Volume xx (200y), Number z, pp. 1–29

Recent trends, applications, and perspectives in 3D shape similarity assessment

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
{andrea.cerri,silvia.biasotti}@ge.imati.cnr.it
[2] School of Electrical Engineering, Tel Aviv University, Israel
bron@eng.tau.ac.il
[3] Institute of Computational Science, University of Lugano (USI), Switzerland
michael.bronstein@usi.ch

Abstract

The recent introduction of 3D shape analysis frameworks able to quantify the deformation of a shape into another in terms of the variation of real functions yields a new interpretation of the 3D shape similarity assessment and opens new perspectives. Indeed, while the classical approaches to similarity mainly quantify it as a numerical score, map based methods also define (dense) shape correspondences. After presenting in detail the theoretical foundations underlying these approaches, we classify them by looking at their most salient features, including the kind of structure and invariance properties they capture, as well as the distances and the output modalities according to which the similarity between shapes is assessed and returned. We also review the usage of these methods in a number of 3D shape application domains, ranging from matching and retrieval to annotation and segmentation. Finally, the most promising directions for future research developments are discussed.

Keywords: 3D shape distances, 3D shape matching, map-based correspondence

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

1. Introduction

For many years, researchers work on different aspects of shape analysis and comparison. For instance, psychologists have studied for many decades how humans perceive a shape, and how perception affects the process of assessing the similarity between shapes [SJ99, Tve77, AP88, Ash92, Koe90]. Thanks to those studies, it is now well known that formalizing the concept of shape similarity is a complex interaction process involving the observer and his/her interpretation of the geometric, structural and semantic properties of shapes [SB11].

More recently, the development of computational models and tools able to provide digital representations of shapes opened the way to the development of 3D shape analysis, thus making the problem of shape similarity a matter of study also for applied mathematicians, computer scientists and other researchers working in the fields of computer vision, computer graphics and pattern recognition, just to cite a few.

Also in this context, modelling shape similarity is definitively not an easy task: first of all, there is neither a single *best* shape representation nor a single *best* similarity measure to be used, and the comparison largely depends on the type of shapes to be analysed and on the properties that are considered relevant in the problem at hand. An intuition of this is given by Figure 1, showing some models belonging to the same "humanoid" class of the SHREC'08 classification benchmark [GM08]. Different notions of similarity might be formulated in this case, by considering in turn functional (semantic), structural or geometric criteria.

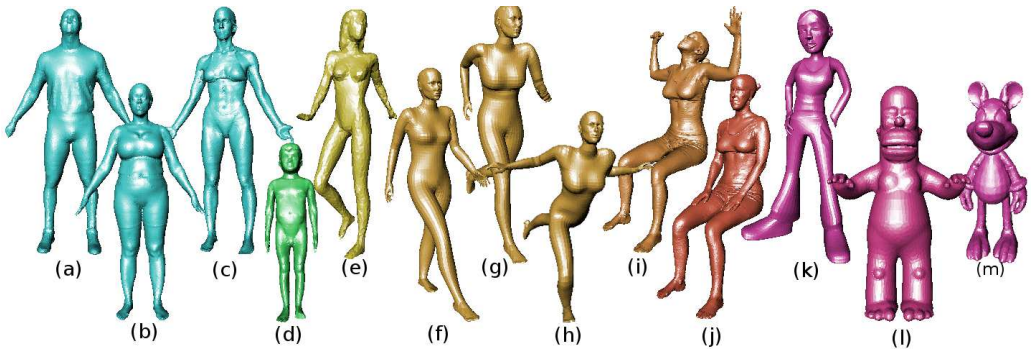

Figure 1: *Representative models of the humanoid class, SHREC'08 classification benchmark [GM08]. Models (a-d) have same pose but (d) differs by scale, (e) is a human model in a different pose and (f-h) are isometric deformations of the same template, (i-j) are two scans of the same model with a significant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

1.1. STAR focus and contribution

In this paper we review methods for the assessment of shape similarity, by specifically targeting the 3D digital world scenario. By shape similarity we mean *quantifying* how much a shape resembles to another, through either some numerical score or by evaluating the distortion of a map between shapes.

In the rapidly growing field of 3D shape analysis, a number of strategies have been proposed so far for shape similarity assessment. 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 the last years the focus is moving to the formalization of new paradigms that allow a larger flexibility in the definition of similarity. In particular, the present work is mainly motivated by the recent introduction, in the fields of computer vision, computer graphics and pattern recognition, of four theoretical and computational frameworks for 3D shape similarity that measure how much shape properties change while deforming one shape into another one. We refer in particular to the minimum-distortion shape embeddings, the Gromov-Hassdorff distance, the functional maps and the natural pseudo-distance frameworks [EK03, MS05, OBCS*12, DF04b]. As we will see later, these approaches quantify similarity not only in terms of a single score, but also define a map between shapes so that, despite the increasing computational complexity, it is possible to infer either a sparse or a dense shape correspondence.

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

- We discuss 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. Examples of such functions are distances from relevant points, heat kernel signatures or channel values in some colour space. On the other hand, distances defined on the model representations provide insights on the corresponding shape distributions, as in the case of geodesics 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.

Methods will be presented by highlighting their salient properties, the specific shape invariants they consider (e.g. rigid and non-rigid transformations), the structure they capture and at which level (e.g., local or global, conformal or diffusion structure) as well as the type of output they provide (e.g. full or partial similarity score, sparse or dense correspondence).

Other surveys on shape similarity have appeared before, reviewing topics partially overlapped to those discussed here. We refer to 3D shape retrieval [TV04, BKS*05, FKMS05, DP06, YLZ07, BKSS07, TV08, LGB*13] or other specific aspects of similarity, such as shape registration [TCL*13], shape correspondence [vKZHCO11], symmetry detection [MPWC13], partial matching [LBZ*13, SPS14] and non-metric distances [SB11]. Differently from those works, this contribution is specifically focused on the process itself of similarity assessment, which can be considered a preparatory step for most of the above topics, clearly conditioning them. We will provide a detailed discussion about the mathematical foundations underlying similarity assessment. This will allow us to present the most recent theoretical developments in the field of 3D shape similarity under an appropriate and rigorous viewpoint, which is indeed the peculiar trait of this survey.

In this view, the present paper can be seen as complementary in spirit to [BDF*08, BFGS12, BCB12], in which the focus is more on shape analysis and description rather than similarity; it also updates and extends [BFF*07] in the reviewed time-lapse, which have seen the introduction of a number of innovative techniques, with particular reference to the four aforementioned frameworks (cf. Section 3).

1.2. Organization

To drive the reader through the bunch of approaches and frameworks revised here, we first introduce the basic notions of mathematical concepts such as topological space, manifold, map, metric and transformation. We also provide a brief overview on 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 start by discussing the basic properties of similarity measures, also describing some representative methods for the algorithmic evaluation of similarity (Section 3.1). Then, we review the main theoretical and computational aspects of four formal frameworks for similarity assessment which have been recently introduced in 3D shape analysis (Sections from 3.2.1 to 3.2.4).

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 along the process. The proposed taxonomy also takes into account the type of input for each method, as well as the the output modalities according to which the similarity between shapes is returned. The aim is to give a multifaceted classification that might help the reader to compare methods not only on the basis of their algorithmic aspects, but also drive him/her in the choice of the method that better fulfils his/her requirements.

In Section 5, we review the available 3D retrieval benchmarks, which are the key to quantitatively evaluate the performance of the methods and help the reader to experiment with the tools seen in the paper. Moreover, a detailed analysis on the application domains for which methods have been proposed is carried out.

Finally, Section 6 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 some approaches may be of interest for part of the research community involved in visualization and computational geometry and topology, beyond people working in shape analysis.

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 *topology*, i.e. a collection of subsets that are referred to as *open sets*. Intuitively, a set U is open if, starting from *any* point in U and following *any* direction, it is possible to move "a little" and stay inside the set. It turns out that the notion of open set provides a fundamental way to speak of *nearness* of points, although without explicitly having a concept of distance defined on the considered topological space. Thus, once a topology has been defined, we are allowed to introduce properties such as continuity, connectedness, closeness, which are all based on some notion of nearness. These properties are in turn key ingredients to model the shape of 3D digital objects, as well as to reason about concepts like robustness and stability of shape analysis methods.

As for *maps*, they can by used to model spatial relations between two (or more) shapes represented by suitable topological spaces. Also, real- or vector-valued maps provide a means 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.

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 every pair of points can be separated by open sets.

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 [†]. 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*.

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 $d(\gamma(t_1),\gamma(t_2)) = \lambda|t_1 - t_2|$ holds for a constant $\lambda \geq 0$.

Transformations. By the term *transformation*, we refer here to *structure-preserving* maps between spaces. Transformations play an essential role in the process of assessing the similarity between shapes, see Figure 2 for an intuition about this. As we will see later, relevant transformations from the viewpoint of shape similarity include isometries, affine transformations and homeomorphisms.

Isometries are distance-preserving maps, taking elements of a metric space to elements of 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 shape is laid out in the Euclidean space.

Affine transformations, or simply *affinities*, preserve straight lines (i.e., all points initially lying on a line 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, ge-

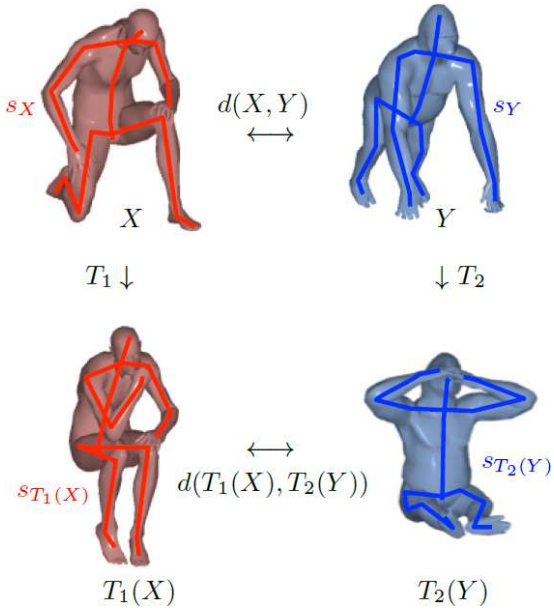

Figure 2: *Conceptual representation of invariant similarity: the distance $d(X,Y)$ between shapes X and Y is measured based on some structure s_X, s_Y extracted from them. The structure remains invariant under the shape transformations $T_i \in \mathcal{T}$ from class $\mathcal{T}$, such that $s_X = s_{T_1(X)}, s_Y = s_{T_2(Y)}$. Consequently, the distance is invariant to transformations, $d(T_1(X),T_2(Y)) = d(X,Y)$.*

[†] Note, however, that this definition of continuous function 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 derivatives of all orders are continuous.

ometric contraction, expansion, homothety, reflection, rotation, scale and compositions of them.

A more flexible 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). From the shape comparison point of view, considering homeomorphisms allows one to deal with more generic deformations, such as those in Figure 3; however, topological invariance is sometimes too coarse, 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 enrich 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] we will discuss later in this paper.

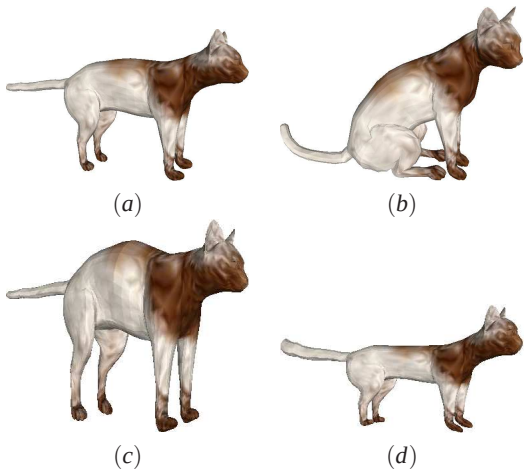

Figure 3: *A cat model* (a) *together with three homeomorphic modifications* $(b-d)$*. All models are from [BCA*14].*

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 to 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 4 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 $F : X \to T(M)$ which assigns 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.

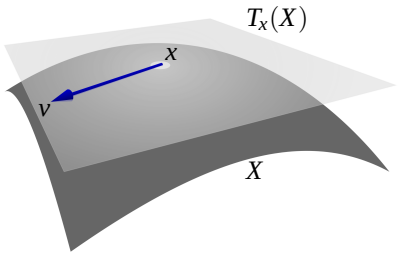

Figure 4: *Tangent plane* $T_x(X)$ *in x. The vector* $v \in T_x(X)$ *is a tangent vector.*

Note that, in practice, a Riemannian metric is a positive definite, symmetric tensor. Indeed, once a local system 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 of 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 in turn closely connected with the notion of *Laplace operator*.

2.4.1. Laplace operator

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

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

where grad and div are the gradient and divergence on the space, and the point $x \in \mathbb{E}^n$ is represented by the Cartesian coordinates $x = (x_1, \ldots, x_n)$. 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 $\Delta f(x)$ measures the difference between $f(x)$ and its average in a small neighbourhood of $x \in \mathbb{E}^n$.

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, $0 \leq \lambda_1 \leq \lambda_2 \leq \ldots$. 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 [VL08, 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. 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 [VL08], allowing for practical computation on a manifold discretization. For example, suppose to have a triangulation T with set of vertices $P := \{p_i, i = 1, \ldots, n\}$. 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, each one 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 $V = \operatorname{diag}(v_1, \ldots, v_n)$ is the diagonal matrix whose elements are $v_i = \sum_{j \in N(i)} w_{ij}$, $W = (w_{ij})$ and $D = \operatorname{diag}(d_1, \ldots, d_n)$, 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

[‡] 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.

described by the partial differential equation:

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

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 heat source point 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 5). By the eigendecomposition of Δ, the heat kernel can be written as

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

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

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

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 of x. Finally, the heat kernel is stable under small perturbations of the underlying manifold. All these properties make the heat kernel a good candidate for the definition of informative functions and distances to be used for shape description, such as the heat kernel signature (HKS) [SOG09, GBAL09] 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.

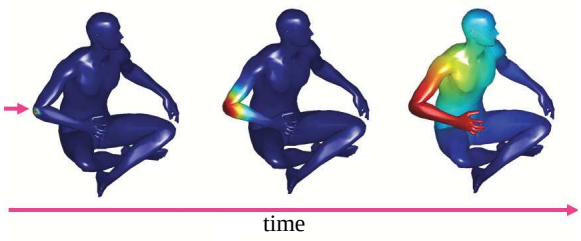

Figure 5: *The heat kernel represents the amount of heat transferred from a source point in time t.*

2.5. Basics of 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.

2.5.1. Basics of Homology

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 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*. A simplicial complex is a topological space that can be obtained by gluing together simple elements, called *simplices*, in a structured way. Figure 6 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). 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. Basics of Morse theory

Morse theory can be seen as the investigation of the relation between functions defined on a manifold and the shape

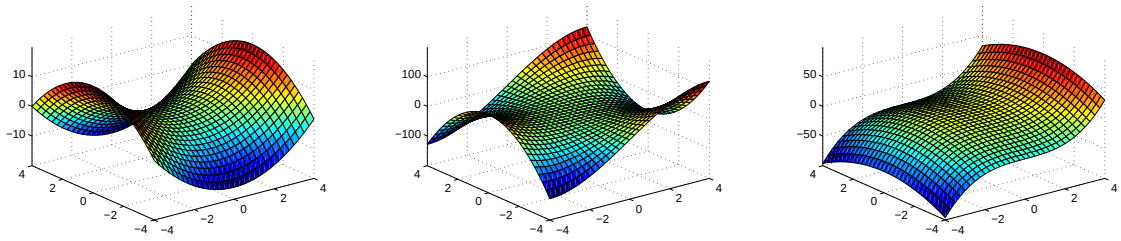

(a) (b) (c)

Figure 7: *(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.*

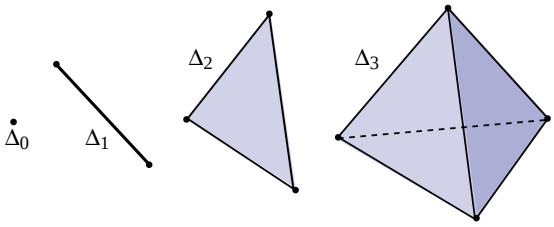

Figure 6: *Examples of 0-, 1-, 2- and 3-simplices.*

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 provides the mathematical background underlying several descriptors, such as Reeb graphs, size functions, persistence diagrams and Morse shape descriptors. A basic reference for Morse theory is [Mil63].

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 x is *non-degenerate* if the determinant of the *Hessian* matrix of f at x,

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

is not zero; otherwise the critical point is *degenerate*. Figure 7 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 $H_f(x)$ of f at x is referred to as the *index* of x. 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* $X_i = \{x \in X \,|\, f(x) \leq t_i\}$. 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]. We will discuss later on (Section 3.2.4) how this ideas have led to the introduction of geometric/topological descriptors for shape analysis.

3. Evaluating similarity between shapes

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 $d(X,Y) \geq d(X,Z) \iff \delta(X,Y) \leq \delta(X,Z), \forall\, X,Y,Z \in \mathcal{X}$.

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 function 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 that 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. $\delta(g(X),Y)) = \delta(X,Y)$).

A common strategy in shape (dis)similarity assessment is to associate the shape of an object with a compact codification of its most salient features, which is usually referred to as *a shape descriptor*. In this way, shape descriptors can be used in place of the whole model representations to derive some (dis)similarity score between the original objects. Nevertheless, a single descriptor might not be enough to get a sufficiently detailed shape characterization. Therefore, batteries of descriptors can be used separately to produce multiple (dis)similarity scores that would be merged a posteriori. This is roughly the rationale behind classical approaches to similarity assessment, and it is exemplified in Section 3.1.

More recently, new emphasis has been given to assess the dissimilarity between 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. The added value in this approach is that similarity can be expressed not only in terms of a single score, but also trough a map between shapes. Despite of the increasing computational complexity, this makes possible to derive either a sparse or a dense shape correspondence, which is particularly useful when analysing variability among shapes. *Minimum-distortion embeddings* and the *Gromov-Hausdorff* frameworks apply the above paradigm by modelling shapes as metric spaces, and measuring the metric distortion when transforming one space into the other. These frameworks are flexible to the choice of the metric, and therefore can handle different invariance requirements. Their output is typically a point-to-point correspondence between samplings of the considered metric spaces, producing a (dis)similarity measure as a by-product. *Functional maps* generalize the point-to-point shape correspondence, which can be computationally expensive for a high number of points, into that between real-valued functions defined on shapes. The main advantage is that a correspondence obtained in this way can be seen as a linear map in the space of functions: hence, a number of tools and techniques from linear algebra can be applied to couple shapes. Methods related to the *natural-pseudo distance* may take advantage of a mathematically sound notion of stability, although the final output is usually only in terms of a (dis)similarity score. Nevertheless, the intrinsic modularity of the framework allows for comparing shapes according to different notions of invariance, which are directly inherited from the functions used to describe the considered shape properties. These four frameworks are detailed in Section 3.2.

3.1. Similarity estimation through shape descriptors

While designing a shape descriptor, the first challenge is to identify the shape properties that better characterize the object under study and are highly discriminative; in all the settings discussed below, 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, rotations and translations in case of shape alignment.

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, as shown

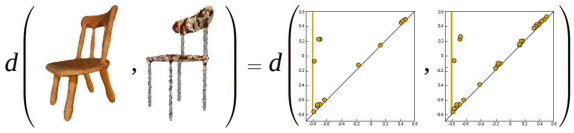

Figure 8: *The (dis)similarity between two objects is computed as a suitable distance between their descriptors.*

in Figure 8, according to a suitable (dis)similarity measure taking into account the above remarks.

The use of shape descriptors is largely acknowledged in the literature and a variety of methods has been proposed so far [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.

Remarks, examples, and applications. Among the number of methods proposed in the literature, we selected some representative ones that are meant to give an overall idea of the variety of descriptions (histograms, matrices, graphs, etc.) and type of information stored (punctual, surface or volumetric, possibly with attributes). What all these methods have in common is that they extract shape information in the form of functions, and use that information to derive shape descriptors. Similarity assessment is then performed by directly comparing descriptors through the use of suitable distances.

An example of descriptor that encodes rigid shape properties is the one proposed by Mademlis et al [MDTS09] in which the potential of a Newtonian field defined in the space outside the shape is adopted. The 3D descriptor is the combination of independent histograms (36 in the paper) related to surface proximity, field intensity and curvature while comparison is performed with ad-hoc distances. 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 while the scale-invariance is achieved through a pre-processing step in which all shapes are normalized and voxelized.

Towards intrinsic invariance, Smeets et al. in [SHVS12] adopt geodesic distances between surface samples. These values are stored in a *geodesic distance matrix* (GDM) that is used to derive further shape histograms to be compared using classical metrics (e.g. χ^2, L^p and Jensen-Shannon divergence). The method is shown to be robust under nearly isometric deformations (examples are provided on articulated objects and faces) although the use of geodesics make it sensitive to topological changes (e.g., open/close mouth, two fingers that touching each other).

The scale-invariant HKS (SI-HKS) was proposed by Bronstein et al. [BK10b] to overcome the scale dependence of the standard HKS and inherits its autodiffusion-related structure. Indeed, scale independence is done in terms of scaling and shift in time: scale is obtained from the logarithm of HKS_t and its discrete derivative with respect to time, while the shift is seen as a different phase that 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 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 different inputs such as point clouds or meshes.

Heat kernels of the Laplace-Beltrami operator are also at the basis of the approach proposed in [BBGO11], in which they are used to derive local feature descriptors for each considered shape. Then, the *bag-of-features* paradigm is deployed for aggregating local (point-wise) descriptors into global (shape-wise) one. Roughly, local features are first used to construct a vocabulary of "geometric words". Then, local shape descriptors are represented by geometric words from the above vocabulary using vector quantization, so that each shape is represented through a "bag of features" by counting the frequency of occurrence (histogram) of vocabulary words. Finally, similarity-sensitive hashing is applied to the bags of features.

Besides the use of histograms, graph-based signatures are well suited when structure and shape parts are relevant for the application. The basic assumption behind this type of methods is that a shape can be decomposed into significant parts, each one described by a local description, and that also the relation between parts is relevant. Then, the comparison is done with graph matching approaches that take into account both local geometrical features and the part hierarchies. As a representative example of this class of methods we highlight the technique for finding corresponding parts in structurally different objects proposed by Shapira et al. [SSS*10]. 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 shape diameter function (SDF) [SSCO08] makes it particularly suitable for articulated shapes.

Zaharescu et al. [ZBVH09, ZBH12] treated the case of textured 3D shapes. Their method builds upon a scale-space derived from different normalized Gaussian deriva-

tives through the Difference-of-Gaussians (DoG) operator [Low04], and incorporates in a unique paradigm geometric 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]). Invariance to the mesh sampling (i.e. the selection of the feature points) is obtained with the normalization of the histograms through L^2 norm, that is also used for their comparison. Depending on the choice of the function (mean curvature, Gaussian, etc.) isometric invariance is satisfied, while the scale-space description guarantees robustness to noise [ZBH12].

Finally, a very recent trend in shape analysis, borrowed from the image processing and computer vision community, is *learning* invariant structure rather than trying to hand-craft them. The main advantage of learning methods is that, instead of trying to model the noise or shape variability axiomatically, one learns them from examples. In particular, learning methods allow for creating class-specific descriptors that address fine-grained differences between shapes in the class. Litman and Bronstein [LB14] proposed a parametric spectral descriptor generalizing the heat and wave kernel signatures [SOG09, GBAL09, ASC11]. Masci et al. [MBBV15, BMM*15] proposed a generalization of the popular convolutional neural networks (CNN) [LBD*89] paradigm to manifolds. Litman et al. [LBBC14] proposed a supervised version of the bag-of-features framework for local descriptor aggregation, allowing to achieve state-of-the-art performance on fine-grained shape classification [PSR*14] and large-scale retrieval [SBS*15] tasks.

3.2. Frameworks for similarity assessment

Recently, the 3D shape analysis community has assisted to the emerging of new frameworks for similarity assessment. The peculiarity of these approaches is to rely on a solid mathematical basement: as a consequence, the associated computational methods may count on a number of formally proven properties involving, e.g., stability against functions (that is to say, shape properties) perturbations and invariance under group of transformations. We revise in what follows the essential traits of such frameworks.

3.2.1. Minimum-distortion embeddings

Generally speaking, approaches for shape similarity through the metric geometry framework model the shapes as metric spaces equipped with some metric. The degree of similarity is quantified as the degree of isometry of the two respective metrics, and the choice of the metric prescribes the invariance of the shape similarity. The simplest choice is the Euclidean metric arising from the space $\mathbb{R}^3$ in which the shape X is embedded, which is invariant to Euclidean isometries (the elements of the isometry group $\mathrm{Iso}(\mathbb{R}^3, d_{\mathbb{R}^3})$ are rigid motions including rotations, translations, and reflections). The metric space $(X, d_{\mathbb{R}^3})$ is a subset of the metric space $(\mathbb{R}^3, d_{\mathbb{R}^3})$. Given two shapes $(X, d_{\mathbb{R}^3})$ and $(Y, d_{\mathbb{R}^3})$ and regarding them as subsets of $(\mathbb{R}^3, d_{\mathbb{R}^3})$, their similarity can be quantified using the Hausdorff distance

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

which expresses the similarity between two subsets of a metric space with metric d [HKR93]. Since the shapes are defined up to Euclidean isometry, one minimizes d_H over all the possible rigid motions,

$$\min_{i\in \mathrm{Iso}(\mathbb{R}^3, d_{\mathbb{R}^3})} d_H(i(X), Y),$$

parametrized by a small number of degrees of freedom (three rotation angles and three translation coordinates). This optimization problem can be regarded as best possible rigid alignment of X and Y in $\mathbb{R}^3$, and is solved efficiently using iterative closest point (ICP) algorithms (provided that good initialization parameters are known) [BM92, CM92].

In a more general setting, we are given two shapes (X, d_X) and (Y, d_Y) with some generic metrics d_X, d_Y (for example, the geodesic or diffusion metrics invariant to isometric deformations of the shapes) that do not arise from a common metric space. In this case, one can either try to compare the metric directly as described in the next section, or alternatively, reduce the problem to the aforementioned setting. For this purpose, one tries to represent the metric d_X (respectively, d_Y) in some fixed metric space (Z, d_Z) by means of an isometric embedding $f : X \to Z$ (respectively, $g : Y \to Z$) satisfying $d_X = d_Z \circ (f \times f)$ (respectively, $d_Y = d_Z \circ (g \times g)$). The images $f(X)$ and $g(Y)$, referred to as *canonical forms* by Elad and Kimmel [EK03], can be compared as subsets of (Z, d_Z) using the Hausdorff distance under the isometries in (Z, d_Z),

$$\min_{i\in \mathrm{Iso}(Z, d_Z)} d_H(i(f(X)), g(Y)),$$

see figure 9 for a graphical intuition. The choice of the embedding space (Z, d_Z) should be such that its isometries can be easily parametrized and searched over. In particular, when $Z = \mathbb{R}^3$, the comparison of canonical forms boils down to the rigid alignment problem.

Unfortunately, isometric embeddings of general metrics into a Euclidean space typically do not exist. It is however possible to find the best possible approximate isometry, by minimizing some error criterion

$$\min_{f: X\to Z} \|d_x - d_Z \circ (f \times f)\|.$$

Remarks, examples, and applications. Elad and Kimmel [EK03] used the L_2 (least-squares) error, finding the approximately isometric embedding by solving the *multidimensional scaling* (MDS) problem [Bor05]. Embeddings into

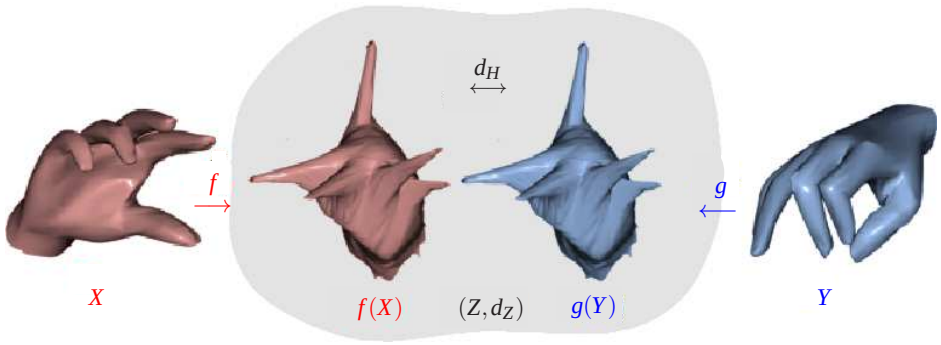

Figure 9: *The canonical forms approach: shapes (X,d_X) and (Y,d_Y) are (approximately) isometrically embedded into a fixed metric space (Z,d_Z) by means of maps f, g. The resulting canonical forms $f(X)$, $g(Y)$ are compared using the Hausdorff distance d_H over all the isometries $Iso(Z,d_Z)$.*

other metric spaces such as the sphere of hyperbolic space were studied in [BBK05, WR02]. More recently [LGX13], a feature-preserved approach has been proposed for non-rigid 3D watertight meshes. The basic idea is to consider MDS embedding results as references and then naturally deform the original meshes against them. In this way, the obtained canonical forms not only have the isometry-invariant property, but also preserve important details on the original surfaces.

It should be noted that the general idea of representing the shape in a space where matching can be described using a small number of degrees of freedom is a common trait to other shape similarity frameworks, such as the line of works of Gu et al. [JWYG04] and Lipman et al. [BLC*11, LD11, LPD13], where instead of looking for an (approximately) isometric embedding, one looks for an angle-preserving (conformal) embedding of the shape into the disk (for shapes with boundary) or sphere (for genus-0 shapes). Such conformal embeddings are known to exist by virtue of the *uniformization theorem* [Poi08]. For two isometric shapes, such conformal maps are defined up to a Möbius transformation.

3.2.2. Gromov-Hausdorff distance

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, see Figure 10. Equivalently, the computation of the Gromov-Hausdorff distance 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

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

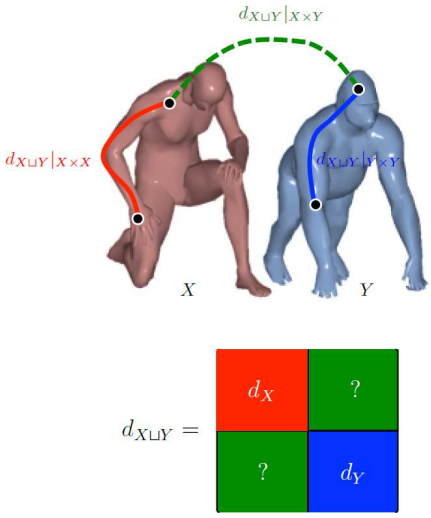

Figure 10: *Gromov-Hausdorff distance as metric coupling: illustratively, the metric $d_{X\sqcup Y}$ on the disjoint union of X and Y can be thought of as a matrix consisting of the blocks $d_{X\sqcup Y|X\times X} = d_X$ (red), $d_{X\sqcup Y|Y\times Y} = d_Y$ (blue) and the unknown blocks $d_{X\sqcup Y|X\times Y}$, $d_{X\sqcup Y|Y\times X}$ (green). Computing the Gromov-Hausdorff distance amounts to determining the green blocks that would result in the smallest d_H.*

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. As for mappings $\phi_{YX} : Y \to X$, 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 dis-

tance d_{GH} 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, M12, M11]. 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 non-rigid 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].

Remarks, examples, and applications. The Gromov-Hausdorff distance was first proposed for deformable shape analysis by Mémoli and Sapiro [MS05] with a probabilistic approximation scheme for discrete spaces. Bronstein et al. [BBK06] proposed the generalized MDS (GMDS) approach, based on a continuous optimization w.r.t. point coordinates on triangular meshes. GMDS allows for computing the distortion terms in a manner similar to MDS in Euclidean spaces, with two main differences: first, the distances have no closed-form expression but are interpolated; and second, the points are represented in local (barycentric) rather than global (Euclidean) coordinates. In [WBBP12], the Gromov-Hausdorff distance was computed in a hierarchical manner using graph labeling methods. Mémoli [M11] generalized the Gromov-Hausdorff construction to metric-measure spaces, introducing the Gromov-Wasserstein distance which generalizes the Wasserstein (earth mover's) distance [RTG00]. An extension of the Gromov-Hausdorff framework to the setting of partial shape matching was proposed in [BB08, BBBK09].

3.2.3. Functional maps

Ovsjanikov et al. [OBCS*12] proposed a functional representation of maps between shapes. Suppose we are given a bijective mapping $\tau : X \to Y$ between the shapes X and Y, and let $f \in \mathcal{F}(X)$ be some real-valued function on X, with $\mathcal{F}(X)$ denoting a space of real-valued functions on X. Then, the function on Y corresponding to f is given by $g = f \circ \tau^{-1} \in \mathcal{F}(Y)$. This correspondence can be represented by means of a linear operator $T : \mathcal{F}(X) \to \mathcal{F}(Y)$ mapping functions between shapes, such that $g = T(f)$. Foregoing the bijective correspondence, one can consider such T as a more generic functional correspondence (the point-wise bijective correspondence is a particular choice of T whereby a delta-function on X is mapped to a delta-function on Y).

Suppose now that we are given two orthonormal bases $\{\phi_i\}_{i\geq 1} \subset \mathcal{F}(X)$ and $\{\psi_j\}_{j\geq 1} \subset \mathcal{F}(Y)$ on X and Y, respectively. Any function $f \in \mathcal{F}(X)$ can be represented through the Fourier series as $f = \sum_{i\geq 1} \langle f, \phi_i \rangle_{\mathcal{F}(X)} \phi_i = \sum_{i\geq 1} a_i \phi_i$. Due to the linearity of T, we have $Tf = T(\sum_{i\geq 1} a_i \phi_i) = \sum_{i\geq 1} a_i T\phi_i$. The functions $T\phi_i \in \mathcal{F}(Y)$ can be then expanded in the basis $\{\psi_j\}_{j\geq 1}$ as $T\phi_i = \sum_{j\geq 1} \langle T\phi_i, \psi_j \rangle_{\mathcal{F}(Y)} \psi_j = \sum_{j\geq 1} c_{ji} \psi_j$, yielding

$$Tf = \sum_{i,j\geq 1} \psi_j c_{ji} \langle f, \phi_i \rangle_{\mathcal{F}(X)}.$$

This way, the functional correspondence is encoded through the coefficients c_{ji}, which determine how Fourier coefficients in the basis $\{\phi_i\}_{i\geq 1}$ are translated into Fourier coefficients in the basis $\{\psi_j\}_{j\geq 1}$. It is readily approximated by taking the first K coefficients in the expansion, $Tf = \sum_{i,j=1}^{K} \psi_j c_{ji} \langle f, \phi_i \rangle_{\mathcal{F}(X)}$.

Discretization. In the discrete setting, the shapes X and Y are represented as discrete spaces with n and m points, respectively, and the functional spaces $\mathcal{F}(X)$ and $\mathcal{F}(Y)$ can be identified with $\mathbb{R}^n$ and $\mathbb{R}^m$, respectively. The functional correspondence T is represented as an $m \times n$ matrix. Denoting by $\Phi = (\phi_1, \ldots, \phi_K)$ and $\Psi = (\psi_1, \ldots, \psi_K)$ the $n \times K$ and $m \times K$ matrices of basis vectors, respectively, we have $T \approx \Psi C^\top \Phi^\top$, where C is the $K \times K$ matrix of coefficients encoding the correspondence. Thus, finding the correspondence boils down to finding the matrix C, which, in turn, appears to be a simple algebraic problem: given a set of q corresponding functions $F = (f_1, \ldots, f_q)$ and $G = (g_1, \ldots, g_q)$ (such that $G \approx TF$), C is computed by solving the system of qK linear equations

$$G^\top \Psi = F^\top \Phi C \tag{3}$$

in K^2 variables. Assuming the columns of F and G are linearly-independent, the system has a unique solution when $q = K$, is under-determined if $q < K$ and over-determined when $q > K$. In the latter case, it is solved in the least-squares sense,

$$\min_{C \in \mathbb{R}^{K\times K}} \|G^\top \Psi - F^\top \Phi C\|_F^2. \tag{4}$$

Figure 11 shows an example of a bijective map between two nearly isometric dog shapes, and the corresponding

functional representation in the form of a 20×20 matrix C, computed according to equation (4). As shown in the Figure, functional maps between near-isometric shapes are represented by nearly-diagonal matrices in the Laplace-Beltrami eigenbases.

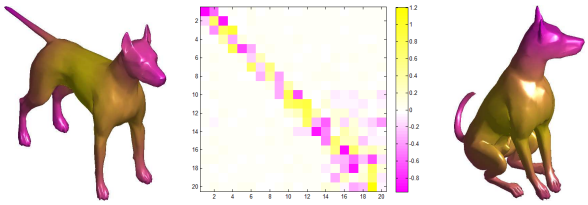

Figure 11: *Two shapes and the isometric mapping between them in the form of its functional representation.*

Choice of F and G. As corresponding functions F and G one may choose any function that can be independently computed on X and Y satisfying at least approximately $G \approx TF$. Ovsjanikov et al. [OBCS*12] used the wave kernel signatures [ASC11]; Kovnatsky et al. [KBB*13] and Pokrass et al. [PBB*13] used indicator functions of stable regions (shape MSERs [LBB11]) which turn out to be very robust. Pokrass et al. [PBB*13] also used point-wise correspondences.

Choice of the basis. It is important to note that C depends on the choice of the bases. Ovsjanikov et al. [OBCS*12] use as the basis functions $\{\phi_i\}_{i\geq 1}$ and $\{\psi_j\}_{j\geq 1}$ the first eigenfunctions of the Laplace-Beltrami operators of X and Y, respectively. Due to analogy with frequency analysis, taking the first K elements of such bases can be interpreted as "low-pass filtering", thus restricts the functional maps to a subset of smooth maps. Using more basis functions (larger K) allows more accurate maps; however, at the same time it requires more corresponding functions to be available (larger q), which in many cases might be difficult to provide.

Kovnatsky et al. [KBB*13, EKB*15] introduced the joint approximate eigenbases, computed as an orthonormal transformation of the Laplacian eigenbases of the form $\hat{\Phi} = \Phi P$ and $\hat{\Psi} = \Psi Q$, where P and Q are orthonormal matrices ($Q^\top Q = P^\top P = I$). The matrices P and Q are found to make sure that the Fourier coefficients of corresponding functions F and G in the respective bases $\hat{\Phi}$ and $\hat{\Psi}$ are approximately equal, making sure at the same time that $\hat{\Phi}$ and $\hat{\Psi}$ approximately diagonalize the respective Laplacians, by solving the optimization problem of the form

$$\min_{P,Q} \quad \|G^\top \Psi Q - F^\top \Phi P\|_F^2 + \mu_1 \mathrm{off}(P^\top \Lambda_X P) + \mu_2 \mathrm{off}(Q^\top \Lambda_Y Q)$$
$$\text{s.t. } P^\top P = Q^\top Q = I,$$

where $\mathrm{off}(A) = \sum_{i\neq j} a_{ij}^2$ is a penalty on a non-diagonal structure, and Λ_X, Λ_Y are diagonal matrices containing the first Laplacian eigenvalues. This optimization is carried out efficiently using manifold optimization techniques [KGB15]. If X and Y are nearly isometric, the correspondence matrix C represented in the joint basis is approximately diagonal, allowing to reduce the system (3) of qK equations in K^2 variables to K variables considering only the diagonal elements of C.

One of the disadvantages of Laplace-Beltrami eigenbasis is poor localization properties, making it hard, in particular, to represent correspondence between shapes with missing parts. As a remedy, Neumann et al. [NVT*14] used the recently introduced *compressed modes* [OLCO14].

Shape similarity through functional maps. The joint diagonalizability of the Laplacians under functional correspondence was used by Kovnatsky et al. [KBB*13] as a criterion for shape similarity (see Figure 12).

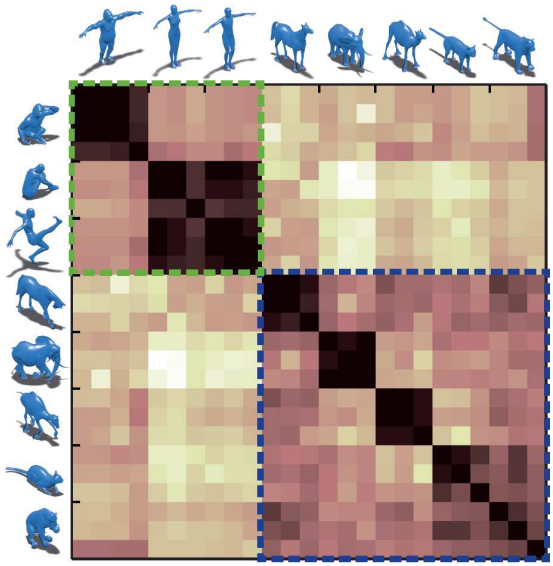

Figure 12: *Quantifying shape similarity using joint diagonalization. Darker colors represent more similar shapes. One can clearly distinguish blocks of isometric shapes. Also, two classes of two- and four-legged shapes (marked with green and blue) are visible. Small figures show representative shapes from each class.*

A framework for capturing fine-grained shape differences was proposed by Rustamov et al. [ROA*13]. The authors notice that typically $\langle f,g\rangle_{\mathcal{F}(X)} \neq \langle Tf,Tg\rangle_{\mathcal{F}(Y)}$ for some choice of the inner products on $\mathcal{F}(X)$ and $\mathcal{F}(Y)$, a pair of functions $f,g \in \mathcal{F}(X)$, and a functional map T. However, by the Riesz representation theorem, there exists a unique self-adjoint linear *shape difference operators* $D : \mathcal{F}(X) \to \mathcal{F}(X)$ "equalizing" the inner products, in the sense

$$\langle f, Dg\rangle_{\mathcal{F}(X)} = \langle Tf, Tg\rangle_{\mathcal{F}(Y)}.$$

The shape difference operator represents the change of inner product from one shape to another under the functional map

and describes local shape differences (see Figure 13). Rustamov et al. use the shape difference operator for the analysis of shape collections, taking the advantage of their linear structure, allowing to perform basic operations such as PCA on shape differences.

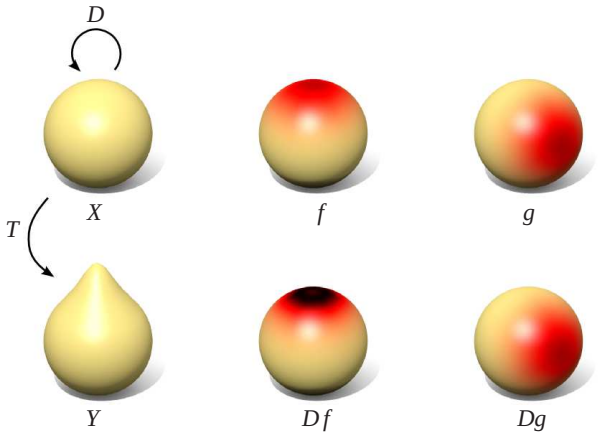

Figure 13: *Shape difference operator allows to represent the difference between two shapes under a functional correspondence.*

Remarks, examples, and applications. Conceptually, functional maps are related to *soft maps* [SNB*12], a probabilistic relaxation of point-to-point shape correspondence. Soft maps can be represented as probability matrices, thus allowing linear algebra tools for their analysis and manipulation. These concepts appears to be one of the leading directions of research in the field of shape analysis, with several important follow-up works. In [OBCCG13], the quality of functional maps is related to the singular vectors and values of the matrix C. In [PBB*13], the problem of finding the matrix C was generalized to the setting when the correspondence between the columns of F and G is unknown, leading to the permuted sparse coding problem. Extension of the functional maps to the setting of partial correspondences was proposed in [RCB*15]. Ovsjanikov et al. studied the use of functional maps in the presence of symmetries [OMPG13]. Huang et al. [HWG14a] showed the application of functional maps to the analysis of shape collections. Shapira and Ben-Chen [SBC14] showed how to use shape difference operators to find correspondence across different collections of shapes. Synthesis of shape analogies by reconstructing a shape embedding from the shape difference operator was studied in [BEKB15].

3.2.4. 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* [DF04b, DF07, DF09]. 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, g), with X and Y homeomorphic, we can imagine to transform one space into the other trough a homeomorphism $h : X \to Y$, 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 supposed to be similar with respect to certain properties if there exists a homeomorphism that preserves the functions conveying those properties.

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 might be the tangent space of S, or a projection of S onto a plane, 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

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

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 sharing the same shape properties if 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]. 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.

Lower bounds for the natural pseudo-distance. The computational issues related to the practical evaluation of the natural pseudo-distance are still an algorithmic bottleneck.

Indeed, a direct computation would require to check all possible homeomorphisms between two spaces X and Y, which is intractable in practice. For this reason, no algorithms have been proposed so far for this task, and the efforts have been rather focused on the definition of computationally efficient approximations of the natural pseudo-distance.

The above issues led to the introduction of *size functions*, shape descriptors providing a lower bound for the natural pseudo-distance [DF04a]. Interestingly, the same result can be read as a stability property under functions' perturbations [dFL10]. Size functions 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, CDF*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 modelled 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, taking inspiration from Morse theory (see Section 2.5.2). 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 14. 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*. Persistence diagrams 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].

Similarly to persistence, also *Reeb graphs* [Ree46] root in Morse theory, but track the evolution of the level sets of a function f rather than its sub-level sets [BGSF08a, DW13]. From the mathematical point of view, 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]. The parametric nature of Reeb graphs with respect to the function f is shown in Figure 14, where the Reeb graphs of a closed surface with respect to two different functions are depicted. Recently, stability results for Reeb graphs under a *functional distortion distance* [BGW14] and an *editing distance* have been proposed [DL12b, DFL14], leading to lower bounds for the natural pseudo-distance.

Thanks to their modularity, persistence diagrams and Reeb graphs provides different shape descriptions simply by changing the considered function. Interestingly, they inherit the invariance properties directly form the considered functions.

Remarks, examples, and applications. A large portion of the persistence applications proposed so far 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-Hausdorff 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. Recently, persistence diagrams have been used in combination with the bag of feature approach, to address shape retrieval and recognition tasks [LOC14].

As for Reeb graphs, they 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 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], more details can be found in the survey paper [BGSF08a]. Also, several graph matching methods have been introduced, ranging from global similarity measures [HSKK01, LMM13] to approximated subgraph matching techniques [BMSF06] and graph kernel approaches [BB13a, BB14].

(a) (b) (c) (d) (e) (f) (g)

Figure 14: *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).*

4. Taxonomy of the methods

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

Type of input. The shape representation format;

Type of invariance. The class(es) of transformations according to which the method is invariant;

Type of output. The modality used to return the shape similarity assessment: either a numerical score, or a shape correspondence, or both;

Type of structure. The particular kind of shape structure that is captured from the chosen shape description method;

Type of distance. The criterion used to assess the (dis)similarity between shape structures. It may refer to either the chosen distance between shape descriptors, or the selected framework to compare shape structures;

Computational cost. We distinguish between the extraction of shape descriptions and the subsequent comparison.

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).

In what follows, different methods are discussed on the basis of the proposed taxonomy, which is then summarized in Tables 1 and 2. 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 possible.

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 [BDF*08]. In the computer graphics community, shapes are traditionally modelled as *surface* models (two-dimensional manifolds representing boundaries of physical 3D objects) or as *volume* models. The most common discretizations of such structures are simplicial meshes (e.g. triangular or tetrahedral meshes) [RBBK10b, Rus10, BLC*11, GL12, DP13, LB14] or regular grids [BCF*08, MDTS09]. 3D grids (voxel) representations, in particular, 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 [TCL*13], which play an essential role in multi-view data fusion. In the analysis of deformable shapes, such representations are less common [MHK*08, MS09a, 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* [KBBK12, ZBH12, BCGS13].

4.2. Type of invariance

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

Intrinsic shape descriptions (such as those based on geodesics and diffusion processes) are invariant to *isometric* shape deformations, which are in general non-rigid but preserve distances computed over the shape surface [BBK09, SSS*10, WZL*10, BLC*11, BB11, BBK*10, RBB*11, BBGO11, FSR11]. Examples of isometric deformations are shape bendings. Recently, isometric deformations preserving volume have been taken onto account, as they better represent shape deformations associated with many natural phenomena [RBBK10b, Rus10, BHKH13].

Other classes of deformations which are relevant for applications include certain classes of non-isometric transformations, which do not preserve the Rienmannian structure of the shape [BCFG11]. Typical examples are shape stretching, scaling and affine transformations [RBB*11, RK14].

4.3. 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 by-product (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*: the former is a global similarity score between shapes, the latter comes from comparing only some of their parts.

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.

4.4. Type of structure

A first classification for shape structures may be done according to the associated invariance. Structures invariant to rigid transformations are referred to as *extrinsic*. Simple rotation- and translation invariant-structures include those based on Euclidean distances (e.g. from the object centre of mass [BCF*08], other points or regions of interest [BK10a, BCFG11] including 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 underlying manifold are referred to as *intrinsic*; this type of invariance is sought in applications involving deformable shapes. Intrinsic structures can be further made invariant to global scale and affine transformations [BK10b, RBB*11]. Among the possible variations in the class of intrinsic structures, we distinguish between *conformal* structures by referring to those based on Gaussian curvature and geodesic distances; *diffusion* structures for those relying on diffusion processes and spectral properties of the Laplace-Beltrami operator [DK10, 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].

Some methods allow for analysing shapes according to the topological exploration of the functions which are used to represent shape properties of interest. Referring to the emerging structures, we will talk about *topological structures*. Examples of methods dealing with this kind of information include curve skeletons [LH13], Reeb graphs [BGSF08a] and persistence-based techniques [BCF*08, CCSG*09, DLL*10, DL12a]. Some of these approaches inherit invariance properties directly from the considered functions, thus leading to intrinsic or extrinsic topological structures.

Structures can be also *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 neighbourhood. For this reason, local structures are typically used for partial similarity assessment. Global structures, on the other hand, capture the properties of the entire shape.

A local structure may be captured in the form of local descriptors. Recent works proposed a plethora of descriptors such as those based on conformal factor [BCG08], autodiffusion [GBAL09] or the heat kernel signature [SOG09] and its scale- [BK10b] and affine-invariant versions [RBB*11, RK14], and the wave kernel signature [ASC11].

Global structures can be obtained by integrating local structure over the entire shape, typically in the form of a single- or multi-dimensional histogram [BBGO11, 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 and their distributions [FSR11], as well as global spectral properties such as the Laplace-Beltrami spectrum 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.

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

Finally, in case of additional *photometric* information (texture), structure is usually captured by embedding shapes in both the Euclidean space and in a colour-based one (such as the RGB or the CIELab colour space) or possibly in a larger one somehow combining the above two. Examples of related methods include those based on colour-aware generalizations of purely geometric approaches such as heat kernel signatures [KBBK12], geodesic distance [BCGS13] and spin-images [PZC13]; other methods take inspiration from techniques initially conceived for 2D images, as in the case of [TSDS11, ZBH12] with the SIFT algorithm [Low04].

4.5. 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.

Minimum-distortion embeddings and the *Gromov-Hausdorff* frameworks allow for comparing shapes in terms of the metric distortion between two metric spaces when transforming one into the other. *Functional maps* extend the similarity problem to the comparison of functions defined on the shapes, returning a correspondence between shapes in terms of a linear map in the space of functions. Methods related to the *natural-pseudo distance* are usually endowed with a mathematically sound notion of stability, and are modular in the choice of functions and invariance properties that are 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 *feature vectors*. Such a description is based on projecting in a k-dimensional vector of the features 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 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 [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.

4.6. Computational cost

Besides accuracy, the computational complexity of a method is a fundamental aspect when dealing with real applications. In our case, extracting and comparing shape descriptions are the two computational issues that each method needs to face.

Table 2 summarizes the computational complexity of the methods listed in this survey. Since the running time complexity is not specified by several of the authors, we classify the methods on the basis of considerations about their theoretical complexity rather than a real implementation. For each method, we indicate the complexity of extracting and comparing the considered shape descriptions. In doing this, we use the term *low* to denote a computational complexity that is sub-quadratic (e.g., strictly smaller than $O(n^2)$ whit n the number of vertices in the input point cloud or triangle mesh), *high* if the complexity is higher than cubic, *medium* otherwise.

5. Applications and benchmarks

The explosive growth in the number of shape analysis techniques, including those for shape similarity assessment, has made acute the need for a widely-accepted performance evaluation protocol. This has led to the introduction of benchmarks, whose variability in the type of 3D content (medical, remote sensing, entertainment, cultural heritage etc.) and its representation (triangle mesh, volume models, point clouds etc.) in part reflects the increased availability of 3D data [GZL*14].

Beyond evaluation purposes, the objective of benchmarks is to provide environments in which methods can be tested on the cutting-edge challenges arising from real-world applications. As for shape similarity assessment, such challenges include, e.g., browsing large-scale shape collections, recognizing shapes in complex or cluttered scenes, dealing with different shape representation formats. Another issue is given by noise, usually introduced in the data acquisition and/or processing phases. We can distinguish between multiple kinds of noise, possibly affecting either the geometric or the topological shape structure, or both. We mention shortcuts, shot noise and missing parts, just to cite a few. Noise may come also in the form of illumination changes, decorative pattern degradation or material deterioration, in case additional texture information is considered.

In what follows we provide an overview of the application domains involving the process of similarity assessment, and the most popular benchmarks that have been released to train and evaluate methods in these tasks.

Applications. In the existing literature, we identify some (partially overlapping and sometimes used synonymously) specific classes of applicative tasks [FKMS05].

Shape matching is usually referred to the task of establishing a correspondence between feature points or regions of different shapes [RBB*11, ASC11, LBB11, ZBVH09]. Often, this is the result of minimizing the distortion of some shape structure, while mapping one shape to another

[KLF11, KBB*13, PBB*13, LB14]. Nevertheless, matching two shapes can be expressed in the form of a global similarity score [BLC*11, LPD13, RBBK10a, BBK*10, BCF*08], possibly (but not necessarily) obtained as the by-product of a correspondence like in the frameworks discussed throughout the paper.

Partial matching is a variant of the shape matching problem [LBZ*13, SPS14], according to which similarity assessment is restricted to shape parts, still in terms of correspondence [SHCB11, vKZH13] or numerical score [SSS*10, DLL*10, TDVC11, WZL*10]. Partial matching and the strictly related complementary matching problem are key issues, e.g., for the reassembly and geometric autocompletion of fragmented CH objects [GSP*14] and protein docking [SAHBZ08, ADPH11].

Symmetry or *self-similarity* detection can be seen as another particular case of shape matching, in which a shape is matched to itself [OMMG10, RBBK10a, LH13]. Also in this case, we can distinguish between *partial* and *full* problems.

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, focusing in particular on the computation of good alignment axes [CVB09]. More recent works considered finding differences between shapes [DP13] and non-rigid registration of deformable shapes [LZSCO09, BHKH13].

Shape retrieval refers to the task of finding the models in a database that best match a given query [TV04, TV08]. Therefore, all method whose output is a similarity score between couples of shapes, can be adopted for 3D content based retrieval [FSR11, DAL12, BB13a, LJ13, BCGS13]. A variation of the retrieval problem consists in looking for partial shape similarity, for example if some of the considered shapes present missing parts or in case part of a shape has to be considered as background and hence discarded [TVD09, DLL*10, AK11, Lav12]. This is usually referred to as *partial shape retrieval* [LBZ*13]. Following a recent trend in shape analysis, great attention is currently paid to the retrieval of deformable shapes [WLZ10, APP*10, BBGO11, STP12, LBH13, LGJ14].

Shape 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 [BDP10, SHVS12, BWdBP13] 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 [tHV10, RBB*11, SHVS12]. The performance improvement in 3D face recognition has led to the application of 3D methods for the identification of facial expressions [MAD*11, BDP13a], also in presence of partial occlusions [BDP13b] and missing data [SKVS13].

Shape classification aims at finding the class the query model belongs to [BGM*06, GM08]. It is closely related to shape retrieval as classification can be induced by a similarity score between shapes [CCSG*09, Bia10, BCA*14]. However, a recent trend is to use machine learning approaches for this task [BB13b, BEGB13].

Shape collection exploration and organization aim to deriving high-level information about objects from their relation with the other objects in the collection [OLGM11, HZG*12, KLM*12, ROA*13]. The goal is to facilitate exploration and content search so as to understand their overall categorization and summarize their content [LMS13]. The key challenge is that shapes can vary in different ways, and users may be interested in different types of variations [KLM*12, KLM*13]. When a collection possesses rich variation and highly dissimilar objects [HZG*12] a single measure is not likely to provide a good organization and quantitative measures may be unreliable, or at least not informative enough: in this case it is possible to use qualitative information derived from multiple quantitative measures or to build a network where shape correspondences are consistently preserved [HWG14b].

Additional application scenarios related to the above tasks are component detection and segmentation [LZSCO09, LBB11], *shape editing* [DP13, KBB*13], attribute transfer [OBCS*12] and semantic annotation [ARSF09, LMS13].

Benchmarks. Among the firsts 3D shape benchmarks is *SHREC* (http://www.aimatshape.net/event/SHREC) that started as a shape retrieval contest [Vt08, VTH09] and grew over the years into additional tracks covering multiple tasks such as shape correspondence [BBC*10a], feature detection [BBB*11, GMP*14] and classification [GM08, MVR*10, PSR*14]. Retrieval of textured shapes was studied in [MFP*13, BCA*14, GFF*15]. Robustness and stability issues were addressed in dedicated tracks [VGB*10, BBC*10b, BBB*12]. Scalability of algorithms was tested in [SBS*15]. Performance of different methods in fine-grained classification settings was assessed in [PSR*14].

Other popular benchmarks devoted to 3D shape similarity are the *Princeton Shape Benchmark*, conceived at first for shape retrieval and now extended to shape correspondence and segmentation [SMKF04, CGF09]; the *McGill 3D Shape Benchmark*, in part building on the Princeton Shape Benchmark but specifically designed for non-rigid shape retrieval [SZM*08]; the *Toyohashi Shape Benchmark* [TKA12], consisting of 10,000 3D models grouped in 352 classes; the NTU 3D model database [CTSO03] consisting of 10,911 models that can be used to create a 3D shape search engine; the 3D architecture shape benchmark (ASB) [WBK09] made of 2257 objects classified in 42 classes and specifically designed for architectural 3D models.

In some cases, benchmarks offer also a number of plot-based and scalar-based measures to asses the performance of the methods, such as [SMKF04, Vt08]. An on-line content-based retrieval evaluation tool (RETRIEVAL 3D [KIPC15]) has been recently offered and is accessible through a dynamic visualisation environment.

Besides benchmarks covering the need of general-purpose test-beds for different aspects of similarity assessment, it is worth to mention datasets targeting specific real-word application scenarios, such as object recognition (e.g., PASCAL [XMS14]), face recognition (e.g., the 3D face database at the University of York[§] and the Bosphorus dataset [SAD*08]), RGB-D vision (e.g., the NYU Dept. [SHKF12], the BigBIRD: (Big) Berkeley Instance Recognition Dataset [SSN*14]) and, in case of LIDAR data acquired with vehicle and airborne laser scanner, remote sensing (e.g., the ISPRS Test Project on Urban Classification and 3D Building Reconstruction[¶] and the iQmulus Processing Contest[‖]).

6. Discussions

The ensemble of the reviewed approaches highlights that in the last years (since 2008) rigid-invariant comparison has come to be considered a quite well established problem. Indeed, the few methods tackling this problem are mainly devoted to the improvement of the accuracy and the computational efficiency of existing methods. In the meanwhile, point clouds have decreased their popularity because strictly related to rigid registration and alignment problems and there is a growing availability of 3D mesh generation algorithms to convert point clouds to meshes, e.g., [MRB09]. Conversely, analysing and comparing non-rigid shapes currently emerge as extremely challenging and addressed issues (involving almost the 90% of the reviewed methods). In this context, defining a suitable notion of similarity appears to be more complicated as the variety of the deformations evolve from rigid to non-rigid. For instance, the classic metric paradigm becomes less effective, and non metric distances come into the play.

In this scenario, we have focused our attention on methods for 3D similarity assessment that abstract the shape properties of interest as functions. We have identified four computational frameworks able to measure the variation of the considered functions, i.e. shape properties, while transforming one shape into another. Being rooted in well-established mathematical theories, these techniques take advantage of results on stability against functions' perturbation and invariance under groups of transformations. Scalability of the methods is still a computational bottleneck of many techniques.

Among all possible non-rigid deformations, the efforts have been mainly focused on (almost) isometric ones: 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 these case, we see the role of functions as crucial to convey the shape properties that one wishes to take into account; however, while results and algorithms are well-established for real-valued functions, it not always the same for the use of multi-variate ones.

In this context, we see the theoretical framework offered by the natural pseudo-distance as an option to overcome the limitations of traditional methods relying on assumptions of rigidity, isometry, or 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.

Finally, there is a need for systematic comparative studies, since several state-of-the-art methods have not yet been compared with each other or their performances are exhibited on different benchmarks; far of thinking there exists the best method for all applications, we hope that further efforts will be devoted to the clarification and evaluation of the performance of the methods with a specific target to application domains.

Future challenges. We conclude this survey by listing a series of topics deserving, in our opinion, further research and efforts:

First, the increasing complexity of the deformations considered in applications suggests the need of more general techniques that can deal with larger groups of transformations. In spite of some progress in this sense [RK14], there is probably a long road ahead. For example, there is still no solid frameworks to compare shapes with different structure or topology but same functionality; or to formulate hybrid methods combining multiple methodological paradigms. e.g. local and global descriptors [LLL*15], or view-based and spectral-based paradigms [LGJ14].

Second, shape descriptions and representations are becoming more and more sophisticated. This is due to a deeper knowledge about properties and physical phenomena that are meaningful for shape characterization. For example, we are witnessing the consolidation of methods based on (auto)diffusion processes [GBAL09, SOG09], quantum mechanics [ASC11] and gravitational laws [MDTS09], as well as the introduction of functions encoding shape properties [SSCO08, KBBK12, GL12] never considered before. In this context, we wish for the development of "shape filtering" theory, according to which suitable functions concur to model shape properties of interest, playing at the

[§] http://www-users.cs.york.ac.uk/∼nep/research/3Dface/tomh/3DFaceDatabase.html
[¶] http://www2.isprs.org/commissions/comm3/wg4/results.html
[‖] http://www.isprs-geospatialweek2015.org/workshops/geobigdata/iqpc.html

some time the role of manifold harmonics in geometric filtering [VL08], that is, enabling a progressive, "property-based" shape representation.

Third, we are now aware that shape characterization benefits from the analysis of semantics and functionality [MS09b, CMSF11, LMS13, ZCOAM14, FAvK*14]. A major issue is then how to model such a knowledge and effectively embed it in the similarity evaluation pipeline. Here, an open and acute issue for research is how to directly encapsulate the implicit knowledge contained in a family of shapes into the shape description. In general, the use of prior knowledge might be inevitable to advance in the direction of bridging the gap between geometrical attributes and high-level semantic information. A possible way to automatically infer knowledge is to introduce in the loop statistical methods such as learning techniques [DAL12,LBBC14] that have achieved state-of-the-art performance in computer vision applications [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, BB13b, BSF13, TDVC13] or design class- or application-specific shape descriptors [LB14, MBBV15, BMM*15]. Another solution might be to consider meta-representations able to characterize the configurations that are common across a family of 3D objects and to consider this knowledge when storing the arrangements among shape parts [FAvK*14] thus making possible to explore multiple shape configurations in parallel and to collectively edit sets of shapes.

Finally, a critical issue 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 applications. For example, a difficulty we see in this step is that methods for capturing 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. A major issue is then how to adapt these approaches to deal with large scale data.

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, and 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, respectively.

References

[ADPH11] AXENOPOULOS A., DARAS P., PAPADOPOULOS G., HOUSTIS E. N.: A shape descriptor for fast complementarity matching in molecular docking. *IEEE-ACM T. Comput. Bi. 8*, 6 (2011), 1441–1457. 20

[AK11] AREEVIJIT W., KANONGCHAIYOS P.: Reeb graph based partial shape retrieval for non-rigid 3D objects. In *10th International Conference on Virtual Reality Continuum and Its Applications in Industry* (2011), ACM, pp. 573–576. 19, 20, 30, 31

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

[APP*10] AGATHOS A., PRATIKAKIS I., PAPADAKIS P., PERANTONIS S., AZARIADIS P., SAPIDIS N.: 3D articulated object retrieval using a graph-based representation. *Visual Comput. 26*, 10 (2010), 1301–1319. 20, 30, 31

[ARSF09] ATTENE M., ROBBIANO F., SPAGNUOLO M., FALCIDIENO B.: Characterization of 3D shape parts for semantic annotation. *Comput. Aided Design 41*, 10 (2009), 756–763. 20

[ASC11] AUBRY M., SCHLICKEWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops, IEEE International Conference on* (2011), pp. 1626–1633. 11, 14, 18, 19, 21, 30, 31

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

[BB08] BRONSTEIN A., BRONSTEIN M.: Not only size matters: Regularized partial matching of nonrigid shapes. In *Computer Vision and Pattern Recognition Workshops, IEEE Computer Society Conference on* (2008), pp. 1–6. 13

[BB11] BRONSTEIN M., BRONSTEIN A.: Shape recognition with spectral distances. *IEEE Trans. Pattern Anal. Mach. Intell. 33*, 5 (2011), 1065–1071. 17, 30, 31

[BB13a] BARRA V., BIASOTTI S.: 3D shape retrieval using Kernels on Extended Reeb Graphs. *Pattern Recogn. 46*, 11 (2013), 2985 – 2999. 16, 18, 19, 20

[BB13b] BARRA V., BIASOTTI S.: Learning Kernels on Extended Reeb Graphs for 3D Shape Classification and Retrieval. In *Eurographics Workshop on 3D Object Retrieval* (2013), The Eurographics Association, pp. 25–32. 19, 20, 22, 30, 31

[BB14] BARRA V., BIASOTTI S.: 3d shape retrieval and classification using multiple kernel learning on extended reeb graphs. *Visual Comput.* (2014), 1–13. 16

[BBB*11] BOYER E., BRONSTEIN A. M., BRONSTEIN M. M., BUSTOS B., DAROM T., HORAUD R., HOTZ I., KELLER Y., KEUSTERMANS J., KOVNATSKY A., LITMAN R., REININGHAUS J., SIPIRAN I., SMEETS D., SUETENS P., VANDERMEULEN D., ZAHARESCU A., ZOBEL V.: SHREC 2011: Robust Feature Detection and Description Benchmark. In *Eurographics Workshop on 3D Object Retrieval* (2011), The Eurographics Association, pp. 71–78. 20

[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 *Eurographics Workshop on 3D Object Retrieval* (2012), The Eurographics Association, pp. 101–107. 20

[BBBK09] BRONSTEIN A. M., BRONSTEIN M. M., BRUCKSTEIN A. M., KIMMEL R.: Partial similarity of objects, or how to compare a centaur to a horse. *Int. J. Comput. Vision 84*, 2 (2009), 163–183. 13

[BBC*10a] BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U., DUBROVINA A., GUIBAS L. J., HORAUD R. P., KIMMEL R., KNOSSOW D., VON LAVANTE E., MATEUS D., OVSJANIKOV M., SHARMA A.: SHREC'10 Track: Correspondence Finding. In *Eurographics Workshop on 3D Object Retrieval* (2010), The Eurographics Association, pp. 87–91. 20

[BBC*10b] 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'10 Track: Robust Shape Retrieval. In *Eurographics Workshop on 3D Object Retrieval* (2010), The Eurographics Association, pp. 71–78. 20

[BBGO11] BRONSTEIN A. M., BRONSTEIN M. M., GUIBAS L. J., OVSJANIKOV M.: Shape Google: Geometric words and expressions for invariant shape retrieval. *ACM T. Graphic. 30*, 1 (2011), 1:1–1:20. 10, 17, 18, 20, 30, 31

[BBK05] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Expression-invariant face recognition via spherical embedding. In *Image Processing, IEEE International Conference on* (2005), vol. 3, pp. III–756–9. 12

[BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proc. Natl. Acad. Sci. U.S.A. 103*, 5 (2006), 1168–1172. 13

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

[BBK09] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Topology-invariant similarity of nonrigid shapes. *Int. J. Comput. Vision 81*, 3 (2009), 281–301. 17, 30, 31

[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. Comput. Vision 89*, 2-3 (2010), 266–286. 17, 18, 20, 30, 31

[BCA*14] BIASOTTI S., CERRI A., ABDELRAHMAN M., AONO M., HAMZA A. B., EL-MELEGY M., FARAG A., GARRO V., GIACHETTI A., GIORGI D., GODIL A., LI C., LIU Y.-J., MARTONO H. Y., SANADA C., TATSUMA A., VELASCO-FORERO S., XU C.-X.: Retrieval and Classification on Textured 3D Models. In *Eurographics Workshop on 3D Object Retrieval* (2014), The Eurographics Association, pp. 111–120. 5, 20

[BCB12] BRONSTEIN M., CASTELLANI U., BRONSTEIN A.: Diffusion Geometry in Shape Analysis. In *Eurographics Tutorials* (2012), The Eurographics Association. 3

[BCF*08] BIASOTTI S., CERRI A., FROSINI P., GIORGI D., LANDI C.: Multidimensional size functions for shape comparison. *J. Math. Imaging Vis. 32* (2008), 161–179. 17, 18, 20, 30, 31

[BCFG11] BIASOTTI S., CERRI A., FROSINI P., GIORGI D.: A new algorithm for computing the 2-dimensional matching distance between size functions. *Pattern Recogn. Lett. 32*, 14 (2011), 1735–1746. 17, 18

[BCG08] BEN-CHEN M., GOTSMAN C.: Characterizing shape using conformal factors. In *Eurographics Conference on 3D Object Retrieval* (2008), The Eurographics Association, pp. 1–8. 18

[BCGS13] BIASOTTI S., CERRI A., GIORGI D., SPAGNUOLO M.: PHOG: Photometric and Geometric Functions for Textured Shape Retrieval. *Comput. Graph. Forum 32*, 5 (2013), 13–22. 17, 18, 20, 30, 31

[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 Comput. Surv. 40*, 4 (2008), 1–87. 3, 9, 10, 17

[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. 20, 30, 31

[BDP13a] BERRETTI S., DEL BIMBO A., PALA P.: Automatic facial expression recognition in real-time from dynamic sequences of 3D face scans. *Visual Comput. 29*, 12 (2013), 1333–1350. 20, 30, 31

[BDP13b] BERRETTI S., DEL BIMBO A., PALA P.: Sparse matching of salient facial curves for recognition of 3-d faces with missing parts. *IEEE T. Inf. Foren. Sec. 8*, 2 (2013), 374–389. 20, 30, 31

[BEGB13] BONEV B., ESCOLANO F., GIORGI D., BIASOTTI S.: Information-theoretic selection of high-dimensional spectral features for structural recognition. *CVIU 117*, 3 (2013), 214–228. 20, 30, 31

[BEKB15] BOSCAINI D., EYNARD D., KOUROUNIS D., BRONSTEIN M. M.: Shape-from-operator: recovering shapes from intrinsic operators. *Comput. Graph. Forum 34*, 2 (2015). 15

[BFF*07] BIASOTTI S., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., MARINI S., PATANÉ G., SPAGNUOLO M.: 3D Shape Description and Matching Based on Properties of Real Functions . In *Eurographics Tutorials* (2007), The Eurographics Association, pp. 949–998. 3

[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 Courses* (2012), ACM, pp. 17:1–17:33. 3

[BGM*06] BIASOTTI S., GIORGI D., MARINI S., SPAGNUOLO M., FALCIDIENO B.: A Comparison Framework for 3D Object Classification Methods. In *Multimedia Content Representation, Classification and Security* (2006), vol. 4105 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 314–321. 20

[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. 16, 18

[BGSF08b] BIASOTTI S., GIORGI D., SPAGNUOLO M., FALCIDIENO B.: Size functions for comparing 3D models. *Pattern Recogn. 41*, 9 (2008), 2855–2873. 16

[BGW14] BAUER U., GE X., WANG Y.: Measuring distance between reeb graphs. In *Symposium on Computational Geometry* (2014), ACM, pp. 464:464–464:473. 16

[BHKH13] BUDD C., HUANG P., KLAUDINY M., HILTON A.: Global non-rigid alignment of surface sequences. *Int. J. Comput. Vision 102*, 1-3 (2013), 256–270. 17, 20

[Bia10] BIASOTTI S.: Shape comparison through mutual distances of real functions. In *ACM Workshop on 3D object retrieval* (2010), ACM, pp. 33–38. 20, 30, 31

[BK10a] BALOCH S., KRIM H.: Object recognition through topo-geometric shape models using error-tolerant subgraph isomorphisms. *IEEE T. Image Process. 19*, 5 (2010), 1191–1200. 17, 18, 19, 30, 31

[BK10b] BRONSTEIN M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Computer Vision and Pattern Recognition, IEEE Conference on* (2010), pp. 1704–1711. 10, 18, 30, 31

[BKS*05] BUSTOS B., KEIM D. A., SAUPE D., SCHRECK T., VRANIĆ D. V.: Feature-based similarity search in 3D object databases. *ACM Comput. Surv. 37*, 4 (2005), 345–387. 2, 10, 19

[BKSS07] BUSTOS B., KEIM D. A., SAUPE D., SCHRECK T.: Content-based 3D object retrieval. *IEEE Comput. Graph. 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. *Proc. Natl. Acad. Sci. U.S.A.* (2011). 12, 17, 20, 30, 31

[BM92] BESL P., MCKAY N. D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell. 14*, 2 (Feb 1992), 239–256. 11

[BMM*03] BIASOTTI S., MARINI S., MORTARA M., PATANÈ G., SPAGNUOLO M., FALCIDIENO B.: 3D shape matching through topological structures. *Lecture Notes in Computer Science 2886* (2003), 194–203. 16

[BMM*15] BOSCAINI D., MASCI J., MELZI S., BRONSTEIN M. M., CASTELLANI U., VANDERGHEYNST P.: Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *Comput. Graph. Forum*, 35 (2015). 11, 22

[BMSF06] BIASOTTI S., MARINI M., SPAGNUOLO M., FALCIDIENO B.: Sub-part correspondence by structural descriptors of 3D shapes. *Comput. Aided Design 38*, 9 (2006), 1002–1019. 16, 19

[Bor05] BORG I.: *Modern Multidimensional Scaling: Theory and Applications*. Springer, New York, NY, 2005. 11

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

[BSF13] BIASOTTI S., SPAGNUOLO M., FALCIDIENO B.: Grouping real functions defined on 3D surfaces. *Comput. Graph. 37*, 6 (2013), 608 – 619. 22

[BSW08] BELKIN M., SUN J., WANG Y.: Discrete Laplace operator for meshed surfaces. In *Symposium on Computational Geometry* (2008), ACM, pp. 278 – 287. 6

[BWdBP13] BERRETTI S., WERGHI N., DEL BIMBO A., PALA P.: Geometric histograms of 3d keypoints for face identification with missing parts. In *Eurographics Workshop on 3D Object Retrieval* (2013), The Eurographics Association, pp. 57–64. 20, 30, 31

[CCFM08] CASTELLANI U., CRISTANI M., FANTONI S., MURINO V.: Sparse points matching by combining 3D mesh saliency with statistical descriptors. *Comput. Graph. Forum 27*, 2 (2008), 643–652. 11, 30, 31

[CCSG*09] CHAZAL F., COHEN-STEINER D., GUIBAS L. J., MÉMOLI F., OUDOT S.: Gromov-Hausdorff stable signatures for shapes using persistence. *Comput. Graph. Forum 28*, 5 (2009), 1393–1403. 16, 18, 20, 30, 31

[CDF*13] CERRI A., DI FABIO B., FERRI M., FROSINI P., LANDI C.: Betti numbers in multidimensional persistent homology are stable functions. *Math. Method. Appl. Sci. 36*, 12 (2013), 1543–1557. 16

[CFG06] CERRI A., FERRI M., GIORGI D.: Retrieval of trademark images by means of size functions. *Graph. Models 68*, 5 (2006), 451–471. 15

[CGF09] CHEN X., GOLOVINSKIY A., FUNKHOUSER T.: A benchmark for 3D mesh segmentation. *ACM T. Graphic. 28*, 3 (2009). 20

[CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *Appl. Comput. Harmon. A. 21*, 1 (2006), 5 – 30. 6

[CM92] CHEN Y., MEDIONI G.: Object modelling by registration of multiple range images. *Image Vision Comput. 10*, 3 (1992), 145–155. 11

[CMS12] CIRESAN D., MEIER U., SCHMIDHUBER J.: Multi-column deep neural networks for image classification. In *Computer Vision and Pattern Recognition, IEEE Conference on* (2012), pp. 3642–3649. 22

[CMSF11] CATALANO C. E., MORTARA M., SPAGNUOLO M., FALCIDIENO B.: Semantics and 3D media: Current issues and perspectives. *Comput. Graph. 35*, 4 (2011), 869–877. 22

[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. 16

[CTSO03] CHEN D.-Y., TIAN X.-P., SHEN Y.-T., OUHYOUNG M.: On visual similarity based 3d model retrieval. *Comput. Graph. Forum 22*, 3 (2003), 223–232. 20

[CVB09] CHAOUCH M., VERROUST-BLONDET A.: Alignment of 3d models. *Graph. Models 71*, 2 (2009), 63 – 76. 20

[DAL12] DARAS P., AXENOPOULOS A., LITOS G.: Investigating the effects of multiple factors towards more accurate 3-d object retrieval. *IEEE T. Multimedia 14*, 2 (2012), 374–388. 20, 22

[DD09] DEZA M. M., DEZA E.: *Encyclopedia of Distances*. Springer Berlin Heidelberg, 2009. 19

[DF04a] DONATINI P., FROSINI P.: Lower bounds for natural pseudodistances via size functions. *Arch. Inequal. Appl. 2* (2004), 1–12. 16

[DF04b] DONATINI P., FROSINI P.: Natural pseudodistances between closed manifolds. *Forum Math. 16*, 5 (2004), 695–715. 2, 15

[DF07] DONATINI P., FROSINI P.: Natural pseudodistances between closed surfaces. *J. Eur. Math. Soc. 9*, 2 (2007), 231–253. 15

[DF09] DONATINI P., FROSINI P.: Natural pseudo-distances between closed curves. *Forum Math. 21*, 6 (2009), 981–999. 15

[dFL10] D'AMICO M., FROSINI P., LANDI C.: Natural pseudo-distance and optimal matching between reduced size functions. *Acta Appl. Math. 109*, 2 (2010), 527–554. 16

[DFL14] DI FABIO B., LANDI C.: Stable shape comparison of surfaces via reeb graphs. In *Discrete Geometry for Computer Imagery*, vol. 8668 of *Lecture Notes in Computer Science*. Springer International Publishing, 2014, pp. 202–213. 16

[DFP04] DIBOS F., FROSINI P., PASQUIGNON D.: The use of size functions for comparison of shapes through differential invariants. *J. Math. Imaging Vis. 21* (2004), 107–118. 15

[DK10] DUBROVINA A., KIMMEL R.: Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *International Symposium on 3D Data Processing Visualization and Transmission* (2010). 18, 30, 31

[DK12] DAROM T., KELLER Y.: Scale-invariant features for 3-d mesh models. *IEEE T. Image Process. 21*, 5 (2012), 2758–2769. 30, 31

[DL12a] DI FABIO B., LANDI C.: Persistent homology and partial similarity of shapes. *Pattern Recogn. Lett. 33*, 11 (2012), 1445 – 1450. 18, 30, 31

[DL12b] DI FABIO B., LANDI C.: Reeb graphs of curves are stable under function perturbations. *Math. Method. Appl. Sci. 35*, 12 (2012), 1456–1471. 16

[DLL*10] DEY T., LI K., LUO C., RANJAN P., SAFA I., WANG Y.: Persistent heat signature for pose-oblivious matching of incomplete models. *Comput. Graph. Forum 29*, 5 (2010), 1545–1554. 15, 18, 20, 30, 31

[DP06] DEL BIMBO A., PALA P.: Content-based retrieval of 3D models. *ACM T. Multim. Comput. 2*, 1 (2006), 20–43. 2, 10

[DP13] DENNING J. D., PELLACINI F.: Meshgit: Diffing and merging meshes for polygonal modeling. *ACM T. Graphic. 32*, 4 (2013), 35:1–35:10. 17, 20

[DW13] DEY T., WANG Y.: Reeb graphs: Approximation and persistence. *Discrete Comput. Geom. 49*, 1 (2013), 46–73. 16

[EH10] EDELSBRUNNER H., HARER J.: *Computational Topology: An Introduction*. American Mathematical Society, 2010. 16

[EHB13] ESCOLANO F., HANCOCK E., BIASOTTI S.: Complexity Fusion for Indexing Reeb Digraphs. In *Computer Analysis of Images and Patterns* (2013), vol. 8047 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 120–127. 30, 31

[EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Anal. Mach. Intell. 25*, 10 (2003), 1285–1295. 2, 11

[EKB*15] EYNARD D., KOVNATSKY A., BRONSTEIN M. M., GLASHOFF K., BRONSTEIN A. M.: Multimodal manifold analysis using simultaneous diagonalization of laplacians. *IEEE Trans. Pattern Anal. Mach. Intell.* (2015). 14

[ELZ02] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. *Discrete Comput. Geom. 28* (2002), 511–533. 5, 16

[FAvK*14] FISH N., AVERKIOU M., VAN KAICK O., SORKINE-HORNUNG O., COHEN-OR D., MITRA N. J.: Meta-representation of shape families. *ACM T. Graphic. 33*, 4 (2014), 34:1–34:11. 22

[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, 19

[FM99] FROSINI P., MULAZZANI M.: Size homotopy groups for computation of natural size distances. *Bull. Belg. Math. Soc. Simon Stevin 6* (1999), 455–464. 5

[FS10] FERRI M., STANGANELLI I.: Size functions for the morphological analysis of melanocytic lesions. *Int. J. Biomed. Imaging* (2010), Article ID 621357. 15

[FSR11] FANG Y., SUN M., RAMANI K.: Temperature distribution descriptor for robust 3d shape retrieval. In *Computer Vision and Pattern Recognition Workshops, IEEE Computer Society Conference on* (2011), pp. 9–16. 17, 18, 20

[GAP*09] GIORGI D., ATTENE M., PATANE 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 *Advanced Video and Signal Based Surveillance, IEEE International Conference on* (2009), pp. 79–84. 20

[GBAL09] GEBAL K., BAERENTZEN J. A., AANAES H., LARSEN R.: Shape Analysis Using the Auto Diffusion Function. *Comput. Graph. Forum* (2009). 6, 7, 11, 18, 21

[GDP*05] GRINSPUN E., DESBRUN M., POLTHIER K., SCHRÖDER P., STERN A.: Discrete differential geometry: An applied introduction. In *ACM SIGGRAPH Courses* (2005), ACM. 6

[GDZ10] GAO Y., DAI Q., ZHANG N.-Y.: 3D model comparison using spatial structure circular descriptor. *Pattern Recogn. 43*, 3 (2010), 1142 – 1151. 17, 19, 30, 31

[GFF*15] GIACHETTI A., FARINA F., FORNASA F., TATSUMA A., SANADA C., AONO M., BIASOTTI S., CERRI A., CHOI S.: Retrieval of non-rigid (textured) shapes using low quality 3d models. In *Eurographics Workshop on 3D Object Retrieval* (2015), The Eurographics Association, pp. 137–144. 20

[GL12] GIACHETTI A., LOVATO C.: Radial symmetry detection and shape characterization with the multiscale area projection transform. *Comput. Graph. Forum 31*, 5 (2012), 1669–1678. 17, 21, 30, 31

[GM08] GIORGI D., MARINI S.: Shape retrieval contest 2008: Classification of watertight models. In *Shape Modeling and Applications, IEEE International Conference on* (2008), pp. 219–220. 1, 2, 20

[GMP*14] GIACHETTI A., MAZZI E., PISCITELLI F., AONO M., HAMZA A. B., BONIS T., CLAES P., GODIL A., LI C., OVSJANIKOV M., PATRAUCEAN V., SHU C., SNYDERS J., SUETENS P., TATSUMA A., VANDERMEULEN D., WUHRER S., XI P.: Automatic Location of Landmarks used in Manual Anthropometry. In *Eurographics Workshop on 3D Object Retrieval* (2014), The Eurographics Association. 20

[Gro99] GROMOV M.: *Metric Structures for Riemannian and Non-Riemannian Spaces Couverture*. Springer, 1999. 13

[GSP*14] GREGOR R., SIPIRAN I., PAPAIOANNOU G., SCHRECK T., ANDREADIS A., MAVRIDIS P.: Towards Automated 3D Reconstruction of Defective Cultural Heritage Objects. In *Eurographics Workshop on Graphics and Cultural Heritage* (2014), The Eurographics Association, pp. 135–144. 20

[GZL*14] GUO Y., ZHANG J., LU M., WAN J., MA Y.: Benchmark datasets for 3d computer vision. In *Industrial Electronics and Applications, IEEE Conference on* (2014), pp. 1846–1851. 19

[HAvL05] HEIN M., AUDIBERT J.-Y., VON LUXBURG U.: From graphs to manifolds - weak and strong pointwise consistency of graph laplacians. In *Conference on Learning Theory* (2005), 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. *Visual Comput. 25*, 5-7 (2009), 667–675. 30, 31

[HHZQ12] HOU T., HOU X., ZHONG M., QIN H.: Bag-of-feature-graphs: A new paradigm for non-rigid shape retrieval. In *Pattern Recognition, International Conference on* (2012), pp. 1513–1516. 30, 31

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

[HKR93] HUTTENLOCHER D. P., KLANDERMAN G. A., RUCKLIDGE W. A.: Comparing images using the Hausdorff distance. *IEEE Trans. Pattern Anal. Mach. Intell. 15*, 9 (1993), 850–863. 11

[HSKK01] HILAGA M., SHINAGAWA Y., KOHMURA T., KUNII T. L.: Topology matching for fully automatic similarity estimation of 3D shapes. In *Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2001), ACM, pp. 203–212. 16

[HWG14a] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM T. Graphic. 33*, 4 (2014), 36. 15

[HWG14b] HUANG Q., WANG F., GUIBAS L.: Functional map

networks for analyzing and exploring large shape collections. *ACM T. Graphic. 33*, 4 (2014), 36:1–36:11. 20

[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 T. Graphic. 31*, 6 (2012), 167:1–167:11. 20, 30, 31

[JWYG04] JIN M., WANG Y., YAU S.-T., GU X.: Optimal global conformal surface parameterization. In *Visualization, 2004. IEEE* (2004), pp. 267–274. 12

[KBB*12] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., RAVIV D., KIMMEL R.: Affine-invariant photometric heat kernel signatures. In *Eurographics Workshop on 3D Object Retrieval* (2012), The Eurographics Association, pp. 39–46. 30, 31

[KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Comput. Graph. Forum 32* (2013), 439–448. 14, 20, 30, 31

[KBBK12] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R.: Photometric heat kernel signatures. In *International Conference on Scale Space and Variational Methods in Computer Vision* (2012), Springer-Verlag, pp. 616–627. 17, 18, 21, 30, 31

[KGB15] KOVNATSKY A., GLASHOFF K., BRONSTEIN M. M.: MADMM: a generic algorithm for non-smooth optimization on manifolds. *ArXiv*, 1505.07676 (2015). 14

[KIPC15] KOUTSOUDIS A., IOANNAKIS G., PRATIKAKIS I., CHAMZAS C.: RETRIEVAL 3D: An On-line Content-Based Retrieval Performance Evaluation Tool. In *Eurographics Workshop on 3D Object Retrieval* (2015), The Eurographics Association. 21

[KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM T. Graphic. 30*, 4 (2011), 79:1–79:12. 20, 30, 31

[KLM*12] KIM V. G., LI W., MITRA N. J., DIVERDI S., FUNKHOUSER T.: Exploring collections of 3D models using fuzzy correspondences. *ACM T. Graphic. 31*, 4 (2012), 54:1–54:11. 18, 20

[KLM*13] KIM V. G., LI W., MITRA N. J., CHAUDHURI S., DIVERDI S., FUNKHOUSER T.: Learning part-based templates from large collections of 3d shapes. *ACM T. Graphic. 32*, 4 (2013), 70:1–70:12. 20

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

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

[Lag10] LAGA H.: Semantics-Driven Approach for Automatic Selection of Best Views of 3D Shapes. In *Eurographics Workshop on 3D Object Retrieval* (2010), The Eurographics Association. 22, 30, 31

[Lav12] LAVOUÉ G.: Combination of bag-of-words descriptors for robust partial shape retrieval. *Visual Comput. 28*, 9 (2012), 931–942. 20, 30, 31

[LB14] LITMAN R., BRONSTEIN A.: Learning spectral descriptors for deformable shape correspondence. *IEEE Trans. Pattern Anal. Mach. Intell. 36*, 1 (2014), 171–180. 11, 17, 20, 22, 30, 31

[LBB11] LITMAN R., BRONSTEIN A. M., BRONSTEIN M. M.: Diffusion-geometric maximally stable component detection in deformable shapes. *Comput. Graph. 35*, 3 (2011), 549–560. 14, 18, 19, 20, 30, 31

[LBBC14] LITMAN R., BRONSTEIN A., BRONSTEIN M., CASTELLANI U.: Supervised learning of bag-of-features shape descriptors using sparse coding. *Comput. Graph. Forum 33*, 5 (2014), 127–136. 11, 22

[LBD*89] LECUN Y., BOSER B., DENKER J. S., HENDERSON D., HOWARD R. E., HUBBARD W., JACKEL L. D.: Backpropagation applied to handwritten zip code recognition. *Neural comput. 1*, 4 (1989), 541–551. 11

[LBH13] LI C., BEN HAMZA A.: A multiresolution descriptor for deformable 3D shape retrieval. *Visual Comput. 29*, 6-8 (2013), 513–524. 20, 30, 31

[LBZ*13] LIU Z., BU S., ZHOU K., GAO S., HAN J., WU J.: A survey on partial retrieval of 3D shapes. *J. Comput. Sci. Technol. 28*, 5 (2013), 836–851. 2, 20

[LD11] LIPMAN Y., DAUBECHIES I.: Conformal Wasserstein distances: Comparing surfaces in polynomial time. *Adv. Math. 227*, 3 (2011), 1047 – 1077. 12

[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 (2013), 449–461. 2

[LGJ14] LI B., GODIL A., JOHAN H.: Hybrid shape descriptor and meta similarity generation for non-rigid and partial 3d model retrieval. *Multimed. Tools Appl. 72*, 2 (2014), 1531–1560. 20, 21, 30, 31

[LGX13] LIAN Z., GODIL A., XIAO J.: Feature-preserved 3d canonical form. *Int. J. Comput. Vision 102*, 1-3 (2013), 221–238. 12

[LH13] LI C., HAMZA A.: Symmetry discovery and retrieval of nonrigid 3D shapes using geodesic skeleton paths. *Multimed. Tools Appl.* (2013), 1–21. 18, 19, 20, 30, 31

[LJ13] LI B., JOHAN H.: 3d model retrieval using hybrid features and class information. *Multimed. Tools Appl. 62*, 3 (2013), 821–846. 20, 30, 31

[LLL*15] LI B., LU Y., LI C., GODIL A., SCHRECK T., AONO M., BURTSCHER M., CHEN Q., CHOWDHURY N. K., FANG B., FU H., FURUYA T., LI H., LIU J., JOHAN H., KOSAKA R., KOYANAGI H., OHBUCHI R., TATSUMA A., WAN Y., ZHANG C., ZOU C.: A comparison of 3d shape retrieval methods based on a large-scale benchmark supporting multimodal queries. *Comput. Vis. Image Und. 131*, C (2015), 1–27. 21

[LMM13] LI P., MA H., MING A.: Combining topological and view-based features for 3D model retrieval. *Multimed. Tools Appl. 65*, 3 (2013), 335–361. 16, 30, 31

[LMS13] LAGA H., MORTARA M., SPAGNUOLO M.: Geometry and context for semantic correspondences and functionality recognition in man-made 3D shapes. *ACM T. Graphic. 32*, 5 (Oct. 2013), 150:1–150:16. 19, 20, 22

[LO07] LING H., OKADA K.: An efficient Earth Mover's distance algorithm for robust histogram comparison. *IEEE Trans. Pattern Anal. Mach. Intell. 29*, 5 (2007), 840–853. 19

[LOC14] LI C., OVSJANIKOV M., CHAZAL F.: Persistence-based structural recognition. In *Computer Vision and Pattern Recognition, IEEE Conference on* (2014), pp. 2003–2010. 16

[Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision 60*, 2 (2004), 91–110. 11, 18

[LPD13] LIPMAN Y., PUENTE J., DAUBECHIES I.: Conformal

Wasserstein distance: II. computational aspects and extensions. *Math. Comput. 82*, 281 (2013). 12, 20

[LZSCO09] LIU R., ZHANG H., SHAMIR A., COHEN-OR D.: A part-aware surface metric for shape analysis. *Comput. Graph. Forum 28*, 2 (2009), 397–406. 20, 30, 31

[MÍ11] MÉMOLI F.: Gromov-Wasserstein Distances and the Metric Approach to Object Matching. *Found. Comput. Math. 11*, 4 (2011), 417–487. 13

[MÍ12] MÉMOLI F.: Some properties of Gromov-Hausdorff distances. *Discrete Comput. Geom.* (2012), 1–25. 13

[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 Recogn. 44*, 8 (2011), 1581–1589. 20, 30, 31

[MBBV15] MASCI J., BOSCAINI D., BRONSTEIN M. M., VANDERGHEYNST P.: ShapeNet: Convolutional neural networks on non-euclidean manifolds. *arXiv*, 1501.06297 (2015). 11, 22

[MBH12] MOHAMED W., BEN HAMZA A.: Reeb graph path dissimilarity for 3D object matching and retrieval. *Visual Comput. 28*, 3 (2012), 305–318. 19, 30, 31

[MDTS09] MADEMLIS A., DARAS P., TZOVARAS D., STRINTZIS M. G.: 3D object retrieval using the 3D shape impact descriptor. *Pattern Recogn. 42*, 11 (2009), 2447 – 2459. 10, 17, 18, 21, 30, 31

[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 *Biometrics: Theory Applications and Systems, IEEE International Conference on* (2010), pp. 1–6. 30, 31

[MFP*13] MACHADO J., FERREIRA A., PASCOAL P. B., ABDELRAHMAN M., AONO M., EL-MELEGY M., FARAG A., JOHAN H., LI B., LU Y., TATSUMA A.: SHREC'13 Track: Retrieval of Objects Captured with Low-Cost Depth-Sensing Cameras. In *Eurographics Workshop on 3D Object Retrieval* (2013), The Eurographics Association. 20

[MHK*08] MATEUS D., HORAUD R., KNOSSOW D., CUZZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Computer Vision and Pattern Recognition, IEEE Conference on* (2008), pp. 1–8. 17, 30, 31

[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 Rev. 45*, 1 (2003), 3–49. 7

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

[MRB09] MARTON Z., RUSU R., BEETZ M.: On fast surface reconstruction methods for large and noisy point clouds. In *Robotics and Automation, IEEE International Conference on* (2009), pp. 3218–3223. 21

[MS05] MÉMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math. 5*, 3 (2005), 313–347. 2, 13

[MS09a] MAHMOUDI M., SAPIRO G.: Three-dimensional point cloud recognition via distributions of geometric distances. *Graph. Models 71*, 1 (2009), 22 – 31. 17, 30, 31

[MS09b] MORTARA M., SPAGNUOLO M.: Semantics-driven best view of 3D shapes. *Comput. Graph. 33*, 3 (2009), 280 – 290. 22

[MVR*10] MAVRIDIS L., VENKATRAMAN V., RITCHIE D. W., MORIKAWA N., ANDONOV R., CORNU A., MALOD-DOGNIN N., NICOLAS J., TEMERINAC-OTT M., REISERT M., BURKHARDT H., AXENOPOULOS A., DARAS P.: SHREC'10 Track: Protein Model Classification. In *Eurographics Workshop on 3D Object Retrieval* (2010), The Eurographics Association, pp. 117–124. 20

[NBPF11] NATALI M., BIASOTTI S., PATANÈ G., FALCIDIENO B.: Graph-based representations of point clouds. *Graph. Models 73*, 5 (2011), 151 – 164. 17, 30, 31

[NVT*14] NEUMANN T., VARANASI K., THEOBALT C., MAGNOR M., WACKER M.: Compressed manifold modes for mesh processing. *Comput. Graph. Forum 33*, 5 (2014), 35–44. 14

[OBCCG13] OVSJANIKOV M., BEN-CHEN M., CHAZAL F., GUIBAS L.: Analysis and visualization of maps between shapes. *Comput. Graph. Forum 32*, 6 (2013), 135–145. 15

[OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM T. Graphic. 31*, 4 (July 2012), 30:1–30:11. 2, 13, 14, 20

[OLCO14] OZOLINS V., LAI R., CAFLISCH R., OSHER S.: Compressed plane waves yield a compactly supported multiresolution basis for the laplace operator. *Proc. Natl. Acad. Sci. U.S.A. 111*, 5 (2014), 1691–1696. 14

[OLGM11] OVSJANIKOV M., LI W., GUIBAS L., MITRA N. J.: Exploration of continuous variability in collections of 3d shapes. *ACM T. Graphic. 30*, 4 (2011), 33:1–33:10. 20

[OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *Comput. Graph. Forum 29*, 5 (2010), 1555–1564. 20, 30, 31

[OMPG13] OVSJANIKOV M., MÉRIGOT Q., PĂTRĂUCEAN V., GUIBAS L.: Shape matching via quotient spaces. *Comput. Graph. Forum 32*, 5 (2013), 1–11. 15

[PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Comput. Graph. Forum 32*, 2pt4 (2013), 459–468. 14, 15, 20, 30, 31

[Poi08] POINCARÉ H.: Sur l'uniformisation des fonctions analytiques. *Acta Math. 31*, 1 (1908), 1–63. 12

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

[PSR*14] PICKUP D., SUN X., ROSIN P. L., MARTIN R., CHENG Z., LIAN Z., AONO M., HAMZA A. B., BRONSTEIN A., BRONSTEIN M., ET AL.: SHREC'14 track: Shape retrieval of non-rigid 3D human models. In *Eurographics Workshop on 3D Object Retrieval* (2014), The The Eurographics Association, pp. 101–110. 11, 20

[PZC13] PASQUALOTTO G., ZANUTTIGH P., CORTELAZZO G. M.: Combining color and shape descriptors for 3d model retrieval. *Signal Process. Image 28*, 6 (2013), 608 – 623. 18, 30, 31

[RBB*11] RAVIV D., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., SOCHEN N.: Affine-invariant geodesic geometry of deformable 3D shapes. *Comput. Graph. 35*, 3 (2011), 692 – 697. 17, 18, 19, 20, 30, 31

[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 (2010), 18–39. 20, 30, 31

[RBBK10b] RAVIV D., BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R.: Volumetric heat kernel signatures. In *ACM Workshop on 3D Object Retrieval* (2010), ACM, pp. 39–44. 17, 30, 31

[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 (2009), 381–390. 6

[RCB*15] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *ArXiv*, 1506.05274 (2015). 15

[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. 16

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

[RK14] RAVIV D., KIMMEL R.: Affine invariant geometry for non-rigid shapes. *Int. J. Comput. Vision* (2014), 1–11. 17, 18, 21

[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 T. Graphic. 32*, 4 (2013), 72:1–72:12. 14, 20, 30, 31

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

[RPSS10] RUGGERI M. R., PATANÈ G., SPAGNUOLO M., SAUPE D.: Spectral-driven isometry-invariant matching of 3D shapes. *Int. J. Comput. Vision 89*, 2-3 (2010), 248–265. 19, 30, 31

[RTG00] RUBNER Y., TOMASI C., GUIBAS L. J.: The Earth Mover's distance as a metric for image retrieval. *Int. J. Comput. Vision 40* (2000), 99–121. 13

[Rus10] RUSTAMOV R. M.: Robust Volumetric Shape Descriptor. In *Eurographics Workshop on 3D Object Retrieval* (2010), The Eurographics Association. 17, 30, 31

[SAD*08] SAVRAN A., ALYÜZ N., DIBEKLIOĞLU H., ÇELIKTUTAN O., GÖKBERK B., SANKUR B., AKARUN L.: Biometrics and identity management. Springer-Verlag, Berlin, Heidelberg, 2008, ch. Bosphorus Database for 3D Face Analysis, pp. 47–56. 21

[SAHBZ08] SHENTU Z., AL HASAN M., BYSTROFF C., ZAKI M. J.: Context shapes: Efficient complementary shape matching for protein-protein docking. *Proteins (Other) 70*, 3 (2008), 1056–1073. 20

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

[SBC14] SHAPIRA N., BEN-CHEN M.: Cross-collection map inference by intrinsic alignment of shape spaces. *Comput. Graph. Forum 33*, 5 (2014), 281–290. 15

[SBS*15] SIPIRAN I., BUSTOS B., SCHRECK T., BRONSTEIN A. M., BRONSTEIN M., CASTELLANI U., CHOI S., LAI L., LI H., LITMAN R., SUN L.: Scalability of Non-Rigid 3D Shape Retrieval. In *Eurographics Workshop on 3D Object Retrieval* (2015), The The Eurographics Association. 11, 20

[SHCB11] SHARMA A., HORAUD R., CECH J., BOYER E.: Topologically-robust 3d shape matching based on diffusion geometry and seed growing. In *Computer Vision and Pattern Recognition, IEEE Conference on* (2011), pp. 2481–2488. 20, 30, 31

[SHKF12] SILBERMAN N., HOIEM D., KOHLI P., FERGUS R.: Indoor segmentation and support inference from rgbd images. In *European Conference on Computer Vision - Volume Part V* (2012), Springer-Verlag, pp. 746–760. 21

[SHVS12] SMEETS D., HERMANS J., VANDERMEULEN D., SUETENS P.: Isometric deformation invariant 3D shape recognition. *Pattern Recogn. 45*, 7 (2012), 2817 – 2831. 10, 20, 30, 31

[Sin06] SINGER A.: From graph to manifold Laplacian: The convergence rate. *Appl. Comput. Harmon. A. 21*, 1 (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 Comput. Graph. 11*, 6 (1991), 44–51. 16

[SKK91] SHINAGAWA Y., KUNII T. L., KERGOSIEN Y. L.: Surface coding based on Morse Theory. *IEEE Comput. Graph. 11* (1991), 66–78. 16

[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. 20, 30, 31

[SMKF04] SHILANE P., MIN P., KAZHDAN M., FUNKHOUSER T.: The princeton shape benchmark. In *Shape Modeling International* (2004), IEEE Computer Society, pp. 167–178. 20, 21

[SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. *Comput. Graph. Forum 31*, 5 (2012), 1617–1626. 15, 30, 31

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. *Comput. Graph. Forum 28*, 5 (2009), 1383–1392. 6, 7, 11, 18, 21, 30, 31

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

[SPS14] SAVELONAS M., PRATIKAKIS I., SFIKAS K.: An overview of partial 3d object retrieval methodologies. *Multimed. Tools Appl.* (2014), 1–26. 2, 20

[SSCO08] SHAPIRA L., SHAMIR A., COHEN-OR D.: Consistent mesh partitioning and skeletonisation using the shape diameter function. *Vis. Comput. 24*, 4 (2008), 249–259. 10, 21

[SSN*14] SINGH A., SHA J., NARAYAN K. S., ACHIM T., ABBEEL P.: Bigbird: A large-scale 3d database of object instances. In *IEEE International Conference on Robotics and Automation* (2014), IEEE, pp. 509–516. 21

[SSS*10] SHAPIRA L., SHALOM S., SHAMIR A., COHEN-OR D., ZHANG H.: Contextual part analogies in 3D objects. *Int. J. Comput. Vision 89*, 2-3 (2010), 309–326. 10, 17, 19, 20, 30, 31

[STP12] SFIKAS K., THEOHARIS T., PRATIKAKIS I.: Non-rigid 3D object retrieval using topological information guided by conformal factors. *Visual Comput. 28*, 9 (2012), 943–955. 19, 20, 30, 31

[SZM*08] SIDDIQI K., ZHANG J., MACRINI D., SHOKOUFANDEH A., BOUIX S., DICKINSON S.: Retrieving articulated 3-d models using medial surfaces. *Machine Vision Appl. 19*, 4 (2008), 261–275. 20

[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 T. Vis. Comput. Gr. 19*, 7 (2013), 1199–1217. 2, 10, 17, 20

[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. 20

[TDVC13] TABIA H., DAOUDI M., VANDEBORRE J.-P., COLOT O.: A parts-based approach for automatic 3D shape categorization using belief functions. *ACM Trans. Intell. Syst. Technol. 4*, 2 (2013), 33. 22

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

[tHV10] TER HAAR F. B., VELTKAMP R. C.: Expression modeling for expression-invariant face recognition. *Comput. Graph. 34*, 3 (2010), 231–241. 20

[TKA12] TATSUMA A., KOYANAGI H., AONO M.: A large-scale shape benchmark for 3d object retrieval: Toyohashi shape benchmark. In *Signal Information Processing Association Annual Summit and Conference, Asia-Pacific* (2012), pp. 1–10. 20

[TS04] TUNG T., SCHMITT F.: Augmented reeb graphs for content-based retrieval of 3d mesh models. In *Shape Modeling Applications* (2004), pp. 157–166. 16

[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. 16

[TSDS11] TOMBARI F., SALTI S., DI STEFANO L.: A combined texture-shape descriptor for enhanced 3d feature matching. In *Image Processing, IEEE International Conference on* (2011), pp. 809–812. 18

[TV04] TANGELDER J., VELTKAMP R.: A survey of content based 3D shape retrieval methods. In *Shape Modeling Applications* (2004), pp. 145–156. 2, 10, 20

[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, 10, 19, 20

[TVD09] TIERNY J., VANDEBORRE J.-P., DAOUDI M.: Partial 3D shape retrieval by Reeb pattern unfolding. *Comput. Graph. Forum 28*, 1 (2009), 41–55. 19, 20, 30, 31

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

[VGB*10] VELTKAMP R. C., GIEZEMAN G.-J., BAST H., BAUMBACH T., FURUYA T., GIESEN J., GODIL A., LIAN Z., OHBUCHI R., SALEEM W.: SHREC'10 Track: Large Scale Retrieval. In *Eurographics Workshop on 3D Object Retrieval* (2010), The Eurographics Association, pp. 101–108. 20

[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. *Comput. Graph. Forum 32*, 6 (2013), 189–200. 18, 20, 30, 31

[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, 10

[VL08] VALLET B., LÉVY B.: Spectral geometry processing with manifold harmonics. *Comput. Graph. Forum 27*, 2 (2008), 251–260. 6, 22

[Vt08] VELTKAMP R., TER HAAR F.: Shape retrieval contest (shrec) 2008. In *Shape Modeling and Applications, IEEE International Conference on* (2008), pp. –. 20, 21

[VTH09] VELTKAMP R. C., TER HAAR F. B.: Shrec 2009 - shape retrieval contest. In *Eurographics Workshop on 3D Object Retrieval* (2009), The The Eurographics Association, pp. 57–59. 20

[WBBP12] WANG C., BRONSTEIN M., BRONSTEIN A., PARAGIOS N.: Discrete minimum distortion correspondence problems for non-rigid shape matching. In *Scale Space and Variational Methods in Computer Vision* (2012), vol. 6667 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 580–591. 13, 30, 31

[WBK09] WESSEL R., BLÜMEL I., KLEIN R.: A 3D Shape Benchmark for Retrieval and Automatic Classification of Architectural Data. In *Eurographics Workshop on 3D Object Retrieval* (2009), The Eurographics Association, pp. 53–56. 20

[WLZ10] WANG X., LIU Y., ZHA H.: Intrinsic spin images: A subspace decomposition approach to understanding 3D deformable shapes. In *International Symposium 3D Data Processing, Visualization and Transmission* (2010), pp. 17 – 20. 20, 30, 31

[WMKG07] WARDETZKY M., MATHUR S., KAELBERER F., GRINSPUN E.: Discrete Laplace operators: No free lunch. In *Eurographics Symposium on Geometry Processing* (2007), The Eurographics Association, pp. 33–37. 6

[WR02] WALTER J. A., RITTER H.: On interactive visualization of high-dimensional data using the hyperbolic plane. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2002), ACM, pp. 123–132. 12

[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 *Computer Vision and Pattern Recognition, IEEE Conference on* (2010), pp. 438–445. 6, 10, 17, 18, 20

[XMS14] XIANG Y., MOTTAGHI R., SAVARESE S.: Beyond pascal: A benchmark for 3d object detection in the wild. In *Applications of Computer Vision, IEEE Winter Conference on* (2014), pp. 75–82. 21

[XSW03] XIAO Y., SIEBERT P., WERGHI N.: A discrete reeb graph approach for the segmentation of human body scans. In *3-D Digital Imaging and Modeling, Fourth International Conference on* (2003), pp. 378–385. 16

[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. 9

[ZBH12] ZAHARESCU A., BOYER E., HORAUD R.: Keypoints and local descriptors of scalar functions on 2D manifolds. *Int. J. Comput. Vision 100*, 1 (2012), 78–98. 10, 11, 17, 18, 30, 31

[ZBVH09] ZAHARESCU A., BOYER E., VARANASI K., HORAUD R.: Surface feature detection and description with applications to mesh matching. In *Computer Vision and Pattern Recognition, IEEE Conference on* (2009), pp. 373–380. 10, 19, 30, 31

[ZCOAM14] ZHENG Y., COHEN-OR D., AVERKIOU M., MITRA N. J.: Recurring Part Arrangements in Shape Collections. *Comput. Graph. Forum 33*, 2 (2014), 115–124. 22

Table 1: *Classification of methods with respect to the type of input, type of invariance and type of output (similarity score and/or correspondence).*

Method (refs.)	Input				Invariance			Similarity		Correspondence		
	point	surface	volume	texture	rigid	isometry	other	full	part.	sparse	dense	soft
Blended Intrinsic Maps [KLF11]		√			√	√	non-isometric	√		√		
Conformal similarity and correspondences [BLC*11]		√			√	√		√		√		
Distances distributions [MS09a]	√				√	√		√		√		
Volume GPS [Rus10]			√		√	√	scale	√				
Heat Kernel Signatures (HKS) [SOG09]		√			√	√		√		√		
Scale Invariant HKS [BK10b]		√			√	√	scale	√		√		
Shape Google [BBGO11]		√			√	√	scale	√		√		
Bag of Feature Graphs [HHZQ12]		√			√	√	scale	√		√		
Topology-Invariant Geometry [BBK09]		√			√	√	scale	√		√		
Topology-robust Diffusion Geometry [BBK*10]		√			√	√	scale	√		√		
Spectral Distances [BB11]		√			√	√	scale	√		√		
Minimum Distortion Correspondences [WBBP12]		√			√	√	scale	√		√		
Soft Maps [SNB*12, HZG*12]		√			√	√	non-isometric	√	√		√	√
(Quasi-)harmonic correspondences [KBB*13, PBB*13]		√			√	√	scale	√	√		√	√
Equi-affine Invariant Geometry [RBB*11]		√	√		√	√	scale, affinity	√		√		
One-point isometric matching [OMMG10]		√			√	√		√	√		√	
Topologically-robust matching [SHCB11]		√			√	√		√		√	√	
Volumetric HKS [RBBK10b]			√		√	√	scale	√		√		
Contextual part analogies [SSS*10]		√			√	√	scale	√	√	√		
Geodesic Distance Matrices [SHVS12]		√			√	√	scale	√				
Intrinsic Shape Differences [ROA*13]		√			√	√		√	√		√	√
Mutual Distance Matrices [Bia10]		√			√	√		√				
Wave Kernel Signature [ASC11]		√			√	√		√		√		
3D MSERs [LBB11]		√			√	√		√	√	√		
Spectral descriptors [LB14]		√			√	√	scale, affinity	√		√		
Spectral graph wavelets [LBH13]		√			√	√	scale	√				
Part-aware metric [LZSCO09]		√			√	√		√	√	√		
3D shape impact [MDTS09]		√	√		√		scale	√				
Spatial circular descriptors [GDZ10]		√			√		scale	√				
meshSIFT [MFK*10, SKVS13]		√			√	√		√			√	
Salient Points matching [CCFM08]		√			√	√	scale	√		√		
Salient spectral features [HH09]		√			√	√	scale	√				
Local spectral descriptors [Lav12]		√			√	√	scale	√	√	√		
Intrinsic spin images [WLZ10]		√			√	√		√		√		
Reduced Laplace-Beltrami eigenfunctions [DK10]		√			√	√		√	√	√		
Spectral isometric matching [RPSS10]		√			√	√	scale	√		√		
Semantic best view selection [Lag10]		√	√					√	√			
Facial sparse Matching [BDP13b]		√			√	√	scale		√	√		
Facial expression recognition [BDP13a]		√			√			√		√		
Isogeodesic stripes [BDP10]		√			√	√	scale	√		√		
Bilateral maps [vKZH13]		√			√	√	scale	√	√	√		
Persistence-based GH signatures [CCSG*09]		√			√	√	scale	√		√		
Persistence-based heat signatures [DLL*10]		√			√	√			√	√		
Multidimensional size functions [BCF*08]		√	√		√	√	scale	√		√		
Extended Reeb graphs [BB13b]		√			√	√	scale	√				
Laplacian eigenfunctions for point registration [MHK*08]	√				√	√		√				
Skeleton paths [LH13]		√	√		√	√	scale	√		√		
Point cloud graphs [NBPF11]	√				√	√	scale	√		√		
Extended Reeb graphs [BEGB13, EHB13]		√			√	√	scale	√		√		
Topo-geometric models [BK10a]		√			√		scale	√	√	√		
Skeletal Reeb graphs [MBH12]		√			√		scale	√	√	√		
Reeb graphs [AK11]		√			√	√		√	√	√		
Reeb graphs and unfolding signatures [TVD09]		√			√	√	scale	√	√	√		
Conformal factors [STP12]		√			√	√		√		√		
Graph-based representations [APP*10]		√			√	√		√		√		
Reeb graphs & view [LMM13]			√	√		√		√				
MeshHOG [ZBVH09, ZBH12]		√		√	√	√		√			√	
Photometric HKS [KBBK12, KBB*12]		√		√	√	√	scale, affinity	√		√		
Multi-scale area projection transform [GL12]		√		√	√		scale	√				
PHOG [BCGS13]		√		√	√	√	scale	√		√		
Non-rigid symmetry detection [RBBK10a]		√			√	√	scale, affinity	√	√	√		
Persistence-based recognition of occluded shapes [DL12a]		√			√			√	√			
Geometric histograms [BWdBP13]		√			√	√		√	√		√	
Local facial patches [MAD*11]		√			√		scale	√		√		
Textured Spin Images [PZC13]		√		√	√		scale	√				
Scale-Invariant Spin Images [DK12]		√			√		scale	√		√		
Geodesic/curvature based features [LGJ14]		√			√		scale	√	√			
Hybrid features and class information [LJ13]		√			√		scale	√	√			

Table 2: *Classification of methods according to the type of captured shape structure, type of distance and computational complexity. In particular, computational cost refers to extracting shape description/comparing descriptions.*

Method (refs.)	Structure	Distance	Computational cost
Blended Intrinsic Maps [KLF11]	Conformal	Wasserstein-based	medium/high
Conformal similarity and correspondences [BLC*11]	Conformal	Minimum-distortion embeddings	Low/high
Distances distributions [MS09a]	Diffusion	L^1, L^2, χ^2, Jensen-Shannon divergence	medium/low
Volume GPS [Rus10]	Diffusion	χ^2	medium/medium
Heat Kernel Signatures (HKS) [SOG09]	Multi-scale autodiffusion	L^1	medium/low
Scale Invariant HKS [BK10b]	Multi-scale autodiffusion	Weighted L^1	medium/low
Shape Google [BBGO11]	Multi-scale autodiffusion	Hamming distance	medium/low
Bag of Feature Graphs [HHZQ12]	Multi-scale autodiffusion	L^2-based	medium/low
Topology-Invariant Geometry [BBK09]	Diffusion	Joint similarity (Gromov-Hausdorff distance)	medium/medium
Topology-Robust Diffusion Geometry [BBK*10]	Diffusion	Gromov-Hausdorff distance	medium/high
Spectral Distances [BB11]	Diffusion	Normalized L^1,L^2,χ^2, Earth Mover's Distance	medium/high
Minimum Distortion Correspondences [WBBP12]	Multi-scale autodiffusion	Gromov-Hausdorff distance	medium/high
Soft Maps [SNB*12, HZG*12]	Local conformal	Functional maps	medium/high
(Quasi-)harmonic correspondences [KBB*13, PBB*13]	Diffusion	Functional maps	medium/medium
Equi-Affine Invariant Geometry [RBB*11]	Diffusion	Gromov-Hausdorff distance	medium/high
One-point isometric matching [OMMG10]	Multi-scale autodiffusion	L^2-based	medium/low
Topologically-robust matching [SHCB11]	Multi-scale autodiffusion	L^2-, L^∞-based	medium/low
Volumetric HKS [RBBK10b]	Multi-scale autodiffusion	L^1	medium/low
Contextual part analogies [SSS*10]	Semi-local conformal	Bipartite graph matching	medium/medium
Geodesic Distance Matrices [SHVS12]	Conformal	χ^2	medium/low
Intrinsic shape differences [ROA*13]	Conformal	Functional maps	medium/medium
Mutual Distance Matrices [Bia10]	Diffusion	Soft Earth Mover's Distance	medium/medium
Wave Kernel Signature [ASC11]	Multi-scale autodiffusion	L^1	medium/low
3D MSERs [LBB11]	Semi-local autodiffusion	Local point & region distance	medium/low
Spectral descriptors [LB14]	Multi-scale autodiffusion	Metric learning	medium/medium
Spectral graph wavelets [LBH13]	Diffusion	Intrinsic spatial pyramid matching	medium/high
Part-aware metric [LZSCO09]	Conformal	χ^2	medium/low
3D shape impact [MDTS09]	Semi-local extrinsic	Normalized L^2 and diffusion distance	low/low
Spatial circular descriptors [GDZ10]	Extrinsic	Hungarian distance	low/medium
meshSIFT [MFK*10, SKVS13]	Multi-scale conformal	Angle distance	medium/low
Salient Points matching [CCFM08]	Multi-scale conformal	Distance on Hidden Markov Models	low/medium
Salient spectral features [HH09]	Diffusion	*Ad-hoc* distance	medium/medium
Local spectral descriptors [Lav12]	Local diffusion	Bag of features	medium/low
Intrinsic spin images [WLZ10]	Multi-scale intrinsic	Eearth Mover's Distance	medium/medium
Reeuced Laplace-Beltrami eigenfunctions [DK10]	Diffusion	*Ad-hoc* (quadratic optimization problem)	medium/medium
Spectral isometric matching [RPSS10]	Diffusion	Bipartite graph matching, L^1-based	medium/medium
Semantic best view selection [Lag10]	Conformal	*Ad-hoc* distance	medium/medium
Facial sparse Matching [BDP13b]	Multi-scale conformal	*Ad-hoc* distance, RANSAC	medium/low
Facial expression recognition [BDP13a]	Multi-scale extrinsic	Distance on Hidden Markov Models	medium/medium
Isogeodesic stripes [BDP10]	Conformal	Weighted graph distance	medium/medium
Bilateral maps [vKZH13]	Semi-local conformal	*Ad-hoc* (functional optimization)	low/high
Persistence-based GH signatures [CCSG*09]	Multi-scale topology	Gromov-Hausdorff distance, natural pseudo-distance	low/medium
Persistence-based heat signatures [DLL*10]	Multi-scale autodiffusion & topology	L^1-based	low/low
Multidimensional size functions [BCF*08]	Multi-scale topology	Natural pseudo-distance	low/medium-high
Extended Reeb graphs [BB13b]	Semi-local topology	Kernel aggregation & learning	low/high
Laplacian eigenfunctions for point registration [MHK*08]	Diffusion	Hungarian distance	medium/medium
Skeleton paths [LH13]	Extrinsic & topology	Endpoint distance	low/low
Point cloud graphs [NBPF11]	Semi-local topology	Spectral graph distance	low/medium
Extended Reeb graphs [BEGB13, EHB13]	Semi-local topology & topology	Spectral graph distance	low/medium
Topo-geometric models [BK10a]	Extrinsic & topology	Maximal sub-graph approximation	low/medium
Skeletal Reeb graphs [MBH12]	Extrinsic topology	Shortest path graph matching	low/medium-high
Reeb graphs [AK11]	Conformal topology	Maximal sub-graph approximation	low/medium-high
Reeb graphs and unfolding signatures [TVD09]	Semi-local conformal & topology	Maximal sub-graph approximation	low/medium
Conformal factors [STP12]	Semi-local conformal & topology	Hungarian distance	low/medium-high
Graph-based representations [APP*10]	Semi-local conformal	Earth Mover's Distance	low/high
Reeb graphs & view [LMM13]	Conformal & Multi-scale topology	Earth Mover's Distance	medium/high
MeshHOG [ZBVH09, ZBH12]	Multi-scale conformal & photometry	L^2	medium/low
Photometric HKS [KBBK12, KBB*12]	Multi-scale autodiffusion & photometry	Bag of features	medium/low
Multi-scale area projection transform [GL12]	Multi-scale extrinsic	Jeffrey divergence	medium/low
PHOG [BCGS13]	Conformal & photometry, multi-scale topology	L^1, natural pseudo-distance	medium/medium-high
Non-rigid symmetry detection [RBBK10a]	Multi-scale autodiffusion	Minimum-distortion embeddings	medium/high
Persistence-based recognition of occluded shapes [DL12a]	Multi-scale extrinsic & topology	Hausdorff and Natural pseudo-distance	low/low-medium
Geometric histograms [BWdBP13]	Multi-scale	RANSAC	medium/low
Local facial patches [MAD*11]	Semi-local	Geodesic in a scale space	low/medium
Textured spin images [PZC13]	Photometry	Weighted statistical function	medium/medium
Scale-Invariant Spin Images [DK12]	Local extrinsic	RANSAC	medium/low
Geodesic/curvature based features [LGJ14]	Semi-local conformal	*Ad-hoc* distance	low/low
Hybrid features and class information [LJ13]	Semi-local conformal	scaled L_1 and Canberra distance	low/medium