of transformations [DFP04, DLL*10], or working with particular classes of objects [CFG06, FS10, CCSG*09]. Nevertheless, the choice of the most appropriate functions for a particular application is not fixed a priori and, as observed for the Gromov-Hausdorff framework, has to be carefully carried out up to the specific application/problem at hand.

3.4.1. Examples

The computational issues related to the practical evaluation of the natural pseudo-distance are still an algorithmic bottleneck: for this reason many efforts focused on the definition of computationally efficient approximations. In addition, the existence of a proper notion of distance between spaces has led to the definition of descriptors that are stable under shape's perturbations.

These research issues led to the introduction of *size functions*, shape descriptors that are proven to be stable under the natural pseudo-distance, and provide a lower bound for it [dFL10]. They were afterwards included in the framework of *Topological Persistence* (hereafter simply *persistence*) [ELZ02, EH10], whose family of theoretical and computational tools, with particular reference to *persistence diagrams* [CSEH07], can be used to derive lower bounds for the natural pseudo-distance [CSEH07, CFF*13] and the Gromov-Hausdorff distance [CCSG*09]; in the latter case, topological spaces are replaced by metric spaces. All these signatures are able to naturally combine the classifying power of topology with the descriptive power of geometry.

Having modeled a given shape as a pair (X, f), with $f : X \to \mathbb{R}$, according to persistence we can consider the sublevel sets of f to define a collection of subspaces $X_u = \{x \in X | f(x) \leq u\}$, $u \in \mathbb{R}$, nested by inclusion, i.e. *a filtration of* X. Homology may then be applied to derive some topological information about the filtration of X. More precisely, the idea is to track how topological features vary in passing from a set of the filtration into a larger one, in much the same way as suggested by Morse theory (see Section 2.5.3). From the homological viewpoint, this can be done in terms of the evolution of the Betti numbers along the filtration, which gives insights, e.g., on the *birth* and the *death* of connected components, tunnels or voids.

The topological evolution of the sublevel sets of f is finally encoded in a persistence diagram $\mathrm{dgm}(f)$. This is a collection of points in the half-plane $\{(u, v) \in \mathbb{R}^2 : u < v\}$. For each point, the u-coordinate represents the birth, in terms of the values of the function f, of a topological feature, whereas the v-coordinate represents its death. A persistence diagram provides a multi-scale description of the shape under study. Indeed, points far from the diagonal $u = v$ represent long-lived features, while points close to the diagonal – they are characterized by a shorter life – stand for noise and details. The paradigm is that long-lived features are more meaningful or coarse for shape description, while short-lived ones stand for noise and details. Examples of persistence diagrams, describing the evolution of connected components along different filtrations, are shown in Figure 9. The (red) vertical line in the four diagrams can be seen as a point at infinity, and represents a topological feature that *will never die*.

Thanks to their modularity, persistence diagrams provides different shape descriptions simply by changing the considered function. Interestingly, they inherit the invariance properties directly form the considered functions. Perhaps more importantly, they are stable under the Hausdorff and *bottleneck distance*, which in turn provide lower bounds for the natural pseudo-distance. In particular, small changes in the input function f produces only small changes in the associated persistence diagram $\mathrm{dgm}(f)$ [CSEH07, dFL10].

Most of the persistence applications fall in the field of shape matching and retrieval: persistence diagrams play the role of shape descriptors, while similarity is derived from a stable distance between them. For example, diameter function, eccentricity function and higher-order eccentricity functions are used in [CCSG*09] to build persistence diagrams on Rips filtrations of finite metric spaces, so to derive stable signatures providing a lower bound for the Gromov-Haussdorf distance, while [BGSF08b] uses size functions, which are roughly the persistence diagrams studying the evolution of connected components, to compare attributed skeletal graphs derived from functions that code extrinsic and intrinsic shape properties.

Similarly to persistence, also *Reeb graphs* [Ree46] root in Morse theory, but track the evolution of the level-set of a function f instead of its sub-level sets [BGSF08a, DW11]. From the mathematical point of view the Reeb graphs can be defined as the quotient space induced by the equivalence relation that identifies the points belonging to the same connected component of level sets of f [Ree46]. Reeb graphs have been introduced in Computer Graphics in the 90's by Shinagawa et al. [SK91, SKK91] while their use for shape matching dates back to 2001 [HSKK01] with the definition of the Multiresolution Reeb graph (MRG). Since then, several variations of the Reeb graph have been introduced to couple the topological information stored in the graph with geometric attributes of the shape parts corresponding to

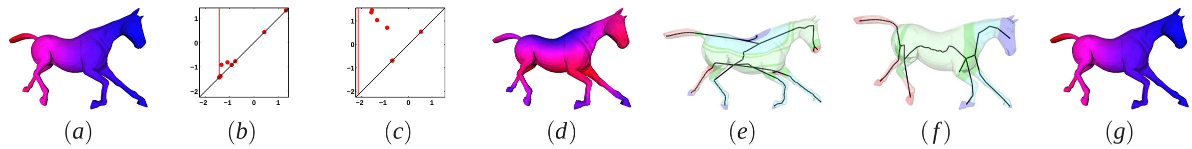

Figure 9: *Persistence diagrams $(b-c)$ and Reeb graphs $(e-f)$ related to different choices of the function f (color coded, increasing values from blue to red).*

nodes and arcs, the most popular being the augmented Multiresolution Reeb graph [TS04, TS05], the Extended Reeb graph [BMM*03] and the Discrete Reeb graph [XSW03]).

Also, several graph matching methods have been introduced, ranging from global similarity measures [HSKK01, LMM13] to approximated sub-graph matching techniques [BMSF06] and graph kernel approaches [BB13a]. Recently, a stability result for Reeb graphs under the *functional distortion distance* has been proposed [BGW14], leading to a lower bound for the natural pseudo-distance.

The parametric nature of Reeb graphs with respect to the function f is shown in the last two rows of Figure 9, where the Reeb graphs of a closed surface with respect to two different functions are depicted. Notice how different functions can give insights on the shape from a different perspective.

Recent advances in the field of topological methods for shape analysis include the extension of the persistence-based descriptions to the use of vector-valued functions [BCF*08, CZ09, CSZ10, CL13, BCGS13], the introduction of statistics in combination with topology [MMH11], and the possible use for shape analysis [GH10], the comparison of set of shapes [PBF07], the use of collections of graphs [BB13a, BEGB13] and learning techniques [BB13b] to automatically infer the most promising shape properties and select the salient features.

4. Taxonomy of the methods

In what follows, we propose a "practical" classifications of the surveyed methods, according to key characteristics that are important in applications:

- *Input Type:* what kind of shape representations is used;
- *Structure type:* what kind of shape structure is captured from the chosen description method;
- *Distance:* the criterion used to assess the (dis)similarity between shape structures;
- *Invariance:* the class(es) of transformations to which the method is invariant;
- *Output type:* the result of the method: either a similarity value, or a shape correspondence, or both.

Obviously, these criteria are inter-related: for example, the input type may put limitations on the kind of the structures that can be computed, and, in turn, the choice of the structure would usually determine the invariance (e.g. if one uses diffusion geometric structures to find correspondence between shapes, such a correspondence would be invariant to isometric deformations).

We discuss different methods from these viewpoints in the following and summarize our taxonomy in Table 1. Have in mind that, in line with the "practical" perspective we have chosen for the proposed taxonomy, and as a convention in this survey, we will stick to the information related to the specific application setting described in the respective paper, though in some cases generalizations might be straightforward.[‡]

4.1. Type of input

The input type is related both to the application from which shapes come and the mathematical model of the shape similarity or correspondence [Req80, BDF*08]. In the computer graphics community, shapes are traditionally modeled as *surface* models (two-dimensional manifolds representing boundaries of physical 3D objects) or as *volume* models [Man88, Mor86]. The most common discretizations of such structures are simplicial meshes (e.g. triangular or tetrahedral meshes) [BDF*08, BBK09, RBBK10b, Rus10, BLC*11, DP13, LB14] or 2D/3D regular grids [LTN06, BCF*08, MDTS09]. 3D grids (voxel) representations are mainly used in medical applications.

On the other hand, in the computer vision community it is common to see *point cloud* representations for 3D data obtained in shape-from-X problems. The recent emergence of 3D acquisition hardware has made these representations popular in rigid matching problems [TV08, TCL*13], which play an essential role in multi-view fusion. In the analysis of deformable shapes, such representations are less common [MHK*08, MS09, NBPF11].

In many situations, additional information can be available in additional to the geometric structure of the shape. A typical example we report here is *texture* [ZBVH09, KBBK11, ZBH12, KBB*12, GL12, KRB*13, BCGS13].

‡ For example, methods based on diffusion geometry are easily applicable to different types of input data, but if the paper shows results on triangular meshes only, we say that the method is applicable to meshes.

4.2. Type of structure

A useful classification of the geometric structures may be done according to their invariance. Structures invariant to rigid transformations are referred to as *extrinsic* (geometry). Simple rotation- and translation invariant-structures include Euclidean distances (e.g. from the object barycenter, lines, planes [BK10a, BCF*08] or shape boundaries [MDTS09, LH13]). Extrinsic structures can be extended to cope with global scale or affine transformations.

Structures invariant to transformations preserving the local metric of the manifold are referred to as *intrinsic* (geometry); this type of invariance is sought in applications involving deformable shapes. In the proposed classification we distinguish between *conformal* structures by referring to those based on Gaussian curvature and geodesic distances; *diffusion* structures for those based on diffusion distances and spectral properties of the Laplace-Beltrami operator [RWP06, Rus07, WZL*10]; *autodiffusion* structures for those built on various types of local spectral descriptors such as the heat and wave kernel signatures [SOG09, GBAL09, ASC11]. Intrinsic structures can be further made invariant to global scale and affine transformations [BK10b, RBB*11].

If one wishes invariance to topological changes, the description must be able to deal with *topology*. Examples of structures able to code topological information are curve skeletons [LH13], Reeb graphs [HSKK01, BGSF08a] and persistence-based representations [EH07, CZ09, DLL*10, BCF*08, BCFG11, DL12].

Structures can be also classified as *local* or *global*. Local structures reflect the properties of the shape in the vicinity of a point of interest and are usually unaffected by the geometry or the topology outside that neighborhood. For that reason, local structures are typically used for partial shape matching. Global structures, on the other hand, capture the properties of the entire shape.

Local structure is usually captured in the form of local descriptors. Recent works proposed a plethora of descriptors such those based as mean and Gaussian curvatures, integral volume descriptors [GMGP05], conformal factor [BCG08], autodiffusion [GBAL09] or the heat kernel signature [SOG09] and its scale- [BK10b] and affine-invariant versions [RBB*11], and the wave kernel signature [ASC11].

Global structure can be obtained by integrating local structure over the entire shape, typically in the form of a single- or multi-dimensional histogram [OBGB09, BB13a]. This is a standard approach in retrieval applications where a holistic description of the entire shape is required. Other inherently global structures include distance functions (e.g., Euclidean, geodesic, or diffusion distances) and their distributions [HSKK01], as well as global spectral properties such as the Laplace-Beltrami spectrum [RWP06] and eigenfunctions.

Some methods combine both local and global properties of the shape producing *semi-local* structures such as the maximally stable extremal regions (MSERs) [LBB11]. These structures arise in the form of hierarchies of stable regions, and are guided both by the behaviour of a local descriptor (e.g., heat kernel signature) and the coarse-scale properties of the shape. Similarly, bilateral maps [vKZH13] provide a medium scale description that depends on the closeness of the base points: the shape portions to be matched must cover at least the 20% or 30% of the overall shape.

Other methods allows for dealing with shape information at different scales, thus providing a unifying interpretation of local and global shape description. We will refer to the related structures as *multi-scale* structures.

In shapes endowed with additional *photometric* information (texture), more complicated approaches such as multidimensional size functions [BCF*08, BCFG11] and combined geometric/photometric descriptors [ZBVH09, ZBH12, GL12, KRB*13, BCGS13] can be used.

4.3. Type of distance

The criteria used to compare shapes depend on both the information enclosed in the emerging structures, and how this information is coded (e.g. histograms, graphs, point-correspondences, etc.). As discussed in Section 3, several frameworks have been introduced for comparing shape structures.

The *Gromov-Hausdorff* distance provides a general framework for comparing shapes in terms of the metric distortion between two metric spaces when transforming one into the other. The framework is flexible to the choice of the metric, and therefore can be used to compare different metric structures, such as diffusion geometry [BBK*10].

Functional maps extend the similarity problem to the comparison of functions defined on the shapes. The main advantage of this problem formulation is that a correspondence is a linear map in the space of functions: hence, a number of tools and techniques from linear algebra can be applied to couple shapes. Soft maps relax the usual point-to-point correspondence through a probabilistic formulation. Since soft maps can be represented in the form of matrices, also in this case their manipulation may be accomplished via the standard linear-algebraic toolkit.

Methods related to the *natural-pseudo distance* may take advantage of a mathematically sound notion of stability. Also, the intrinsic modularity of the framework allows for comparing shapes according to different invariance properties, which are directly inherited from the functions used to describe the considered shape properties.

Beyond the above frameworks, a number of methods considered here represent different interpretations of a standard approach for similarity assessment, based on the computation of suitable distances between shape descriptors.

For instance, a simple, yet effective way of globally describing a shape is to use a *bag of features*. Such a description is the projection in a k-dimensional vector of the feature detected. Feature vector distances are a well-known issue in shape retrieval [BKS*05, TV08]. Traditionally, solutions to this item are provided by the *Minkowski* L^p family of distances. Examples include the Manhattan distance ($p = 1$); the usual Euclidean distance ($p = 2$); the maximum distance ($p = \infty$), also called Chebyshev or chessboard metric. Other distances provided by statistics and information theory are χ^2-statistics, the *Hamming distance*, the *Jeffrey divergence*, the *Jensen-Shannon (J-S) divergence*, the *Wasserstein* distance also known as the *Earth Mover's distance (EMD)* in the discrete settings [LO07].

In case the structure is coded in a graph, many distances have been introduced, each one depending on the type of information stored in the graph and its hierarchical nature. Examples are the *approximation of the maximum common subgraph* [BMSF06, TVD09, BK10a, AK11], *path matching* [SSS*10, MBH12, RPSS10, LH13], *Hungarian* distance [SPT11, STP12, GDZ10] and *graph kernels* [BB13a, BB13b, LMS13].

Many other distances may be listed, which in some cases have been proposed as *ad-hoc* similarity measures between shape descriptors, see [DD09] for more details.

4.4. Invariance

The type of invariance is strictly interlinked to the structure the method captures. For instance, extrinsic geometric structures are invariant to *rigid* transformations (rotations, translations, and reflections) [MDTS09, GDZ10, BK10a, DP13].

Intrinsic structures (such as geodesic and diffusion distances, heat kernels signatures, etc.) are invariant to *isometric* shape deformations, which are in general non-rigid but preserve the Riemannian structure of the shape [BBK09, SSS*10, WZL*10, LD11, BLC*11, BB11, BBK*10, RBB*11, BBGO11, LPD13] or its volume [RBBK10b, Rus10].

Other classes of deformations which are relevant for applications include, e.g, scaling and affine transformations [RBB*11].

Besides invariance, methods may exhibit robustness with respect to small perturbations and variations, such as model sampling [HH09], shot noise [BCF*08, BBB*12, BBC*10, BBB*10] and shortcuts [Rus10, BBC*10, BBB*10, SHCB11, vKZH13]. Topological changes (possibly resulting, e.g., from acquisition artifacts or meshing) can greatly alter global intrinsic structures such as geodesic distances [TVD09, DLM09, APP*10, BCFG11, DL12, GL12, BCGS13]. Multi-scale intrinsic structures such as heat kernels are typically less sensitive to changes in topology, and so are extrinsic geometric structures [BB11, BBK*10, LGB*13, BBC*10, BBB*10].

4.5. Type of output

The process of comparing two shapes may result in either a numerical assessment of their *similarity*, or a *correspondence* between the two shapes, or both.

Algorithms aiming at the computation of correspondence usually produce a quantitative measure of similarity as a byproduct (e.g., metric distortion in [BBK*10]). On the other hand, numerous shape retrieval approaches based on holistic descriptors produce similarity only. The measure of similarity itself can be either *full* or *partial*.

Methods computing correspondence can be further subdivided into those finding full or partial correspondence. Both cases can be classified as *sparse* (correspondence is computed only between a small subset of feature points detected on the shapes being matched), or *dense* (typically represented by providing for each vertex or each triangle on one shape its image under the correspondence on the other). Dense correspondences can be alternatively represented by a smooth approximation of the continuous map in some basis using the functional correspondence formalism. Such correspondences are usually referred to as *soft*. Methods computing fuzzy correspondence abandon the representation of the latter as a function, allowing a single point on one shape to be mapped to a distribution on the other [KLM*12]. Despite the superficial similarity to the functional representation, the underlying details differ substantially.

5. Applications

Application domains of shape similarity and correspondence problems range from mechanical CAD to entertainment, forensic security, and molecular biology. In the literature, it is common to refer to the following (partially overlapping and sometimes used synonymously) classes of applications [FKMS05].

Shape matching is often considered as a synonym of geometric similarity and correspondence. Among the methods that evaluate similarity, we distinguish between those that generically refer to the generic field of matching [RBB*11, ASC11, LBB11, ZBVH09] or aim at quantifying similarity [BLC*11, LD11, LPD13, RBBK10a, BB11, BBK*10, BCF*08] and those that have more specific targets. Often, similarity can be inferred from correspondence methods based on structure distortion minimization [KLF11, LD11, KBB*13, PBB*13, ROA*13, LB14] or for a sufficiently large set of shape correspondences [vKZHCO11].

An extension of the matching problem to the setting where one seeks similarity [OMMG10, SSS*10, DLL*10, AK11, TVD09, DLM09, WZL*10, DL12, BWDP13a, vKZH13] or correspondence [MGP06, GCO06, FS06, DL11, TDVC11, SHCB11] between parts of shapes is referred to as *partial* matching (see, for instance [LBZ*13] for a recent survey). However, the peculiarity of similarity measures for partial matching is that generally they do not satisfy the triangle inequality because parts may match well even if the whole objects significantly differ (consider, for example, the human, centaur and horse shapes [TV04]).

A particular setting of the shape matching problem is matching the shape to itself for the detection of self-similarity or symmetry [RBBK10a, LH13]; symmetry detection, in turn, can be intrinsic or extrinsic, partial or full.

Registration refers to the alignment of the components of two or more shapes [TCL*13]. The problem originated from the need of rigidly aligning point clouds acquired by multi-view 3D scanners. More recent works considered finding differences between shapes [DP13] and non-rigid registration of deformable shapes [LZSCO09].

Shape retrieval refers to the problem of finding the models in a database that best match a given query. Therefore, all method whose output is a similarity score between couples of shapes, can be adopted for 3D content based retrieval, see Table 1.

Shape classification aims at finding the class the query model belongs to [LJ07, CCSG*09, Rus10, Bia10]. A good criterion of similarity between shapes allows performing shape classification if one has examples of shapes belonging to the given classes [BGM*06]; however, the trend is to use machine learning approaches for this task [BB13b, BEGB13].

Recognition is a particular case of retrieval; given a query and a database the problem is to determine if the query is in that dataset or not, and in case the answer is affirmative, to identify the query. A popular application is face recognition [tHV10, BDP10, SHVS12, BWDP13a] for security purposes [GAP*09, BDP13b]. Since facial deformations are almost isometric and some landmarks may be easily identified (for instance, the tip of the nose), methods for face recognition take advantage of the use of intrinsic structures such as geodesic distances from feature points [SHVS12, tHV10, BBK09, RBB*11] or curvature [BK10a]. The improvement in terms of performance of 3D face recognition has led to application of 3D methods to the identification of facial expressions [MAD*11, BDP13a], also in presence of partial occlusions [BDP13a].

Finally, we can identify additional applications that are indirectly related to shape similarity and correspondence, such as *exploration* of shape collections [HZG*12, ROA*13, LGPC13], *component detection* [LBB11, LZSCO09], *segmentation* [LZSCO09], simultaneous *shape editing* [KBB*13], *attribute transfer* [KBB*13], 3D animation *tracking* [SHCB11] and *semantic annotation* [ARSF09, LMS13].

6. Datasets and benchmarks

The explosive growth of the number of shape matching techniques has made acute the need for a widely-accepted performance evaluation protocol. Among the firsts of such benchmarks is SHREC [VRS*06] that started as a shape retrieval contest and grew over the years into additional tracks covering correspondence and feature detection tasks. Another popular benchmark is the Princeton Shape benchmark [SMKF04], which includes retrieval and, recently, correspondence and segmentation tasks. However, existing benchmarks are far from capturing the challenges arising in real applications. Such challenges include, for example, truly large scales involving millions of shapes, noisy, cluttered, and partially occluded data coming from real range sensors, the great variety of shape representations, and the need for domain-specific evaluation criteria. While the recent editions of the SHREC [BBC*10, BBB*10, BBB*12, CBA*13] have tried to address some of these issues by proposing a multi-track structure of the benchmark, a comprehensive benchmark relevant to real-world problems is still acutely needed.

7. Discussions and Future Trends

In this review we have presented methods for 3D similarity, emphasizing the properties of shape descriptors and related distances, yet discussing their theoretical aspects and possible applications. Among the whole literature on the topic, we have restricted our attention to approaches that abstract shape properties through the use of functions and/or distances, and quantify shape similarity in terms of maps between spaces. Most of these techniques fall in well-known mathematical theories (e.g. diffusion geometry, Morse theory, differential topology) and frameworks (Gromov-Hausdorff, functional maps and natural pseudo-distance); thus they borrow from the theory results on invariance, robustness and stability. A first contribution of the survey is the summary proposed in Section 3 of the most recent trends for similarity assessments.

The ensemble of the approaches highlights that defining a suitable similarity distance becomes more complicated as the shape descriptions increase their complexity and the variety of the deformations evolve from rigid to intrinsic or generic deformations. In these cases, the classic metric paradigm becomes less effective and non metric distances come into the play.

On the basis of the taxonomy proposed in Section 4 and Table 1, it is now possible to answer to questions that are crucial when selecting a technique for shape similarity assessment: e.g. is that method mainly targeted for global or partial matching? Which kind of input does it need? Is it suitable for non-rigid shape comparison? Is it possible to find sparse or dense shape correspondences?

The survey highlights that in the last years (since 2008), rigid-invariant comparison is considered a quite well established problem; the few methods tackling this problem are mainly devoted to the improvement of the accuracy and the computational efficiency of existing methods. At the same time, also point clouds have decreased their popularity because strictly related to rigid registration and alignment problems. On the contrary, defining a meaningful measure of the similarity under non-rigid deformations is still an acute issue.

On the one hand, research in the last five years has mainly focused on isometric deformations: the Gromov-Hausdorff and the functional maps frameworks have been successfully applied and are becoming the standard *de facto* for intrinsic similarity. On the other hand, the problem is still open when dealing with generic deformations that includes non-isometric changes or topological variations. In addition, there is the feeling that dealing with some classes of shapes and deformations is inherently more difficult than other ones; for instance, it has been shown there is a link between the stability of any shape matching technique and shape symmetries [OHG11].

The natural pseudo-distance has been introduced to overcome the limitation of the traditional methods which rely on assumptions of rigidity, isometry, or sufficient geometric similarity between corresponding parts; however, the direct evaluation of the natural pseudo-distance is still a computational bottleneck, although some methods may be used to get computable lower bounds.

7.1. Research trends and challenges

Finally, we list a series of topics deserving further research and efforts:

- The increasing complexity of the deformations considered in the applications calls for the development of new *formal theories*. Indeed, there is no a theoretically well-established framework when comparing two shapes with different structure or topology but same functionality. In our opinion, it is necessary to develop a theory for 'geometric filtration' similar to the Fourier analysis so that is possible to have a progressive representation of the shape properties under continuous functions.
- Moreover, it is necessary to specifically investigate the *role of invariance* with respect to larger families of transformation groups (i.e., shape deformations) and how to balance the use of geometrical and topological information for accurate shape descriptions.
- Beside the current theoretical limitations, there are also computational bottlenecks. In the framework proposed in the survey, the role of maps is crucial to convey the shape properties that one wishes take into account; however, while results and algorithms are well-established for single valued functions, it not always the same for the use of multi-variate functions. This is due to the fact that a complete understanding of the overall structure of the description space is still missing, resulting in the lack of efficient algorithms, such as in the case of persistence for dimension higher than two [BCFG11, BCGS13].
- To effectively quantifying similarity it is also necessary to consider the context in which the object is embedded (i.e. its semantics or functionality). A major issue is how to model such a knowledge and effectively embed it in the similarity evaluation pipeline. From our perspective, a possible solution is either to include a-priori knowledge or to co-analyse shapes in the same class or database.
- To automatically infer knowledge, we foresee the introduction in the framework of statistical methods such as learning techniques. Indeed, experiments for image classification using multicolumn neural networks reached near-human performance [CMS12]. These techniques represent a possible solution to automatically determine the weights of the different shape features on the basis of context (e.g., the shape classes of a database) [Lag10, LMS13, BSF13, TDVC13, BB13b].
- Another critical point to be tackled in the near future is *scalability*: indeed we expect that methods for similarity quantification will handle very large volumes of data, in the form of parallel approaches or on-line application. For example, a difficulty we foresee in this step is that methods that are able to capture the shape topology and deal with generic shape deformations are often time consuming, refer to quite complex distances and depend on the overall structure of the shape.

Acknowledgements

We would like acknowledge the support of the Shape Modelling Group at CNR-IMATI, and in particular Michela Spagnuolo and Bianca Falcidieno; we also thank Daniela Giorgi, Patrizio Frosini, Claudia Landi for the fruitful discussions and insights on the similarity problem. SB and AC are partially supported by the CNR research activity ICT.P10.009; the Italian CNR Flagship project INTEROMICS: InterOmics PB05, research unit WP 15; the EU projects VISIONAIR (EU FP7-INFRASTRUCTURES-262044); and IQmulus (EU FP7-ICT-2011-318787). AB and MB are supported by the ERC Starting Grants Nos. 335491 and 307047.

References

[AK11] AREEVIJIT W., KANONGCHAIYOS P.: Reeb graph based partial shape retrieval for non-rigid 3D objects. In *VRCAI '11* (New York, NY, USA, 2011), ACM, pp. 573–576. 16, 25

[AP88] ASHBY F. G., PERRIN N. A.: Toward a unified theory of similarity and recognition. *Psychological Review 95* (1988), 124–150. 1, 9

[APP*09] AGATHOS A., PRATIKAKIS I., PAPADAKIS P., PERANTONIS S., AZARIADIS P., SAPIDIS N. S.: Retrieval of 3D Articulated Objects Using a Graph-based Representation. In *3DOR 2009* (2009), Spagnuolo M., Pratikakis I., Veltkamp R., Theoharis T., (Eds.), pp. 29–36. 25

[APP*10] AGATHOS A., PRATIKAKIS I., PAPADAKIS P., PERANTONIS S., AZARIADIS P., SAPIDIS N.: 3D articulated object retrieval using a graph-based representation. *The Visual Computer 26*, 10 (2010), 1301–1319. 16, 25

[ARSF09] ATTENE M., ROBBIANO F., SPAGNUOLO M., FALCIDIENO B.: Characterization of 3D shape parts for semantic annotation. *CAD 41*, 10 (Oct. 2009), 756–763. 17

[ASC11] AUBRY M., SCHLICKEWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *ICCV Workshops 2011* (Nov 2011), pp. 1626–1633. 15, 16, 25

[Ash92] ASHBY F. G.: *Multidimensional models of perception and cognition*. Lawrence Erlbaum Associates, NJ, 1992. 1

[Bad92] BADDELEY A. J.: Errors in binary images and an Lp version of the Hausdorff metric. *Nieuw Arch. Wiskunde 10* (1992), 157–183. 10

[BB11] BRONSTEIN M., BRONSTEIN A.: Shape recognition with spectral distances. *IEEE Trans. Pattern Anal. Mach. Intell. 33*, 5 (2011), 1065–1071. 16, 25

[BB13a] BARRA V., BIASOTTI S.: 3D shape retrieval using Kernels on Extended Reeb Graphs. *Pattern Recognition 46*, 11 (2013), 2985 – 2999. 14, 15, 16, 25

[BB13b] BARRA V., BIASOTTI S.: Learning kernels on extended Reeb graphs for 3D shape classification and retrieval. In *3DOR 2013* (2013), pp. 25–32. 14, 16, 17, 18, 25

[BBB*10] BRONSTEIN A. M., BRONSTEIN M., BUSTOS B., CASTELLANI U., CRISANI M., FALCIDIENO B., GUIBAS L. J., SIPIRAN I., KOKKINOS I., MURINO V., OVSJANIKOV M., PATANÉ G., SPAGNUOLO M., SUN J.: SHREC 2010: Robust feature detection and description benchmark. In *3DOR 2010* (2010), pp. 87–91. 16, 17

[BBB*12] BIASOTTI S., BAI X., BUSTOS B., CERRI A., GIORGI D., LI L., MORTARA M., SIPIRAN I., ZHANG S., SPAGNUOLO M.: SHREC'12 Track: Stability on abstract shapes. In *3DOR 2012* (Cagliari, Italy, 2012), Spagnuolo M., Bronstein M., Bronstein A., Ferreira A., (Eds.), pp. 101–107. 16, 17

[BBC*10] BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U., FALCIDIENO B., FUSIELLO A., GODIL A., GUIBAS L. J., KOKKINOS I., LIAN Z., OVSJANIKOV M., PATANÉ G., SPAGNUOLO M., TOLDO R.: SHREC 2010: Robust large-scale shape retrieval benchmark. In *3DOR 2010* (2010), pp. 71–78. 16, 17

[BBGO11] BRONSTEIN A. M., BRONSTEIN M. M., GUIBAS L. J., OVSJANIKOV M.: Shape Google: Geometric words and expressions for invariant shape retrieval. *ACM Trans. Graph. 30*, 1 (Feb. 2011), 1:1–1:20. 16, 25

[BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*, 1 ed. Springer Publishing Company, Incorporated, 2008. 2, 6

[BBK09] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Topology-invariant similarity of nonrigid shapes. *Int. J. of Comput. Vision 81*, 3 (Mar. 2009), 281–301. 14, 16, 17, 25

[BBK*10] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *Int. J. of Comput. Vision 89*, 2-3 (2010), 266–286. 2, 15, 16, 25

[BCB12] BRONSTEIN M., CASTELLANI U., BRONSTEIN A.: Diffusion geometry in shape analysis. In *EG 2012 Tutorial* (2012). 2

[BCF*08] BIASOTTI S., CERRI A., FROSINI P., GIORGI D., LANDI C.: Multidimensional size functions for shape comparison. *J. Math. Imaging Vision 32* (2008), 161–179. 2, 14, 15, 16, 25

[BCFG11] BIASOTTI S., CERRI A., FROSINI P., GIORGI D.: A new algorithm for computing the 2-dimensional matching distance between size functions. *Pattern Recognition Letters 32*, 14 (2011), 1735–1746. 15, 16, 18, 25

[BCG08] BEN-CHEN M., GOTSMAN C.: Characterizing shape using conformal factors. In *3DOR 2008* (2008), Eurographics Association, pp. 1–8. 10, 15

[BCGS13] BIASOTTI S., CERRI A., GIORGI D., SPAGNUOLO M.: PHOG: Photometric and Geometric Functions for Textured Shape Retrieval. *Computer Graphics Forum 32*, 5 (2013), 13–22. 14, 15, 16, 18, 25

[BDF*08] BIASOTTI S., DE FLORIANI L., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., PAPALEO L., SPAGNUOLO M.: Describing shapes by geometrical-topological properties of real functions. *ACM Computing Surveys 40*, 4 (2008), 1–87. 2, 9, 14

[BDP10] BERRETTI S., DEL BIMBO A., PALA P.: 3D face recognition using isogeodesic stripes. *IEEE Trans. Pattern Anal. Mach. Intell. 32*, 12 (2010), 2162–2177. 17, 25

[BDP13a] BERRETTI S., DEL BIMBO A., PALA P.: Automatic facial expression recognition in real-time from dynamic sequences of 3D face scans. *The Visual Computer 29*, 12 (2013), 1333–1350. 17, 25

[BDP13b] BERRETTI S., DEL BIMBO A., PALA P.: Sparse matching of salient facial curves for recognition of 3-d faces with missing parts. *IEEE Trans. on Information Forensics and Security 8*, 2 (2013), 374–389. 17, 25

[BEGB13] BONEV B., ESCOLANO F., GIORGI D., BIASOTTI S.: Information-theoretic selection of high-dimensional spectral features for structural recognition. *Comput. Vis. Image Und. 117*, 3 (2013), 214–228. 14, 17, 25

[BFF*07] BIASOTTI S., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., PATANÈ G., SPAGNUOLO M.: 3D shape description and matching based on properties of real functions. In *EG 2007 Tutorial* (2007), pp. 949–998. 2

[BFGS12] BIASOTTI S., FALCIDIENO B., GIORGI D., SPAGNUOLO M.: The Hitchhiker's guide to the galaxy of mathematical tools for shape analysis. In *ACM SIGGRAPH 2012 Courses* (New York, NY, USA, 2012), SIGGRAPH '12, ACM, pp. 17:1–17:33. 2

[BGM*06] BIASOTTI S., GIORGI D., MARINI S., SPAGNUOLO M., FALCIDIENO B.: A Comparison Framework for 3D Object Classification Methods. In *MRCS* (2006), Günsel B., Jain A. K., Tekalp A. M., Sankur B., (Eds.), vol. 4105 of *Lecture Notes in Computer Science*, Springer, pp. 314–321. 17

[BGSF08a] BIASOTTI S., GIORGI D., SPAGNUOLO M., FALCIDIENO B.: Reeb graphs for shape analysis and applications. *Theor. Comput. Sci. 392*, 1-3 (2008), 5–22. 13, 15

[BGSF08b] BIASOTTI S., GIORGI D., SPAGNUOLO M., FALCIDIENO B.: Size functions for comparing 3D models. *Pattern Recognition 41*, 9 (2008), 2855–2873. 13

[BGW14] BAUER U., GE X., WANG Y.: Measuring distance between Reeb graphs. In *SoCG 2014* (June 2014). 14

[Bia10] BIASOTTI S.: Shape comparison through mutual distances of real functions. In *ACM Work. on 3D object retrieval* (New York, NY, USA, 2010), ACM, pp. 33–38. 17, 25

[BK10a] BALOCH S., KRIM H.: Object recognition through topo-geometric shape models using error-tolerant subgraph isomorphisms. *IEEE Trans. on Image Processing 19*, 5 (2010), 1191–1200. 15, 16, 17, 25

[BK10b] BRONSTEIN M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *CVPR 2010* (June 2010), pp. 1704–1711. 15

[BK10c] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *CVPR* (2010), IEEE, pp. 1704–1711. 10, 25

[BKS*05] BUSTOS B., KEIM D. A., SAUPE D., SCHRECK T., VRANIĆ D. V.: Feature-based similarity search in 3D object databases. *ACM Computing Surveys 37*, 4 (December 2005), 345–387. 2, 9, 16

[BKSS07] BUSTOS B., KEIM D. A., SAUPE D., SCHRECK T.: Content-based 3D object retrieval. *IEEE Computer Graphics and Applications 27*, 4 (2007), 22–27. 2

[BLC*11] BOYER D. M., LIPMAN Y., CLAIR E. S., PUENTE J., PATEL B. A., FUNKHOUSER T., JERNVALL J., , DAUBECHIES I.: Algorithms to automatically quantify the geometric similarity of anatomical surfaces. *PNAS* (Oct. 2011). 14, 16, 25

[BMM*03] BIASOTTI S., MARINI S., MORTARA M., PATANÉ G., SPAGNUOLO M., FALCIDIENO B.: 3D shape matching through topological structures. *Lecture Notes in Computer Science 2886* (2003), 194–203. 14

[BMSF06] BIASOTTI S., MARINI M., SPAGNUOLO M., FALCIDIENO B.: Sub-part correspondence by structural descriptors of 3D shapes. *Computer Aided Design 38*, 9 (September 2006), 1002–1019. 14, 16

[BS07] BOBENKO A. I., SPRINGBORN B. A.: A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete Comput. Geom. 38*, 4 (Dec. 2007), 740–756. 6

[BSF13] BIASOTTI S., SPAGNUOLO M., FALCIDIENO B.: Grouping real functions defined on 3D surfaces. *Computers & Graphics 37*, 6 (2013), 608 – 619. 18

[BSW08] BELKIN M., SUN J., WANG Y.: Discrete Laplace operator for meshed surfaces. In *Symposium on Computational Geometry* (2008). 6

[BWDP13a] BERRETTI S., WERGHI N., DEL BIMBO A., PALA P.: Geometric histograms of 3D keypoints for face identification with missing parts. In *3DOR 2013* (2013), pp. 57–64. 16, 17

[BWDP13b] BERRETTI S., WERGHI N., DEL BIMBO A., PALA P.: Matching 3D face scans using interest points and local histogram descriptors. *Computers & Graphics 37*, 5 (2013), 509–525. 25

[CBA*13] CERRI A., BIASOTTI S., ABDELRAHMAN M., ANGULO J., BERGER K., CHEVALLIER L., EL-MELEGY M. T., FARAG A. A., LEFEBVRE F., GIACHETTI A., GUERMOUD H., LIU Y.-J., VELASCO-FORERO S., VIGOUROUX J.-R., XU C.-X., ZHANG J.-B.: SHREC'13 Track: Retrieval on textured 3D models. In *3DOR 2013* (2013), pp. 73–80. 17

[CCFM08] CASTELLANI U., CRISTANI M., FANTONI S., MURINO V.: Sparse points matching by combining 3D mesh saliency with statistical descriptors. *Computer Graphics Forum 27*, 2 (2008), 643–652. 10, 25

[CCSG*09] CHAZAL F., COHEN-STEINER D., GUIBAS L. J., MÉMOLI F., OUDOT S.: Gromov-Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum 28*, 5 (2009), 1393–1403. 13, 17, 25

[CFF*13] CERRI A., FABIO B. D., FERRI M., FROSINI P., LANDI C.: Betti numbers in multidimensional persistent homology are stable functions. *Mathematical Methods in the Applied Sciences 36*, 12 (2013), 1543–1557. 2, 13

[CFG06] CERRI A., FERRI M., GIORGI D.: Retrieval of trademark images by means of size functions. *Graph. Models 68*, 5 (2006), 451–471. 13

[CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *Applied and Computational Harmonic Analysis 21*, 1 (2006), 5 – 30. 5

[CL13] CERRI A., LANDI C.: The persistence space in multidimensional persistent homology. In *Discrete Geometry for Computer Imagery*, Gonzalez-Diaz R., Jimenez M.-J., Medrano B., (Eds.), vol. 7749 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2013, pp. 180–191. 14

[CMS12] CIRESAN D., MEIER U., SCHMIDHUBER J.: Multi-column deep neural networks for image classification. In *CVPR 2012* (2012), pp. 3642–3649. 18

[CMV69] CODY W., MEINARDUS G., VARGA R.: Chebyshev rational approximations to exp(-z) in $[0, +\infty)$ and applications to heat-conduction problems. *Journal of Approximation Theory 2*, 1 (1969), 50 – 65. 7

[CSEH07] COHEN-STEINER D., EDELSBRUNNER H., HARER J.: Stability of persistence diagrams. *Discr. Comput. Geom. 37*, 1 (2007), 103–120. 13

[CSZ10] CARLSSON G. E., SINGH G., ZOMORODIAN A.: Computing multidimensional persistence. *JoCG 1*, 1 (2010), 72–100. 14

[CZ09] CARLSSON G., ZOMORODIAN A.: The theory of multidimensional persistence. *Discr. Comput. Geom. 42*, 1 (2009), 71–93. 2, 14, 15

[DD09] DEZA M. M., DEZA E.: *Encyclopedia of Distances*. Springer Berlin Heidelberg, 2009. 9, 16

[DF04] DONATINI P., FROSINI P.: Natural pseudodistances between closed manifolds. *Forum Mathematicum 16*, 5 (2004), 695–715. 2, 3, 12

[DF07] DONATINI P., FROSINI P.: Natural pseudodistances between closed surfaces. *J. of the European Math. Society 9*, 2 (2007), 231–253. 2, 12

[dFL10] D'AMICO M., FROSINI P., LANDI C.: Natural pseudo-distance and optimal matching between reduced size functions. *Acta Applicandae Mathematicae 109*, 2 (2010), 527–554. 13

[DFP04] DIBOS F., FROSINI P., PASQUIGNON D.: The use of size functions for comparison of shapes through differential invariants. *J. of Math. Imaging and Vision 21* (2004), 107–118. 13

[DK10] DUBROVINA A., KIMMEL R.: Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *3DPVT* (2010). 25

[DL11] DI FABIO B., LANDI C.: A Mayer-Vietoris Formula for Persistent Homology with an Application to Shape Recognition in the Presence of Occlusions. *Foundations of Computational Mathematics 11*, 5 (2011), 499–527. 16

[DL12] DI FABIO B., LANDI C.: Persistent homology and partial similarity of shapes. *Pattern Recognition Letters 33*, 11 (2012), 1445 – 1450. Computational Topology in Image Context. 15, 16, 25

[DLL*10] DEY T., LI K., LUO C., RANJAN P., SAFA I., WANG Y.: Persistent heat signature for pose-oblivious matching of incomplete models. *Computer Graphics Forum 29*, 5 (2010), 1545–1554. 13, 15, 16, 25

[DLM09] DI FABIO B., LANDI C., MEDRI F.: Recognition of occluded shapes using size functions. In *ICIAP 2009* (2009), Foggia P., Sansone C., Vento M., (Eds.), vol. 5716 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 642–651. 16

[DP06] Del Bimbo A., Pala P.: Content-based retrieval of 3D models. *ACM Trans. on Multimedia Computing, Communications and Applications 2*, 1 (2006), 20–43. 2, 9

[DP13] Denning J. D., Pellacini F.: Meshgit: Diffing and merging meshes for polygonal modeling. *ACM Trans. Graph. 32*, 4 (July 2013), 35:1–35:10. 14, 16, 17

[DW11] Dey T. K., Wang Y.: Reeb graphs: approximation and persistence. In *SoCG '11* (2011), ACM, pp. 226–235. 13

[EH07] Edelsbrunner H., Harer J.: Persistent homology – a survey. *Contemporary Mathematics* (2007), 257–282. 15

[EH10] Edelsbrunner H., Harer J.: *Computational Topology: An Introduction*. American Mathematical Society, 2010. 2, 13

[EHB13] Escolano F., Hancock E. R., Biasotti S.: Complexity fusion for indexing Reeb digraphs. In *CAIP (1)* (2013), Wilson R. C., Hancock E. R., Bors A. G., Smith W. A. P., (Eds.), vol. 8047 of *Lecture Notes in Computer Science*, Springer, pp. 120–127. 25

[ELZ02] Edelsbrunner H., Letscher D., Zomorodian A.: Topological persistence and simplification. *Discrete Computational Geometry 28* (2002), 511–533. 5, 13

[FKMS05] Funkhouser T., Kazhdan M., Min P., Shilane P.: Shape-based retrieval and analysis of 3D models. *Commun. ACM 48*, 6 (June 2005), 58–64. 2, 16

[FM99] Frosini P., Mulazzani M.: Size homotopy groups for computation of natural size distances. *Bulletin of the Belgian Mathematical Society 6* (1999), 455–464. 5, 12

[FS06] Funkhouser T., Shilane P.: Partial matching of 3D shapes with priority-driven search. In *SGP '06* (Cagliari, Sardinia, Italy, 2006), Eurographics Association, pp. 131–142. 16

[FS10] Ferri M., Stanganelli I.: Size functions for the morphological analysis of melanocytic lesions. *International Journal of Biomedical Imaging 2010* (2010), Article ID 621357. 13

[GAP*09] Giorgi D., Attene M., Patanè G., Marini S., Pizzi C., Biasotti S., Spagnuolo M., Falcidieno B., Corvi M., Usai L., Roncarolo L., Garibotto G.: A Critical Assessment of 2D and 3D Face Recognition Algorithms. In *AVSS* (2009), Tubaro S., Dugelay J.-L., (Eds.), IEEE Computer Society, pp. 79–84. 17

[GBAL09] Gębal K., Bærentzen J. A., Aanæs H., Larsen R.: Shape analysis using the auto diffusion function. In *SGP '09* (2009), Eurographics Association, pp. 1405–1413. 5, 15

[GCO06] Gal R., Cohen-Or D.: Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph. 25*, 1 (Jan. 2006), 130–150. 16

[GDP*05] Grinspun E., Desbrun M., Polthier K., Schröder P., Stern. A.: Discrete differential geometry: An applied introduction. In *ACM SIGGRAPH 2005 Courses* (New York, NY, USA, 2005), SIGGRAPH '05, ACM. 6

[GDZ10] Gao Y., Dai Q., Zhang N.-Y.: 3D model comparison using spatial structure circular descriptor. *Pattern Recognition 43*, 3 (2010), 1142 – 1151. 16, 25

[GH10] Gamble J., Heo G.: Exploring uses of persistent homology for statistical analysis of landmark-based shape data. *Journal of Multivariate Analysis 101*, 9 (2010), 2184 – 2199. 14

[GL12] Giachetti A., Lovato C.: Radial symmetry detection and shape characterization with the multiscale area projection transform. *Computer Graphics Forum 31*, 5 (2012), 1669–1678. 14, 15, 16, 25

[GM08] Giorgi D., Marini S.: Shape retrieval contest 2008: Classification of watertight models. In *Shape Modeling International 2008* (June 4-6 2008), IEEE, pp. 219–220. 2, 3

[GMGP05] Gelfand N., Mitra N. J., Guibas L. J., Pottmann H.: Robust global registration. In *SGP '05* (Aire-la-Ville, Switzerland, Switzerland, 2005), Eurographics Association. 15

[Gro99] Gromov M.: *Metric Structures for Riemannian and Non-Riemannian Spaces Couverture*. Springer, 1999. 10, 11

[HAvL05] Hein M., Audibert J.-Y., von Luxburg U.: From graphs to manifolds - weak and strong pointwise consistency of graph laplacians. In *COLT 2005* (2005), Auer P., Meir R., (Eds.), vol. 3559 of *Lecture Notes in Computer Science*, Springer, pp. 470–485. 6

[HH09] Hu J., Hua J.: Salient spectral geometric features for shape matching and retrieval. *The Visual Computer 25*, 5-7 (2009), 667–675. 16, 25

[HHZQ12] Hou T., Hou X., Zhong M., Qin H.: Bag-of-feature-graphs: A new paradigm for non-rigid shape retrieval. In *ICPR 2012* (Nov 2012), pp. 1513–1516. 25

[Hir97] Hirsch M. W.: *Differential Topology*. Springer, 1997. 5, 8

[HKR93] Huttenlocher D. P., Klanderman G. A., Rucklidge W. A.: Comparing images using the Hausdorff distance. *IEEE Trans. Pattern Anal. Mach. Intell. 15*, 9 (Sept. 1993), 850–863. 10

[HSKK01] Hilaga M., Shinagawa Y., Kohmura T., Kunii T. L.: Topology matching for fully automatic similarity estimation of 3D shapes. In *SIGGRAPH '01: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (Los Angeles, CA, August 2001), ACM Press, pp. 203–212. 13, 14, 15

[HZG*12] Huang Q.-X., Zhang G.-X., Gao L., Hu S.-M., Butscher A., Guibas L.: An optimization approach for extracting and encoding consistent maps in a shape collection. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 167:1–167:11. 17

[KBB*12] Kovnatsky A., Bronstein M. M., Bronstein A. M., Raviv D., Kimmel R.: Affine-invariant photometric heat kernel signatures. In *3DOR 2012* (2012), Spagnuolo M., Bronstein M. M., Bronstein A. M., Ferreira A., (Eds.), Eurographics Association, pp. 39–46. 14, 25

[KBB*13] Kovnatsky A., Bronstein M. M., Bronstein A. M., Glashoff K., Kimmel R.: Coupled quasi-harmonic bases. *Computer Graphics Forum 32* (2013), 439–448. 16, 17, 25

[KBBK11] Kovnatsky A., Bronstein M. M., Bronstein A. M., Kimmel R.: Photometric heat kernel signatures. In *SSVM* (2011), Bruckstein A. M., ter Haar Romeny B. M., Bronstein A. M., Bronstein M. M., (Eds.), vol. 6667 of *Lecture Notes in Computer Science*, Springer, pp. 616–627. 14, 25

[KLF11] Kim V. G., Lipman Y., Funkhouser T.: Blended intrinsic maps. *ACM Trans. Graph. 30*, 4 (July 2011), 79:1–79:12. 16, 25

[KLM*12] Kim V. G., Li W., Mitra N. J., DiVerdi S., Funkhouser T.: Exploring collections of 3D models using fuzzy correspondences. *ACM Trans. Graph. 31*, 4 (July 2012), 54:1–54:11. 16

[Koe90] Koenderink J. J.: *Solid shape*. MIT Press, Cambridge, MA, USA, 1990. 1

[KRB*13] Kovnatsky A., Raviv D., Bronstein M. M., Bronstein A. M., Kimmel R.: Geometric and photometric data fusion in non-rigid shape analysis. *NM-TMA 6*, 1 (2013), 199–222. 14, 15, 25

[Kru78] KRUMHANSL C.: Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review 85*, 5 (1978), 445–463. 9

[Lag10] LAGA H.: Semantics-driven approach for automatic selection of best views of 3D shapes. In *3DOR 2010* (2010), Eurographics Association, pp. 15–22. 18, 25

[Lav11] LAVOUÉ G.: Bag of Words and Local Spectral Descriptor for 3D Partial Shape Retrieval. In *3DOR 2011* (Llandudno, Wales, UK, 2011), Laga H., Schreck T., Ferreira A., Godil A., Pratikakis I., Veltkamp R., (Eds.), Eurographics Association, pp. 41–48. 25

[Lav12] LAVOUÉ G.: Combination of bag-of-words descriptors for robust partial shape retrieval. *The Visual Computer 28*, 9 (2012), 931–942. 25

[LB14] LITMAN R., BRONSTEIN A.: Learning spectral descriptors for deformable shape correspondence. *IEEE Trans. Pattern Anal. Mach. Intell. 36*, 1 (Jan 2014), 171–180. 14, 16, 25

[LBB11] LITMAN R., BRONSTEIN A. M., BRONSTEIN M. M.: Diffusion-geometric maximally stable component detection in deformable shapes. *Comput. Graph. 35*, 3 (June 2011), 549–560. 15, 16, 17, 25

[LBH13] LI C., BEN HAMZA A.: A multiresolution descriptor for deformable 3D shape retrieval. *The Visual Computer 29*, 6-8 (2013), 513–524. 25

[LBZ*13] LIU Z., BU S., ZHOU K., GAO S., HAN J., WU J.: A survey on partial retrieval of 3D shapes. *Journal of Computer Science and Technology 28*, 5 (2013), 836–851. 2, 17

[LD11] LIPMAN Y., DAUBECHIES I.: Conformal Wasserstein distances: Comparing surfaces in polynomial time. *Advances in Mathematics 227*, 3 (2011), 1047 – 1077. 16, 25

[LGB*11] LIAN Z., GODIL A., BUSTOS B., DAOUDI M., HERMANS J., KAWAMURA S., KURITA Y., LAVOUÉ G., NGUYEN H. V., OHBUCHI R., OHKITA Y., OHISHI Y., PORIKLI F., REUTER M., SIPIRAN I., SMEETS D., SUETENS P., TABIA H., VANDERMEULEN D.: SHREC'11 Track: Shap retrieval on non-rigid 3D watertight meshes. In *3DOR 2011* (2011), Laga H., Schreck T., Ferreira A., Godil A., Pratikakis I., Veltkamp R. C., (Eds.), Eurographics Association, pp. 79–88. 2

[LGB*13] LIAN Z., GODIL A., BUSTOS B., DAOUDI M., HERMANS J., KAWAMURA S., KURITA Y., LAVOUÉ G., VAN NGUYEN H., OHBUCHI R., OHKITA Y., OHISHI Y., PORIKLI F., REUTER M., SIPIRAN I., SMEETS D., SUETENS P., TABIA H., VANDERMEULEN D.: A comparison of methods for non-rigid 3D shape retrieval. *Pattern Recogn. 46*, 1 (Jan. 2013), 449–461. 2, 16

[LGPC13] LOMBAERT H., GRADY L., POLIMENI J. R., CHERIET F.: FOCUSR: Feature oriented correspondence using spectral regularization–A method for precise surface matching. *IEEE Trans. Pattern Anal. Mach. Intell. 35*, 9 (2013), 2143–2160. 17

[LH13] LI C., HAMZA A.: Symmetry discovery and retrieval of nonrigid 3D shapes using geodesic skeleton paths. *Multimedia Tools and Applications* (2013), 1–21. 15, 16, 17, 25

[LJ07] LING H., JACOBS D. W.: Shape classification using the inner-distance. *IEEE Trans. Pattern Anal. Mach. Intell. 29*, 2 (Feb. 2007), 286–299. 17

[LMM13] LI P., MA H., MING A.: Combining topological and view-based features for 3D model retrieval. *Multimedia Tools and Applications 65*, 3 (2013), 335–361. 14, 25

[LMS13] LAGA H., MORTARA M., SPAGNUOLO M.: Geometry and context for semantic correspondences and functionality recognition in man-made 3D shapes. *ACM Trans. Graph. 32*, 5 (Oct. 2013), 150:1–150:16. 16, 17, 18

[LO07] LING H., OKADA K.: An efficient Earth Mover's distance algorithm for robust histogram comparison. *IEEE Trans. Pattern Anal. Mach. Intell. 29*, 5 (May 2007), 840–853. 16

[Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *Int. J. of Comput. Vision 60*, 2 (2004), 91–110. 10

[LPD13] LIPMAN Y., PUENTE J., DAUBECHIES I.: Conformal Wasserstein distance: II. computational aspects and extensions. *Math. Comput. 82*, 281 (2013). 16, 25

[LTN06] LAGA H., TAKAHASHI H., NAKAJIMA M.: Spherical parameterization and geometry image-based 3D shape similarity estimation (cgs 2004 special issue). *The Visual Computer 22*, 5 (2006), 324–331. 14

[LV08] LÉVY B., VALLET B.: Spectral geometry processing with manifold harmonics. *Computer Graphics Forum 2*, 27 (2008). 6

[LZSCO09] LIU R., ZHANG H., SHAMIR A., COHEN-OR D.: A part-aware surface metric for shape analysis. *Computer Graphics Forum 28*, 2 (2009), 397–406. 17, 25

[MÍ11] MÉMOLI F.: Gromov-Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of Computational Mathematics* (2011), 1–71. 10.1007/s10208-011-9093-5. 11

[MÍ12] MÉMOLI F.: Some properties of Gromov-Hausdorff distances. *Discrete & Computational Geometry* (2012), 1–25. 11

[MAD*11] MAALEJ A., AMOR B. B., DAOUDI M., SRIVASTAVA A., BERRETTI S.: Shape analysis of local facial patches for 3D facial expression recognition. *Pattern Recognition 44*, 8 (2011), 1581–1589. 17

[Man88] MANTYLA M.: *Introduction to Solid Modeling*. WH Freeman & Co. New York, NY, USA, 1988. 7, 14

[MBH12] MOHAMED W., BEN HAMZA A.: Reeb graph path dissimilarity for 3D object matching and retrieval. *The Visual Computer 28*, 3 (2012), 305–318. 16, 25

[MDTS09] MADEMLIS A., DARAS P., TZOVARAS D., STRINTZIS M. G.: 3D object retrieval using the 3D shape impact descriptor. *Pattern Recognition 42*, 11 (2009), 2447 – 2459. 9, 14, 15, 16, 25

[MFK*10] MAES C., FABRY T., KEUSTERMANS J., SMEETS D., SUETENS P., VANDERMEULEN D.: Feature detection on 3D face surfaces for pose normalisation and recognition. In *BTAS 2010* (Sept 2010), pp. 1–6. 25

[MGP06] MITRA N. J., GUIBAS L., PAULY M.: Partial and approximate symmetry detection for 3D geometry. *ACM Trans. Graph. 25*, 3 (2006), 560–568. 16

[MHK*08] MATEUS D., HORAUD R., KNOSSOW D., CUZZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *CVPR 2008* (June 2008), pp. 1–8. 14, 25

[Mil63] MILNOR J.: *Morse Theory*. Princeton University Press, New Jersey, 1963. 5, 8

[Mil65] MILNOR J.: *Lectures on h-cobordism*. Princeton University Press, 1965. 8

[ML03] MOLER C., LOAN C. V.: Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review 45*, 1 (2003), 3–49. 7

[MMH11] MILEYKO Y., MUKHERJEE S., HARER J.: Probability measures on the space of persistence diagrams. *Inverse Problems 27*, 12 (2011), 124007. 14

[Mor86] MORTENSON M. E.: *Geometric Modeling*. John Wiley & Sons, 1986. 14

[MPWC13] MITRA N. J., PAULY M., WAND M., CEYLAN D.: Symmetry in 3D geometry: Extraction and applications. *Computer Graphics Forum 32*, 6 (2013), 1–23. 2

[MS05] MÉMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math. 5*, 3 (July 2005), 313–347. 3, 11

[MS09] MAHMOUDI M., SAPIRO G.: Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models 71*, 1 (2009), 22 – 31. 14, 25

[Mun00] MUNKRES J.: *Topology*. Prentice Hall, 2000. 7

[NBPF11] NATALI M., BIASOTTI S., PATANÃĹ G., FALCIDIENO B.: Graph-based representations of point clouds. *Graphical Models 73*, 5 (2011), 151 – 164. 14, 25

[NSCE02] N. A., SANTA-CRUZ D., EBRAHIMI T.: Mesh: measuring errors between surfaces using the hausdorff distance. In *ICME '02* (2002), vol. 1, pp. 705–708 vol.1. 10

[OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM Trans. Graph. 31*, 4 (July 2012), 30:1–30:11. 2, 3, 11, 12, 25

[OBGB09] OVSJANIKOV M., BRONSTEIN A. M., GUIBAS L. J., BRONSTEIN M. M.: Shape google: a computer vision approach to invariant shape retrieval. In *Proc. NORDIA* (2009). 15, 25

[OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Trans. Graph. 21*, 4 (Oct. 2002), 807–832. 10

[OHG11] OVSJANIKOV M., HUANG Q.-X., GUIBAS L.: A condition number for non-rigid shape matching. *Computer Graphics Forum 30*, 5 (2011), 1503–1512. 18

[OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *Computer Graphics Forum 29*, 5 (2010), 1555–1564. 16, 25

[OMPG13] OVSJANIKOV M., MÉRIGOT Q., PĂTRĂUCEAN V., GUIBAS L.: Shape matching via quotient spaces. *Computer Graphics Forum 32*, 5 (2013), 1–11. 12

[PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum 32*, 2pt4 (2013), 459–468. 16, 25

[PBF07] PARABOSCHI L., BIASOTTI S., FALCIDIENO B.: Comparing sets of 3D digital shapes through topological structures. In *0GbR'7* (2007), Escolano F., Vento M., (Eds.), vol. 4538 of *Lecture Notes in Computer Science*, pp. 114–125. 14

[PS13] PATANÉ G., SPAGNUOLO M.: Heat diffusion kernel and distance on surface meshes and point sets. *Computers & Graphics 37*, 6 (2013), 676 – 686. 7

[RBB*11] RAVIV D., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., SOCHEN N.: Affine-invariant geodesic geometry of deformable 3D shapes. *Computers & Graphics 35*, 3 (2011), 692 – 697. 2, 15, 16, 17, 25

[RBBK10a] RAVIV D., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Full and partial symmetries of non-rigid shapes. *Int. J. Comput. Vision 89*, 1 (Aug. 2010), 18–39. 16, 17

[RBBK10b] RAVIV D., BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R.: Volumetric heat kernel signatures. In *Proceedings of the ACM Workshop on 3D Object Retrieval* (New York, NY, USA, 2010), 3DOR 2010, ACM, pp. 39–44. 14, 16, 25

[RBG*09] REUTER M., BIASOTTI S., GIORGI D., PATANÈ G., SPAGNUOLO M.: Discrete laplace-beltrami operators for shape analysis and segmentation. *Comput. Graph. 33*, 3 (June 2009), 381–390. 6

[Ree46] REEB G.: Sur les points singuliers d'une forme de Pfaff complètement intégrable ou d'une fonction numérique. *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences 222* (1946), 847–849. 13

[Req80] REQUICHA A.: Representations of rigid solids: Theory, methods and systems. *ACM Computing Surveys 12*, 4 (1980), 437–464. 7, 14

[Reu06] REUTER M.: *Laplace Spectra for Shape Recognition*. Books on Demand, ISBN 3-8334-5071-1, 2006. 6

[ROA*13] RUSTAMOV R. M., OVSJANIKOV M., AZENCOT O., BEN-CHEN M., CHAZAL F., GUIBAS L.: Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graph. 32*, 4 (July 2013), 72:1–72:12. 12, 16, 17, 25

[Ros97] ROSENBERG S.: *The Laplacian on a Riemannian Manifold*. Cambridge University Press, 1997. 6

[RPSS10] RUGGERI M. R., PATANÉ G., SPAGNUOLO M., SAUPE D.: Spectral-driven isometry-invariant matching of 3D shapes. *Int. J. of Comput. Vision 89*, 2-3 (2010), 248–265. 16, 25

[Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *SGP'07* (2007), Eurographics Association, pp. 225–233. 15

[Rus10] RUSTAMOV R. M.: Robust volumetric shape descriptor. In *3DOR 2010* (2010), Daoudi M., Schreck T., Spagnuolo M., Pratikakis I., Veltkamp R. C., Theoharis T., (Eds.), Eurographics Association, pp. 1–5. 14, 16, 17, 25

[RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-Beltrami spectra as "Shape-DNA" of surfaces and solids. *Computer-Aided Design 38*, 4 (2006), 342–366. 15

[SB11] SKOPAL T., BUSTOS B.: On nonmetric similarity search problems in complex domains. *ACM Comput. Surv. 43*, 4 (Oct. 2011), 34:1–34:50. 1, 2, 8, 9

[SHCB11] SHARMA A., HORAUD R., CECH J., BOYER E.: Topologically-robust 3D shape matching based on diffusion geometry and seed growing. *CVPR 2013* (2011), 2481–2488. 16, 17, 25

[SHVS12] SMEETS D., HERMANS J., VANDERMEULEN D., SUETENS P.: Isometric deformation invariant 3D shape recognition. *Pattern Recognition 45*, 7 (2012), 2817 – 2831. 9, 17, 25

[Sin06] SINGER A.: From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis 21*, 1 (July 2006), 128–134. 6

[SJ99] SANTINI S., JAIN R.: Similarity measures. *IEEE Trans. Pattern Anal. Mach. Intell. 21*, 9 (Sept. 1999), 871–883. 1, 9

[SK91] SHINAGAWA Y., KUNII T. L.: Constructing a Reeb Graph automatically from cross sections. *IEEE Computer Graphics and Applications 11*, 6 (nov 1991), 44–51. 13

[SKK91] SHINAGAWA Y., KUNII T. L., KERGOSIEN Y. L.: Surface coding based on Morse theory. *IEEE Computer Graphics and Applications 11* (1991), 66–78. 13

[SKVS13] SMEETS D., KEUSTERMANS J., VANDERMEULEN D., SUETENS P.: meshSIFT: Local surface features for 3D face recognition under expression variations and partial data. *Comput. Vis. Image Und. 117*, 2 (2013), 158 – 169. 25

[SMKF04] SHILANE P., MIN P., KAZHDAN M., FUNKHOUSER T.: The Princeton shape benchmark. In *SMI 2004* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 167–178. 17

[SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. *Computer Graphics Forum 31*, 5 (2012), 1617–1626. 12, 25

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum 28*, 5 (2009), 1383–1392. 2, 5, 7, 15, 25

[Spa66] SPANIER E. H.: *Algebraic Topology*. McGraw Hill, 1966. 7

[SPT11] SFIKAS K., PRATIKAKIS I., THEOHARIS T.: Contopo: Non-rigid 3D object retrieval using topological information guided by conformal factors. In *3DOR 2011* (Llandudno, UK, 2011), Eurographics Association, pp. 25–32. 16, 25

[SSCO08] SHAPIRA L., SHAMIR A., COHEN-OR D.: Consistent mesh partitioning and skeletonisation using the shape diameter function. *Vis. Comput. 24*, 4 (Mar. 2008), 249–259. 10

[SSS*10] SHAPIRA L., SHALOM S., SHAMIR A., COHEN-OR D., ZHANG H.: Contextual part analogies in 3D objects. *Int. J. of Comput. Vision 89*, 2-3 (2010), 309–326. 10, 16, 25

[STP12] SFIKAS K., THEOHARIS T., PRATIKAKIS I.: Non-rigid 3D object retrieval using topological information guided by conformal factors. *The Visual Computer 28*, 9 (2012), 943–955. 16, 25

[TCL*13] TAM G., CHENG Z.-Q., LAI Y.-K., LANGBEIN F., LIU Y., MARSHALL D., MARTIN R., SUN X.-F., ROSIN P.: Registration of 3D point clouds and meshes: A survey from rigid to nonrigid. *IEEE Trans. on Vis. and Computer Graphics 19*, 7 (July 2013), 1199–1217. 2, 9, 14, 17

[TDVC11] TABIA H., DAOUDI M., VANDEBORRE J.-P., COLOT O.: A new 3D-matching method of nonrigid and partially similar models using curve analysis. *IEEE Trans. Pattern Anal. Mach. Intell. 33*, 4 (2011), 852–858. 16

[TDVC13] TABIA H., DAOUDI M., VANDEBORRE J.-P., COLOT O.: A parts-based approach for automatic 3D shape categorization using belief functions. *ACM TIST 4*, 2 (2013), 33. 18

[TG82] TVERSKY A., GATI I.: Similarity, separability and the triangle inequality. *Psychological Review 89* (1982), 123–154. 9

[tHV10] TER HAAR F. B., VELTKAMP R. C.: Expression modeling for expression-invariant face recognition. *Computers & Graphics 34*, 3 (2010), 231–241. 17

[TS04] TUNG T., SCHMITT F.: Augmented Reeb graphs for content-based retrieval of 3D mesh models. In *SMI '04* (Los Alamitos, June 2004), IEEE Computer Society Press, pp. 157–166. 14

[TS05] TUNG T., SCHMITT F.: The Augmented Multiresolution Reeb Graph approach for content-based retrieval of 3D shapes. *Int. J. of Shape Modelling 11*, 1 (June 2005), 91–120. 14

[TV04] TANGELDER J., VELTKAMP R.: A survey of content based 3D shape retrieval methods. In *Proceedings of Shape Modeling Applications, 2004* (2004), pp. 145–156. 2, 9, 17

[TV08] TANGELDER J. W. H., VELTKAMP R. C.: A survey of content based 3D shape retrieval methods. *Multimedia Tools Appl. 39*, 3 (2008), 441–471. 2, 9, 14, 16

[TVD09] TIERNY J., VANDEBORRE J.-P., DAOUDI M.: Partial 3D shape retrieval by Reeb pattern unfolding. *Comput. Graph. Forum 28*, 1 (2009), 41–55. 16, 25

[Tve77] TVERSKY A.: Features of similarity. *Psychological Review 84* (1977), 327–352. 1, 9

[VH01] VELTKAMP R. C., HAGENDOORN M.: State-of-the-Art in shape matching. In *Principles of Visual Information Retrieval*, Lew M., (Ed.). Springer-Verlag, 2001, pp. 87–119. 2

[vKZH13] VAN KAICK O., ZHANG H. R., HAMARNEH G.: Bilateral maps for partial matching. *Computer Graphics Forum 32*, 6 (2013), 189–200. 15, 16, 25

[vKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Comput. Graph. Forum 30*, 6 (2011), 1681–1707. 2, 9, 16

[VRS*06] VELTKAMP R., RUIJSENAARS R., SPAGNUOLO M., VAN ZWOL R., TER HAAR F.: *SHREC2006: 3D Shape Retrieval Contest* . Tech. Rep. CS-2006-030, UU, 2006. 17

[WBBP12] WANG C., BRONSTEIN M. M., BRONSTEIN A. M., PARAGIOS N.: Discrete minimum distortion correspondence problems for non-rigid shape matching. In *SSVM'11* (Berlin, Heidelberg, 2012), Springer-Verlag, pp. 580–591. 11, 25

[WLZ10] WANG X., LIU Y., ZHA H.: Intrinsic spin images: A subspace decomposition approach to understanding 3D deformable shapes. In *Proc. 3DPVT* (2010), pp. 17 – 20. 25

[WMKG07] WARDETZKY M., MATHUR S., KÄLBERER F., GRINSPUN E.: Discrete Laplace operators: No free lunch. In *SGP'07* (2007), pp. 33–37. 6

[WZL*10] WU H.-Y., ZHA H., LUO T., WANG X.-L., MA S.: Global and local isometry-invariant descriptor for 3D shape comparison and partial matching. In *CVPR 2010* (June 2010), pp. 438–445. 6, 9, 15, 16

[XSW03] XIAO Y., SIEBERT P., WERGHI N.: A discrete Reeb graph approach for the segmentation of human body scans. In *3DIM 2003* (Los Alamitos, 2003), IEEE Computer Society Press, pp. 378–385. 14

[YLZ07] YANG Y., LIN, ZHANG: Content-Based 3-D Model Retrieval: A Survey. *IEEE Trans. on Sys., Man and Cybernetics, Part C: Appl. and Reviews 37*, 6 (Nov. 2007), 1081–1098. 2

[ZADB06] ZEZULA P., AMATO G., DOHNAL V., BATKO M.: *Similarity Search: The Metric Space Approach*, vol. 32 of *Advances in Database Systems*. Springer, 2006. 8

[ZBH12] ZAHARESCU A., BOYER E., HORAUD R.: Keypoints and local descriptors of scalar functions on 2D manifolds. *Int. J. of Comput. Vision 100*, 1 (2012), 78–98. 10, 14, 15, 25

[ZBVH09] ZAHARESCU A., BOYER E., VARANASI K., HORAUD R.: Surface feature detection and description with applications to mesh matching. In *CVPR 2009* (June 2009), pp. 373–380. 10, 14, 15, 16, 25

[Zha94] ZHANG Z.: Iterative point matching for registration of free-form curves and surfaces. *Int. J. Comput. Vision 13*, 2 (Oct. 1994), 119–152. 10

Table 1: *Classification with respect to the type of structure, distance, input, invariance, output (similarity and/or correspondence) characterizing each method. G-H and NPs abbreviate Gromov-Hausdorff and natural pseudo-distance, respectively.*

Method (refs.)	Structure	Distance	Input				Invariance			Similarity		Correspondence		
			point	surface	volume	texture	rigid	isometry	other	full	part.	sparse	dense	soft
Blended Intrinsic Maps [KLF11]	Conformal	Wasserstein-based		√			√	√	non-isometric	√		√		
Conformal Wasserstein [LD11, LPD13]	Conformal	Wasserstein-based		√			√	√		√		√		
Geometric similarity [BLC*11]	Conformal	Wasserstein		√			√	√		√		√		
Diffusion distance distr. [MS09]	Diffusion	L^1, L^2, χ^2, J-S divergence	√				√	√		√		√		
Volume GPS [Rus10]	Diffusion	χ^2			√		√	√	scale	√				
Heat Kernel Signature [SOG09]	Multi-scale autodiffusion	L^1		√			√	√		√		√		
Scale Invariant HKS [BK10c]	Multi-scale autodiffusion	Weighted L^1		√			√	√	scale	√		√		
Shape Google [OBGB09, BBGO11]	Multi-scale autodiffusion	Hamming distance		√			√	√	scale	√		√		
Bag of Feature Graphs [HHZQ12]	Multi-scale autodiffusion	L^2-based		√			√	√	scale	√		√		
Topology-Invariant Geometry [BBK09]	Diffusion	Pareto distance		√			√	√	scale	√		√		
G-H framework & diffusion [BBK*10]	Diffusion	G-H framework		√			√	√	scale	√		√		
Spectral Distances [BB11]	Diffusion	Normalized L^1, L^2, χ^2, EMD		√			√	√	scale	√		√		
Min. Distortion Corr. [WBBP12]	Multi-scale autodiffusion	G-H framework		√			√	√	scale	√		√		
Soft Maps [SNB*12]	Local geometry	Generalized EMD		√			√	√	non-isometric	√	√		√	√
Appr. joint corr. [KBB*13, PBB*13]	Diffusion	Functional map		√			√	√	scale	√	√		√	√
Equi-affine inv. geometry [RBB*11]	Diffusion	G-H distance		√	√		√	√	scale, affinity	√		√		
One-point isometry match. [OMMG10]	Multi-scale autodiffusion	L^2-type		√			√	√		√	√		√	
Topologically-robust matching [SHCB11]	Multi-scale autodiffusion	L^2, L^∞ type		√			√	√		√		√	√	
Volumetric HKS [RBBK10b]	Multi-scale autodiffusion	L^1			√		√	√	scale	√		√		
Contextual part analogies [SSS*10]	Semilocal conformal	Bipartite graph matching		√			√	√	scale	√	√	√		
Isometric deformation inv. [SHVS12]	Conformal	χ^2		√			√	√	scale	√				
Intrinsic Diff. [OBCS*12, ROA*13]	Conformal	Functional map		√			√	√		√	√		√	√
MDM [Bia10]	Diffusion	EMD		√			√	√		√				
Wavelet Kernel Signature [ASC11]	Multi-scale autodiffusion	L^1		√			√	√		√		√		
3D MSERs [LBB11]	Semi-local autodiffusion	Local point & region distance		√			√	√		√	√	√		
Spectral descriptors [LB14]	Multi-scale autodiffusion	Metric learning		√			√	√	scale, affinity	√		√		
Spectral graph wavelets [LBH13]	Diffusion	Intrinsic spatial pyramid match.		√			√	√	scale	√				
Part-aware metric [LZSCO09]	Conformal	χ^2		√			√	√		√	√	√		
3D shape impact [MDTS09]	Semi-local extrinsic	Normalized L^2 and diffusion distance		√	√		√		scale	√				
Spatial circular descriptor [GDZ10]	Extrinsic	Hungarian distance		√			√		scale	√				
meshSIFT [MFK*10, SKVS13]	Multi-scale geometry	Angle distance		√			√	√		√			√	
Salient Points [CCFM08]	Multi-scale geometry	Distance on Hidden Markov Models		√			√	√	scale	√		√		
Salient spectral features [HH09]	Diffusion	Ad-hoc distance		√			√	√	scale	√				
Bag of words [Lav11, Lav12]	Local diffusion	L^1-based		√			√	√	scale	√	√	√		
Intrinsic spin images [WLZ10]	Multi-scale geometry	EMD		√			√	√		√		√		
LB eigendecomposition [DK10]	Diffusion	Min. of L^1 & geodesic dist.		√			√	√		√	√	√		
Spectral isometry match. [RPSS10]	Diffusion	Bipartite graph matching, L^1-type		√			√	√	scale	√		√		
Semantic best view selection [Lag10]	Conformal	Ad hoc distance		√	√					√	√			
Interest points histograms [BWDP13b]	Multi-scale conformal	χ^2, RANSAC		√			√	√	scale	√		√		
Sparse Matching [BDP13b]	Multi-scale geometry	Ad-hoc distance, RANSAC		√			√	√	scale		√	√		
Facial expression recogn. [BDP13a]	Multi-scale geometry	Distance on Hidden Markov Models		√			√			√		√		
Isogeodesic stripes [BDP10]	Conformal	Weighted graph distance		√			√	√	scale	√		√		
Bilateral maps [vKZH13]	Semi-local conformal	L^1-based		√			√	√	scale	√	√	√		
Persistence [CCSG*09]	Multi-scale topology	G-H framework		√			√	√	scale	√		√		
Persistence [DLL*10]	Multi-scale autodiffusion & topology	L^1-based		√			√	√			√	√		
Size functions [BCF*08, BCFG11]	Multi-scale geometry & topology	(NPs) framework		√	√		√	√	scale	√		√		
Persistence [DL12]	Extrinsic & multi-scale topology	(NPs) framework		√			√			√	√	√		
Extended Reeb graph [BB13a, BB13b]	Semi-local geometry & topology	Kernel aggregation & learning		√			√	√	scale	√				
LB eigenfunction match [MHK*08]	Diffusion	Hungarian distance	√				√							
Skeleton path [LH13]	Extrinsic & topology	Endpoint distance		√	√		√	√	scale	√		√		
Point cloud graph [NBPF11]	Semi-local geometry & topology	Spectral graph distance	√				√	√	scale	√		√		
Extended Reeb graph [BEGB13, EHB13]	Semi-local geometry & topology	Spectral graph distance		√			√	√	scale	√		√		
Topo-geometric model [BK10a]	Extrinsic & topology	Max sub-graph approx.		√			√		scale	√	√	√		
Skeletal Reeb graphs [MBH12]	Extrinsic & topology	Shortest path graph matching		√			√		scale	√	√	√		
Reeb graph [AK11]	Conformal & topology	Max sub-graph approx.		√			√	√		√		√		
Reeb graph [TVD09]	Semi-local conformal & topology	Max sub-graph approx.		√			√	√	scale	√	√	√		
ComTopo [SPT11, STP12]	Semi-local conformal & topology	Hungarian distance		√			√	√		√		√		
Graph based match. [APP*09, APP*10]	Semi-local conformal	EMD		√			√	√		√		√		
Reeb graph & view [LMM13]	Conformal & multi-scale topology	EMD		√			√	√		√		√		
MeshHOG [ZBVH09, ZBH12]	Multi-scale conformal & photometry	L^2		√		√	√	√		√			√	
Photometric HKS [KBBK11, KBB*12, KRB*13]	Multi-scale autodiffusion & photometry	L^1		√		√	√	√	scale, affinity	√		√		
Multi-scale projection transf. [GL12]	Photometry & multi-scale geometry	Jeffrey divergence		√		√	√		scale	√				
PHOG [BCGS13]	Conformal & photometry, multi-scale topology	L^1, (NPs) framework		√		√	√	√	scale	√		√		