

Correspondence-Free Region Localization for Partial Shape Similarity via Hamiltonian Spectrum Alignment

Arianna Rampini
Sapienza University of Rome
rampini@di.uniroma1.it

Irene Tallini
Sapienza University of Rome
tallini.1608460@studenti.uniroma1.it

Maks Ovsjanikov
LIX, École polytechnique
maks@lix.polytechnique.fr

Alex M. Bronstein
Technion
bron@cs.technion.ac.il

Emanuele Rodolà
Sapienza University of Rome
rodola@di.uniroma1.it

Abstract

We consider the problem of localizing relevant subsets of non-rigid geometric shapes given only a partial 3D query as the input. Such problems arise in several challenging tasks in 3D vision and graphics, including partial shape similarity, retrieval, and non-rigid correspondence. We phrase the problem as one of alignment between short sequences of eigenvalues of basic differential operators, which are constructed upon a scalar function defined on the 3D surfaces. Our method therefore seeks for a scalar function that entails this alignment. Differently from existing approaches, we do not require solving for a correspondence between the query and the target, therefore greatly simplifying the optimization process; our core technique is also descriptor-free, as it is driven by the geometry of the two objects as encoded in their operator spectra. We further show that our spectral alignment algorithm provides a remarkably simple alternative to the recent shape-from-spectrum reconstruction approaches. For both applications, we demonstrate improvement over the state-of-the-art either in terms of accuracy or computational cost.

1. Introduction

Assessing similarity between non-rigid shapes is an active research topic in computer vision, pattern recognition and graphics [5]. At the heart of such methods lies the definition of shape *descriptors*, characterizing the shape either locally or globally (e.g., via the Bag-of-Words [38] paradigm or via deep learning [27]). Deformation-invariant shape descriptors often require careful tuning and hand-crafting, or sufficient training examples to enable learned-based methods. While similarity by itself is usually expressed by a numerical score, the problem as a whole is

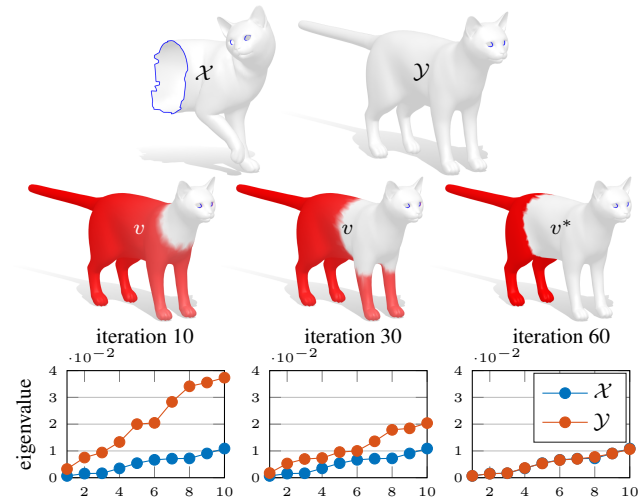


Figure 1. Given a non-rigid partial 3D query \mathcal{X} and a full shape \mathcal{Y} , we propose a way to locate \mathcal{X} in \mathcal{Y} without having to compute a correspondence $\pi : \mathcal{X} \rightarrow \mathcal{Y}$. We do so by looking for an indicator function v on \mathcal{Y} . From this function a Laplacian-like matrix on \mathcal{Y} can be constructed, whose eigenvalues coincide with those of the classical Laplacian on \mathcal{X} . The optimal v^* attaining the eigenvalue alignment localizes the sought region. Here we show the evolution of v (middle row) and of the corresponding eigenvalues (bottom row) across our iterative algorithm until convergence.

also strictly related to (and often solved in tandem with) the complementary problem of shape *correspondence*. In this setting, local shape descriptors are used as ‘probe’ quantities to employ in more sophisticated pipelines to infer a functional [28] or point-to-point [7] correspondence.

Even more challenging is the setting of *partial* shape similarity, which is the focus of this paper. In this case, only some *regions* of the shapes are expected to be similar due to the presence of clutter or missing geometry. Partial similarity arises in numerous practical tasks, ranging from

object reassembly [18, 24] to protein docking [36] and deformable object-in-clutter detection [15] among others. In the partial setting, in addition to defining a similarity metric, one also has to identify the relevant similar regions of the given shapes, which is a difficult problem in itself (see Figure 1). As a result, while global similarity has been extensively studied, only few methods exist to address the (especially non-rigid) partial 3D similarity problem.

A seemingly unrelated problem concerns the recovery of the geometry of an unknown shape (*i.e.*, whose point coordinates in space are not given) from a minimal amount of input data. Notably, it has been recently shown [13] that, in some cases, one can reconstruct the geometry of a shape even when an incomplete sequence of its Laplacian eigenvalues (its spectrum) are given as an input – a problem widely known in mathematical physics as “hearing the shape of the drum” [21]. In vision and graphics, this problem finds applications in non-isometric shape matching as well as style transfer [13] and acoustics design [4].

In this paper we show that partial shape similarity and shape-from-spectrum recovery, while apparently very distinct, can be phrased as instances of the same mathematical problem. We formulate this problem by exploiting the notion of the Hamiltonian operator on manifolds, and reduce it to an unconstrained continuous optimization problem over the vector space of real-valued functions defined over the surfaces; see Figure 1 for an illustration of our methodology on a real example.

1.1. Related work

Partial similarity of 3D shapes has long been researched by the vision and geometry processing communities, with the vast majority tackling the *rigid* setting. Since the focus of this work is on *non-rigid* structures, we discuss methods addressing deformable objects here, and refer the interested reader to the recent surveys [37, 31, 29] for the former case.

Partial shape retrieval. At a high level our task is related to shape retrieval, where the goal is to produce a *ranking* of most similar shapes according to a similarity score. Even in this simpler task, surprisingly few methods have attempted to address the problem. We mention here the recent SHREC’17 challenge [34], highlighting how deep learning-based methods may work well for producing rankings, provided that enough training examples are available (the organizers provided around two thousands).

Our task is more challenging, as we aim at identifying the regions that two non-rigid objects have in common.

Partial shape correspondence. When a map $\pi : \mathcal{X} \rightarrow \mathcal{Y}$ between shapes \mathcal{X} and \mathcal{Y} is available, one may simply look at the image $\pi(\mathcal{X}) \subset \mathcal{Y}$ in order to determine the corresponding region (for simplicity, here we assume that \mathcal{X} is a deformed region of \mathcal{Y} as in Figure 1). Thus, any partial

shape matching method implicitly solves for the common region as a side product, at the cost of solving a computationally heavy correspondence problem.

To date, solving a full correspondence problem is therefore the dominant approach. Bronstein *et al.* [7] proposed to explicitly solve for a point-to-point map π and a membership function μ (or “fuzzy part”, measuring the degree of inclusion of each point into the subset), optimizing for π and μ in an alternating fashion. Optimization for π was posed as a NP-hard quadratic assignment problem [9], while solving for μ required solving a non-convex quadratic problem due to the presence of region regularizers. More recently, Rodolà *et al.* [33] followed a similar approach where the correspondence step was rephrased using the functional map representation [28]. A membership function is still optimized for, with the same regularizers as those in [7]. The framework was extended in [24] to address the case with $n > 2$ regions (“non-rigid puzzle”), leading to a multi-label segmentation problem where n functional maps are optimized simultaneously. The presence of clutter was investigated in [15] using the same formulation of [33], but where the local descriptors (which are used in all the methods above to drive the correspondence) are learned via deep metric learning on a representative dataset. Finally, Brunton *et al.* [8] proposed an iterative approach where a partial map is recovered by an isometric growing procedure starting from sparse pre-matched features over the surfaces.

In contrast to these methods, our approach does not require solving for dense correspondences, thus essentially eliminating the need for local descriptors, and does not make use of complex regularizers. Instead, we phrase the optimization problem in terms of solving for a real-valued (indicator) function aimed at identifying similar regions, and devise a correspondence-free metric that allows us to measure and optimize for the similarity of such functions directly.

A related approach was previously introduced by Pokrass *et al.* [30] for identifying region-wise similarity. However, unlike our method, that approach relies on the comparison of *descriptor* statistics over the surface, thus strongly depending on the specific choice of descriptors.

Isospectralization. From an algebraic standpoint, ours may be seen as an “inverse eigenvalue problem” [11]; despite the name, problems in this domain typically require knowledge of the eigenvectors of the sought operator, while our method does not. Perhaps, the most closely related to ours is the recent work on *isospectralization* [13] where the authors deform the shapes with the goal of aligning their Laplacian eigenvalues. We also use a similar observation that in practice these eigenvalues provide a powerful characterization of the shape’s geometry. However, we focus on the partial setting and propose to solve for a characteristic function of a region, thus avoiding any explicit shape deformation.

1.2. Contribution

The key contribution of this paper is to show that for many practical problems, it is possible to side-step the need for a correspondence by looking at the eigenvalues of certain differential operators constructed on each surface. This leads to tangible consequences that we leverage in two challenging applications. Specifically,

- We introduce a new method for detecting similar regions among deformable shapes; for the first time, the optimization is completely *correspondence-free* and, at its core, descriptor-free;
- Our optimization problem uses a single objective, thus avoiding possibly complex regularizers, and can be solved *efficiently* with off-the-shelf differentiable programming libraries;
- We provide a remarkably simple alternative to recent *isospectralization* techniques [13], while at the same time yielding qualitatively better results at a fraction of the time cost.

2. Background

In this paper we model shapes as 2-dimensional Riemannian manifolds \mathcal{X} (possibly with boundary $\partial\mathcal{X}$) equipped with the area element dx . We denote by $\text{int}(\mathcal{X})$ the interior of \mathcal{X} , namely the set of points $\mathcal{X} \setminus \partial\mathcal{X}$.

Laplacian. The positive semi-definite Laplace-Beltrami operator $\Delta_{\mathcal{X}}$ (or manifold Laplacian) generalizes the basic differential operator from Euclidean analysis to Riemannian manifolds. It admits a spectral decomposition

$$\begin{aligned} \Delta_{\mathcal{X}}\phi_i(x) &= \mu_i\phi_i(x) & x \in \text{int}(\mathcal{X}) & \quad (1) \\ \phi_i(x) &= 0 & x \in \partial\mathcal{X}, & \quad (2) \end{aligned}$$

with homogeneous Dirichlet boundary conditions (2), where $\{0 < \mu_1 \leq \mu_2 \leq \dots\}$ are the eigenvalues (each with finite multiplicity) and ϕ_i are the associated eigenfunctions. For the remainder of the paper, we always assume eigenvalues to be increasingly ordered.

The *spectrum* of $\Delta_{\mathcal{X}}$ is the set of all eigenvalues, forming a discrete subset of \mathbb{R}_+ . It encodes geometric and topological properties of \mathcal{X} ; for instance, it has been used as a basic global shape descriptor for 3D shape retrieval [32].

A direct consequence of Eq. (1) is that each eigenvalue μ_i is associated to an eigenfunction ϕ_i by the equality¹

$$\mu_i = \int_{\mathcal{X}} \|\nabla_{\mathcal{X}}\phi_i(x)\|_2^2 dx \quad (3)$$

¹In this equality we use the intrinsic gradient operator $\nabla_{\mathcal{X}}$ acting on scalar functions $f : \mathcal{X} \rightarrow \mathbb{R}$, which defines the direction (tangent to the surface) of the steepest change of f around each point $x \in \mathcal{X}$.

after normalization so that $\int_{\mathcal{X}} \phi_i^2(x)dx = 1$. Equation (3) makes the *global* nature of the eigenvalues clear, since they are computed as integrals over the entire surface.

Hamiltonian operator. Given a scalar *potential* function $v : \mathcal{X} \rightarrow \mathbb{R}_+$, one can define a Hamiltonian $H_{\mathcal{X}} = \Delta_{\mathcal{X}} + v$, operating on scalar functions f as follows:

$$H_{\mathcal{X}}f = \Delta_{\mathcal{X}}f + vf, \quad (4)$$

where the product vf is taken point-wise. Trivially, for the null potential where $v(x) = 0 \forall x \in \mathcal{X}$, the Hamiltonian $H_{\mathcal{X}}$ simplifies to $\Delta_{\mathcal{X}}$. The operator $H_{\mathcal{X}}$ also admits a spectral decomposition:

$$H_{\mathcal{X}}\psi_i(x) = \lambda_i\psi_i(x), \quad (5)$$

with boundary conditions when necessary. Hamiltonian eigenfunctions have been used in shape analysis applications in [10, 25, 35].

Discretization. In the discrete setting, \mathcal{X} is approximated by a triangle mesh with n vertices $v_i \in V_{\text{int}} \cup V_{\text{bdr}}$, and where each edge $e_{ij} \in E_{\text{int}} \cup E_{\text{bdr}}$ belongs to at most two triangle faces F_{ijk} and F_{jih} (see inset for notation). We use the subscripts *int* and *bdr* on vertices and edges to denote interior and boundary, respectively.

The Laplacian is defined in terms of two $n \times n$ matrices \mathbf{W} and \mathbf{A} , where \mathbf{A} is a diagonal matrix of local area elements a_i and \mathbf{W} is a symmetric matrix of edge-wise weights (also known as cotangent formula, see *e.g.* [26]):

$$w_{ij} = \begin{cases} -(\cot \alpha_{ij} + \cot \beta_{ij})/2 & e_{ij} \in E_{\text{int}}; \\ 0 & (i \neq j) \wedge (i \in V_{\text{bdr}} \vee j \in V_{\text{bdr}}); \\ \sum_{k \neq i} (\cot \alpha_{ik} + \cot \beta_{ik})/2 & (i = j) \wedge v_i \in V_{\text{int}}; \\ 1 & (i = j) \wedge v_i \in V_{\text{bdr}}; \end{cases} \quad (6)$$

$$a_i = \begin{cases} \frac{1}{3} \sum_{j,k:i,j,k \in F} A_{ijk} & v_i \in V_{\text{int}}; \\ 0 & v_i \in V_{\text{bdr}}; \end{cases} \quad (7)$$

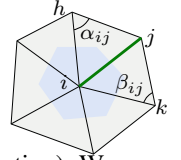
where A_{ijk} is the area of triangle F_{ijk} . In the formulas above, the special treatment of boundary edges and vertices imposes the Dirichlet boundary conditions of Eq. (2).

A generalized eigenproblem $\mathbf{W}\Phi = \mathbf{A}\Phi\text{diag}(\boldsymbol{\mu})$ is solved for computing the Laplacian eigenvalues and eigenvectors. In the case of the Hamiltonian, it takes the form

$$(\mathbf{W} + \mathbf{A}\text{diag}(\mathbf{v}))\Psi = \mathbf{A}\Psi\text{diag}(\boldsymbol{\lambda}). \quad (8)$$

Above, \mathbf{v} is a n -dimensional vector containing the values of the potential v at each vertex, Φ and Ψ are matrices containing the first k eigenvectors as their columns, and $\boldsymbol{\lambda}, \boldsymbol{\mu}$ are the corresponding k -dimensional vectors of eigenvalues.

For brevity, we will write $\boldsymbol{\lambda}(\Delta_{\mathcal{X}} + \text{diag}(\mathbf{v}))$ to denote the eigenvalues appearing in Eq. (8).



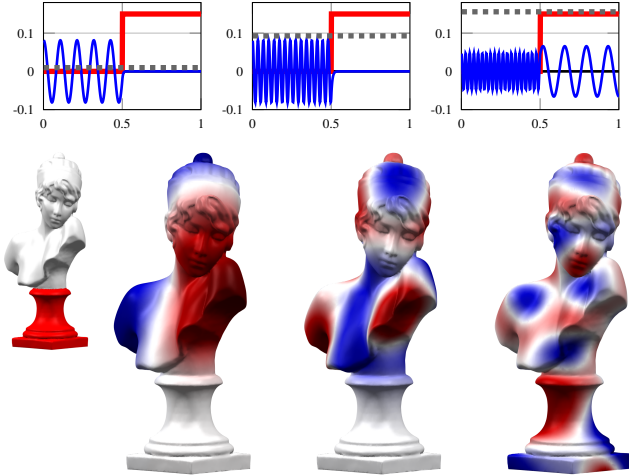


Figure 2. *Top*: A step potential v_τ (solid red) confines eigenfunction ψ_i (solid blue) within the region where $v_\tau(x) = 0$, as long as for the associated eigenvalue λ_i (dashed gray) it holds $\lambda_i < \tau$. As soon as $\lambda_i \geq \tau$, a “diffusion” effect across the potential barrier takes place (rightmost plot). We illustrate this in \mathbb{R} with three eigenfunctions at increasing eigenvalue (left to right). *Bottom*: We show the same phenomenon on a surface. Here, the potential with height τ is depicted on the left and is supported on the pedestal; observe the diffusion in the rightmost figure, due to the eigenvalue of the plotted eigenfunction being higher than τ .

3. Summary of the reasoning

A classical result in mathematical physics [19, Ch. 2] shows that the eigenfunctions ψ_i of $H_{\mathcal{X}}$ exhibit *localization* depending not only on the geometry of the domain \mathcal{X} [20], but also on the shape of the potential v . In this paper, we consider step potentials with finite height $\tau > 0$:

$$v_\tau(x) = \begin{cases} 0 & x \in \mathcal{R} \\ \tau & x \notin \mathcal{R} \end{cases} \quad (9)$$

for a given region $\mathcal{R} \subseteq \mathcal{X}$. The following lemma on the “confinement” property of step potentials is instrumental to our approach:

Lemma 1 [19, Ch. 2.6]) *Let $\mathcal{R} \subseteq \mathcal{X}$ be a region of \mathcal{X} , and let $v_\tau : \mathcal{X} \rightarrow \{0, \tau\}$ be a finite step potential defined as in Eq. (9). Then, for the Hamiltonian $\Delta_{\mathcal{X}} + v_\tau$ all eigenfunctions ψ_i with eigenvalue $\lambda_i < \tau$ are subject to exponential decay outside \mathcal{R} .*

According to Lemma 1, for large enough τ , the Hamiltonian eigenfunctions will be approximately supported within \mathcal{R} . This result holds both on the real line (where the eigenfunctions are Fourier harmonics) and on manifolds². We illustrate real examples of these facts in Figure 2.

²Berger [3, p.403] described this analogy quite vividly, by depicting manifolds as “quantum mechanical worlds”.

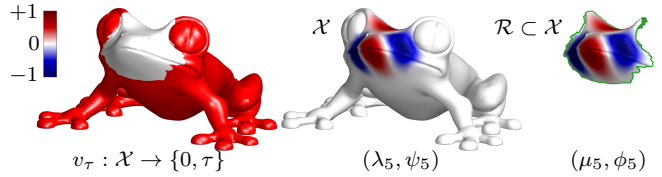


Figure 3. *Left*: Given a region \mathcal{R} (the face of the frog), we show the step potential v_τ such that $v_\tau(x) = 0$ for $x \in \mathcal{R}$, and $v_\tau(x) = \tau$ otherwise. *Middle*: The eigenfunctions ψ_i of the Hamiltonian $\Delta_{\mathcal{X}} + v_\tau$ are localized on \mathcal{R} , meaning that $\psi_i(x) \approx 0$ for $x \notin \mathcal{R}$. *Right*: If the region \mathcal{R} is taken as a separate mesh with boundary $\partial\mathcal{R}$ (depicted in green), the Dirichlet eigenfunction ϕ_i of $\Delta_{\mathcal{R}}$ has the same values as ψ_i at corresponding points. As a consequence, $\lambda_i = \mu_i$.

Key to our approach is the realization that the confined behavior of Hamiltonian eigenfunctions can be exploited by optimizing over the space of potentials v_τ . We therefore establish the following lemma:

Lemma 2 (Eigenvalue equivalence) *Let $\mathcal{R} \subseteq \mathcal{X}$ be a region of \mathcal{X} , and let $v_\tau : \mathcal{X} \rightarrow \{0, \tau\}$ be a finite step potential defined as in Eq. (9). Now let $\{(\lambda_i, \psi_i)\}_{i=1}^k$ be k eigenpairs of $\Delta_{\mathcal{X}} + v_\tau$, increasingly ordered and such that $\lambda_k < \tau$. If (μ_i, ϕ_i) are the eigenpairs of $\Delta_{\mathcal{R}}$ computed with Dirichlet boundary conditions, then $\mu_i \simeq \lambda_i$ for all $i = 1, \dots, k$.*

Sketch of proof. For any eigenpair (λ_i, ψ_i) of $\Delta_{\mathcal{X}} + v_\tau$ with $\lambda_i < \tau$, the following equality holds pointwise:

$$\lambda_i \psi_i(x) = \Delta_{\mathcal{X}} \psi_i(x) + v_\tau(x) \psi_i(x) \quad (10)$$

$$\simeq \Delta_{\mathcal{R}} \psi_i(x) \quad (11)$$

$$= \mu_i \phi_i(x). \quad (12)$$

Eq. (11) follows directly from Lemma 1; Eq. (12) stems from the homogeneous Dirichlet condition $\phi_i(x) = 0$ for $x \in \partial\mathcal{R}$. \square

An illustration of this result is given in Figure 3. Lemma 2 establishes that one can measure the Dirichlet eigenvalues of a partial shape \mathcal{R} equivalently by measuring Hamiltonian eigenvalues over the complete shape \mathcal{X} – provided that an indicator for the corresponding region on \mathcal{X} is given. This gives us a straightforward way to implement localization via optimization over the cone of non-negative real-valued functions on \mathcal{X} . Problem (13) in the next Section formalizes precisely this idea.

4. Our method

We first present and discuss our general formulation, and then demonstrate its application to two separate tasks in Sections 5.1 and 5.2. The main idea behind our approach is to model regions $\mathcal{R} \subset \mathcal{X}$ via potential functions, relying upon the empirical observation that, in many practical

cases, the Laplacian eigenvalues encode the geometry of the domain.

Optimization problem. Our input is a given shape \mathcal{X} and, separately, an increasing sequence of k Dirichlet eigenvalues stored in a vector $\boldsymbol{\mu} \in \mathbb{R}^k$. The precise source of these eigenvalues is *not known*, but they are assumed to come from a manifold Laplacian. It is our task to identify the shape of this domain.

To do so, we consider the following *eigenvalue alignment* problem:

$$\min_{\mathbf{v} \geq 0} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}_{\mathcal{X}} + \text{diag}(\mathbf{v})) - \boldsymbol{\mu}\|_w^2, \quad (13)$$

where we seek for an optimal alignment (according to a weighted norm defined below) between the input eigenvalues $\boldsymbol{\mu}$ and the Hamiltonian eigenvalues $\boldsymbol{\lambda}$; the minimization is carried out over the space of non-negative potentials on \mathcal{X} , and the w -norm is a weighted L_2 norm:

$$\|\boldsymbol{\lambda} - \boldsymbol{\mu}\|_w^2 = \sum_{i=1}^k \frac{1}{\mu_i^2} (\lambda_i - \mu_i)^2, \quad (14)$$

where the weights avoid unduly penalization of lower frequencies. Problem (13) therefore models the input sequence $\boldsymbol{\mu}$ as the first k eigenvalues of a Hamiltonian (whose potential v we seek) constructed on \mathcal{X} . This allows us to model very general tasks in practice, as we show in our experiments.

Remark. We emphasize that our objective is phrased entirely in terms of eigenvalues, while the eigenfunctions never appear explicitly.

Algorithm. Although the optimization problem in Eq. (13) is non-convex, it is still differentiable and we observed very good local optima by using standard numerical optimization methods. For example, projected gradient descent with step $\alpha > 0$ can be easily implemented by the recursive relations:

$$\begin{aligned} \mathbf{v}^{(t)} &= \mathbf{v}^{(t-1)} - \Pi \left(\alpha \nabla \|\boldsymbol{\lambda}(\boldsymbol{\Delta} + \text{diag}(\mathbf{v}^{(t-1)})) - \boldsymbol{\mu}\|_w^2 \right) \\ &= \mathbf{v}^{(t-1)} - \Pi \left(2\alpha (\boldsymbol{\Phi}^{(t-1)} \circ \boldsymbol{\Phi}^{(t-1)}) ((\boldsymbol{\lambda}^{(t-1)} - \boldsymbol{\mu}) \circ \boldsymbol{\mu}^2) \right) \end{aligned}$$

Here, $\Pi(\mathbf{x}) \equiv \max\{\mathbf{x}, 0\}$, with element-wise max, is a projector onto \mathbb{R}_+^n ; element-wise product and division are denoted by \circ and \oslash respectively. The matrix $\boldsymbol{\Phi}^{(t-1)}$ contains the eigenvectors of $\boldsymbol{\Delta} + \text{diag}(\mathbf{v}^{(t-1)})$ as its columns; note that we do *not* optimize over these eigenvectors, but they appear naturally in the eigenvalue derivatives [2].

To simplify even further, we consider the *unconstrained* problem

$$\min_{\mathbf{v} \in \mathbb{R}^n} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}_{\mathcal{X}} + \text{diag}(\sigma(\mathbf{v}))) - \boldsymbol{\mu}\|_w^2, \quad (15)$$

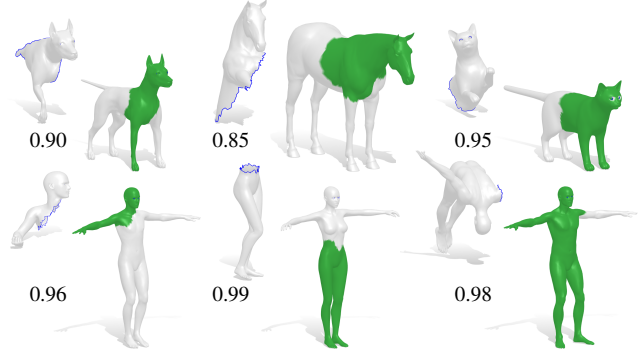


Figure 4. Examples of correctly detected regions using our method on 6 different classes from the SHREC'16 dataset [14]. For each example we show the partial 3D query and its optimal localization on the full model (in green). Below each partial shape we also report the intersection-over-union area ratio between the detected region and the ground truth (ideally 1.0).

where $\sigma(\mathbf{v}) = \frac{\tau}{2} (\tanh(\mathbf{v}) + 1)$ is a saturation function, acting element-wise, which maps the values of \mathbf{v} to within the range $(0, \tau)$. This has the effect of promoting step potentials with height τ , ensuring that all eigenfunctions with associated eigenvalue $\lambda_i < \tau$ will be confined within the region where $\sigma(v(x)) = 0$ (by Lemma 1). In our tests we simply set $\tau = 10 \mu_k$. We minimize problem (15) by a standard trust-region method implemented in Matlab.

5. Results

We showcase our eigenvalue alignment method on two distinct tasks. As we present these results, we also highlight some key properties of our technique that can be exploited in several applications.

5.1. Partial shape localization

As a first application, we consider the task of finding a region of a non-rigid shape from a partial (and possibly deformed) 3D query. Differently from prior work, we aim to do this without having to compute a map between full and partial shape at the same time.

Let \mathcal{X} and \mathcal{Y} be the full and partial shapes, respectively, with the Laplacians $\boldsymbol{\Delta}_{\mathcal{X}}$ and $\boldsymbol{\Delta}_{\mathcal{Y}}$. Further, let $\{\mu_i\}_{i=1}^k$ be the first k eigenvalues of $\boldsymbol{\Delta}_{\mathcal{Y}}$. Our task is to find a region $\mathcal{R} \subseteq \mathcal{X}$ that corresponds to the partial query \mathcal{Y} up to isometry.

To solve this problem, we minimize Eq. (15) directly by using $\boldsymbol{\Delta}_{\mathcal{X}}$ and $\{\mu_i\}$ as input. A local solution will be a potential function $v_\tau : \mathcal{X} \rightarrow (0, \tau)$ indicating the sought region $\mathcal{R} \subseteq \mathcal{X}$; see Figure 4 for qualitative examples.

Property 1 *Our method is **correspondence-free**, since we never need to invoke the notion of a map $\pi : \mathcal{X} \rightarrow \mathcal{Y}$ in order to find the region $\mathcal{R} \subseteq \mathcal{X}$. In particular, we only operate with eigenvalue sequences $\{\lambda_i\}, \{\mu_i\}$, which are invariant to point ordering (see Eq. (3)).*

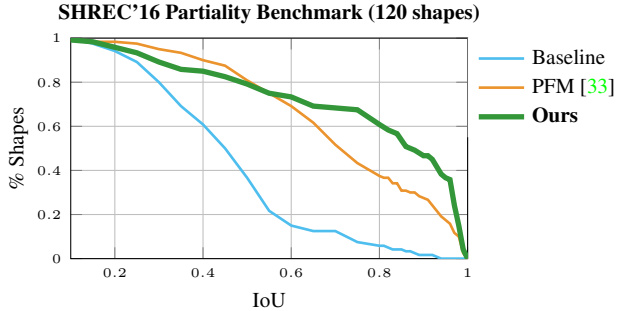


Figure 5. Partial similarity comparisons. We compare with partial functional maps (PFM), a state-of-the-art partial shape matching method. We measure the intersection-over-union (IoU) between the regions identified by each method and the ground truth regions, and plot the cumulative scores over a set of 120 shape pairs. The performance of the baseline is also shown.

It is important to remark that the eigenvalues alignment does not imply that the correct region is always recovered, as there could exist non-isometric regions sharing the same spectrum [17]. However, in our experiments we observed that the output potential v_τ correctly identifies the sought region whenever the eigenvalues align, implying the following property:

Property 2 Empirically, *eigenfunction alignment* is induced by our eigenvalue alignment. Formally, this means that $\lambda_i = \mu_i$ implies $\psi_i(x) = \phi_i(\pi(x))$ for all $x \in \mathcal{R}$ and an isometric map $\pi : \mathcal{R} \rightarrow \mathcal{Y}$.

We stress that the map π is *not* solved for, and is not available to us after optimization. A similar property was also observed in [13], and might be a feature common to eigenvalue alignment approaches in practice. A deeper comparison to [13] will be provided in Section 5.2. Finally, invariance to deformations is put in evidence below:

Property 3 Since the eigenvalues of $\Delta_{\mathcal{X}}$ and $\Delta_{\mathcal{Y}}$ are *isometry-invariant*, so are all solutions to problem (15).

The latter property suggests that other isometry-invariant methods from the shape analysis area might also be used for the same task. We now discuss the most relevant methods.

Comparisons. Partial shape matching methods such as [7, 33, 24] iteratively compute a map and a region. Such methods make use of regularization functionals for the region in order to reach a good solution. While our potential v_τ plays a similar role as the membership functions of these methods, we neither use any region regularizer, nor do we need to solve for a map. The fully spectral approach of Litany *et al.* [23] seeks for an optimal alignment between *eigenfunctions* rather than eigenvalues, thereby solving a pointwise matching problem in the process.

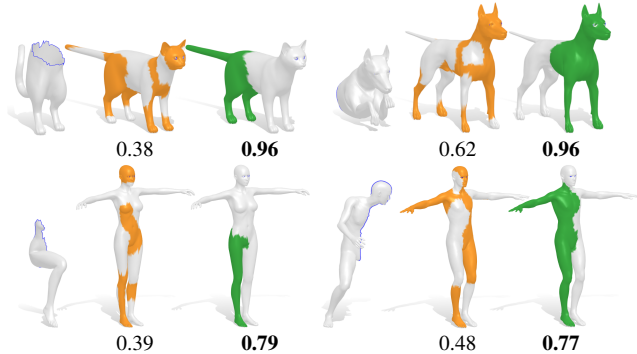


Figure 6. Comparisons with PFM on the SHREC'16 benchmark. For each example we show the partial query, the PFM solution (orange), and our solution (green). The numerical score below each solution denotes the intersection-over-union with respect to the ground truth mask. Note how PFM tends to mix-up the bilateral symmetries of the shapes (left-right and front-back), since it relies on local descriptors that can not discriminate orientation.

We investigated extensively the effectiveness of our method in comparison with the state-of-the-art shape matching approach Partial Functional Maps (PFM) [33] on the SHREC'16 Partiality benchmark [14], where PFM showed top performance. The dataset consists of 120 partial deformable shapes with large missing parts ('cuts' challenge), belonging to 8 sub-classes of humans and animals. The challenge is to locate each of the 120 partial shapes within a full template of the respective class (e.g., a dog in neutral pose). To make the setting even more challenging, we remeshed all the full models \mathcal{X} to half of their density ($\sim 10K$ to $\sim 5K$ vertices) via edge collapse [16], while keeping the partial shapes \mathcal{Y} at their original resolution. This way, we eliminate the possibility for any method to work well by virtue of similar tessellation.

For PFM we used public code released by the authors with default parameters. As a baseline, we also evaluate the co-segmentation approach of [38] with SHOT descriptors [39]. Quantitative and qualitative comparisons are reported in Figure 5 and Figure 6 respectively.

Robustness to sampling. Our method exhibits robustness to changes in resolution (see Figure 8 for a full evaluation). We attribute this to the fact that eigenvalues are integral quantities (see Eq. (3)). In contrast, PFM fully relies upon a data term based on local point descriptors, which are extremely sensitive to local meshing.

Multiple pieces. We highlight another useful property derived from working with potentials. In Figure 7 we show examples where the input 3D query consists of *multiple* disconnected parts (the "non-rigid puzzle" setting with non-overlapping pieces [24]). By taking their union and treating them as one single shape \mathcal{Y} in our optimization problem, we are still able to correctly identify the distinct portions on the

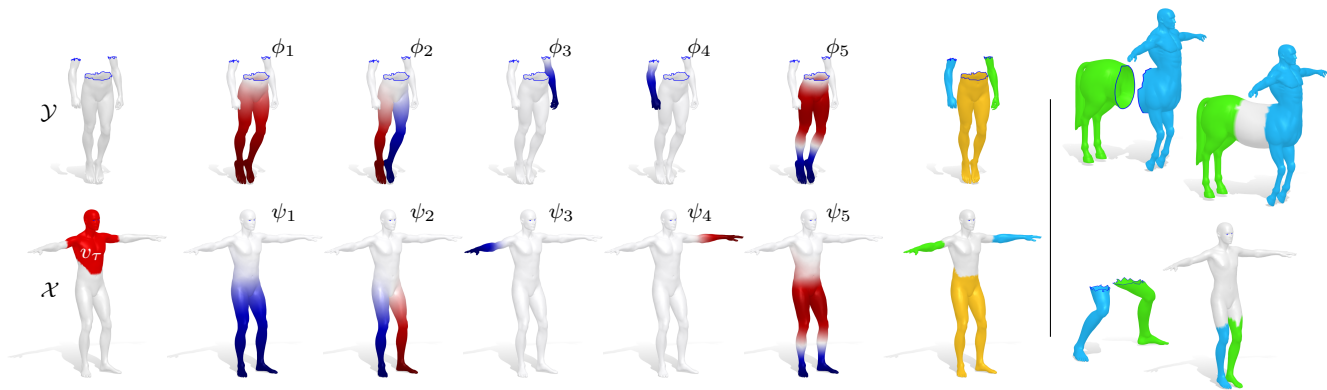


Figure 7. Simultaneous localization of multiple disconnected regions. Given a query shape \mathcal{Y} (here consisting of 3 pieces), our method solves for a potential $v_\tau : \mathcal{X} \rightarrow (0, \tau)$ which isolates the regions on \mathcal{X} corresponding to \mathcal{Y} up to isometry (leftmost column). By Property 2, the found v_τ induces alignment of the eigenfunctions (intermediate columns); we exploit this to easily associate a different label to each region. On the right we plot the identified regions with different colors. The rightmost column shows additional multi-piece examples.

full shape without any special adjustment to our algorithm.

Implementation details. In all our tests we use $k = 20$ eigenvalues. As initialization for \mathbf{v} , we use Gaussians centered around $m = 20$ Euclidean farthest samples on \mathcal{X} , with variance equal to $\sqrt{\text{Area}(\mathcal{X})}$ and $2\sqrt{\text{Area}(\mathcal{X})}$, leading to $2m$ different initializations, which we optimize in parallel. To these, we also add the solutions produced by the baseline. To choose the final solution we rely on Property 2. For each found v , we compute projections of SHOT descriptors $F : \mathcal{X} \rightarrow \mathbb{R}$ over the first 5 Hamiltonian eigenfunctions (squared), $c_i = \int_{\mathcal{X}} \psi_i^2(x) F(x) dx$ for $i = 1, \dots, 5$. We then compute coefficients d_i on \mathcal{Y} similarly by using the Dirichlet eigenfunctions, and keep the v which yields minimum distance $\sum_{i=1}^5 (d_i - c_i)^2$. By Property 2, we expect these coefficients to be the same for the correct solution.

5.2. Isospectralization

As a second application we address the shape-from-spectrum problem. This task was recently considered in [13] and is phrased as follows: given as input a short sequence of Laplacian eigenvalues $\{\mu_i\}_{i=1}^k$ of an *unknown* shape \mathcal{Y} , recover a geometric embedding of \mathcal{Y} . We stress that the Laplacian $\Delta_{\mathcal{Y}}$ is *not* given; if given, it would lead to a shape-from-operator problem [6]. Instead, the goal here is to recover the shape directly from the eigenvalues alone.

We model this problem by minimizing Eq. (15) with the following input. As full domain \mathcal{X} we consider a finite portion of the plane $\mathcal{X} \subset \mathbb{R}^2$ of size $n \times n$, sampled on a Cartesian grid. The input eigenvalues $\{\mu_i\}$ are assumed to be coming from an unknown region $\mathcal{R} \subset \mathcal{X}$, which we identify by the potential v_τ on \mathcal{X} . Therefore, our optimization variables are simply $\mathbf{v} \in (0, \tau)^{n \times n}$ – essentially, an image.

Property 4 We can recover shapes with *unknown topology*, i.e., no restricting assumptions are made on the topology of the sought shape.

For example, if the shape to be recovered has holes, we do not need to know this fact beforehand to properly initialize the optimization.

Comparisons. In Figure 9 we compare directly with the isospectralization technique of Cosmo *et al.* [13]; to the best of our knowledge, this is the only existing approach attempting to solve shape reconstruction purely from eigenvalues. Our method is different in many respects.

The most crucial difference lies in the fact that in [13] the authors optimize for a *deformation field* to apply to an initial template shape, which is assumed to be given as an input. In contrast, with Problem (15) we simply minimize over a vector space of scalar functions. The initial shape used in [13] is also required to have the correct topology (e.g., disc-like or annulus-like), while by Property 4 our optimization

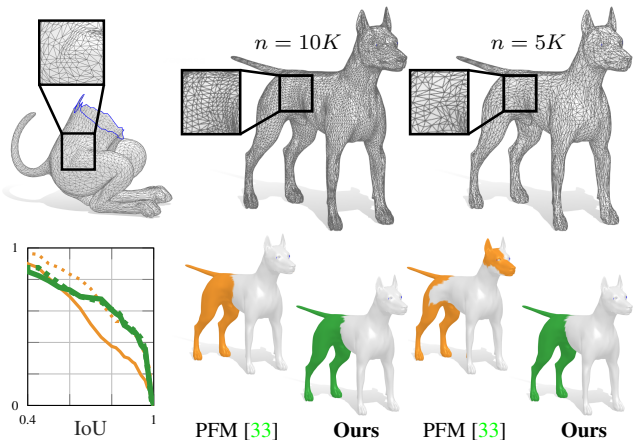


Figure 8. Robustness to changes in vertex density. We show the results of PFM (in orange) and our method (in green) over two versions of the entire SHREC’16 dataset: a version where all meshes have similar resolution (dotted curves in the plots), and one where mild remeshing has been applied (solid curves). PFM shows a performance drop of 20%, while our method remains stable.

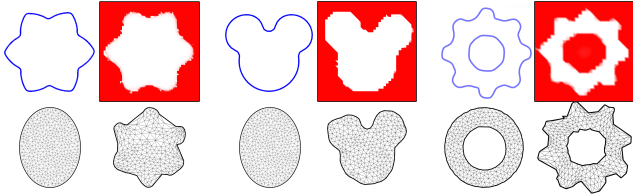


Figure 9. Shape recovery comparisons with [13]. We show the ground truth shapes (blue outline) and our solutions (potential on the square, growing from white to red), all obtained in 2-3 minutes from $k = 50$ input eigenvalues. We always use a constant initialization. On the bottom we show results for [13] with their initial templates, which are required to always have the correct topology.

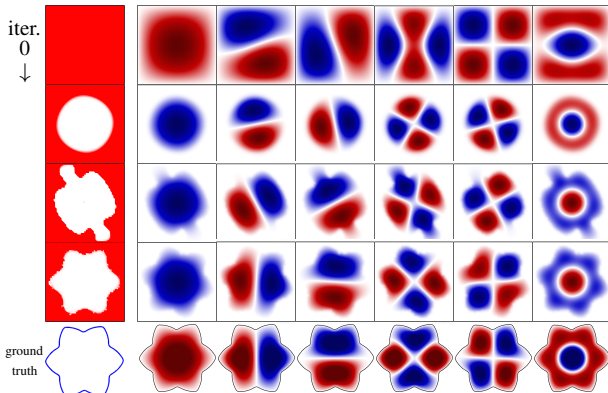


Figure 10. Shape recovery across iterations. Starting from a sequence of just $k = 5$ Laplacian eigenvalues, our method recovers the correct shape in less than 1 minute. The first column shows our solution, evolving from top to bottom. On the right we also show how the eigenfunctions (one per column) change across the iterations; observe how at convergence, they align with those of the ground truth shape (up to rotation).

is completely oblivious to the topology of the shape to recover. We illustrate this with an example in the rightmost column of Figure 9, where we directly compare with [13].

Finally, optimizing over embedding coordinates as in [13] requires additional regularizers to avoid flipped triangles, and careful parameter tuning to avoid rough shapes and collapsing edges; this makes their energy more difficult to minimize, which is done using stochastic optimization tools [22] and automatic differentiation as implemented in TensorFlow [1]. By optimizing over scalar functions instead of deforming a given mesh, we do not encounter any of these issues; for this reason, our optimization can be performed much simpler as detailed in Section 4.

Implementation details. The planar domain \mathcal{X} is discretized as a 30×30 grid with uniform triangles. We initialize simply with the constant potential $\mathbf{v} = \frac{\tau}{4}\mathbf{1}$, in order to emphasize the ability of our method to recover shape entirely from scratch irrespective of topology. We put this in contrast with [13], which needs an initial template as initialization, as mentioned above; see Figure 10.

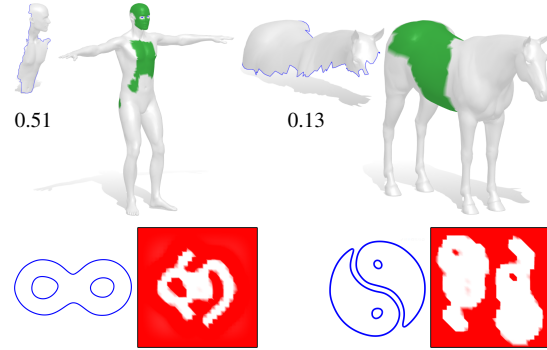


Figure 11. Failure cases. *Top*: On the left example, the eigenvalues did not align well at the local optimum; as a result, the partial shape was localized only roughly. On the right, the found region is wrongly classified as correct since it has similar eigenvalues to the partial input within the limited bandwidth. In both cases we report the IoU. *Bottom*: Shape-from-spectrum recovery of shapes with non-trivial topology can be difficult to achieve, due to the narrow bandwidth and first-order discretization of the operators.

6. Conclusions

We presented a new approach for partial shape similarity, based on aligning short sequences of eigenvalues computed on the given shapes. Our approach is correspondence-free, deformation-invariant, it can be solved efficiently, and only makes use of descriptors for post-processing refinement. Our formulation is also general enough for other relevant tasks in vision and graphics. We demonstrated this on a challenging shape-from-spectrum recovery problem, where it compares favorably with respect to a very recent method.

Limitations and future work. We did observe some failure cases during our tests (see Figure 11). In the partial similarity setting, this typically happens whenever the chosen bandwidth (*i.e.*, the number k of input eigenvalues) is too small for discriminating regions that only differ at medium-high frequencies; *e.g.*, a human arm might be confused with a leg. This can be remedied to some extent by increasing the value of k . For large values ($k > 100$), however, a more accurate discretization of the operator Δ should be used in place of the cotangent formulas of Eq. (6), to avoid artifacts related to mesh tessellation. Using second- or third-order FEM discretization [12] might be a promising solution.

Finally, while we only showed 2D examples of shape-from-spectrum recovery, the same algorithm can be used for reconstructing the geometry of 3D (volumetric) shapes. However, this would require solving for a potential defined in \mathbb{R}^3 , for which efficient data structures could be necessary. We leave this exciting possibility to future work.

Acknowledgments

AR and ER are supported by the ERC StG no. 802554 (SPECGEO). MO is supported by the KAUST OSR Award no. CRG-2017-3426, a gift from NVIDIA Corporation and the ERC StG no. 758800 (EXPROTEA).

References

- [1] M. Abadi, A. Agarwal, P. Barham, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. [8](#)
- [2] K. T. Abou-Moustafa. On derivatives of eigenvalues and eigenvectors of the generalized eigenvalue problem. Technical report, McGill University, May 2009. [5](#)
- [3] M. Berger. *A panoramic view of Riemannian geometry*. Springer Science & Business Media, 2012. [4](#)
- [4] G. Bharaj, D. I. Levin, J. Tompkin, Y. Fei, H. Pfister, W. Matusik, and C. Zheng. Computational design of metallophone contact sounds. *TOG*, 34(6), 2015. [2](#)
- [5] S. Biasotti, A. Cerri, A. Bronstein, and M. Bronstein. Recent trends, applications, and perspectives in 3d shape similarity assessment. *Computer Graphics Forum*, 35(6):87–119, 2016. [1](#)
- [6] D. Boscaini, D. Eynard, D. Kourounis, and M. M. Bronstein. Shape-from-operator: Recovering shapes from intrinsic operators. *Computer Graphics Forum*, 34(2):265–274, 2015. [7](#)
- [7] A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Partial similarity of objects, or how to compare a centaur to a horse. *International Journal of Computer Vision*, 84(2):163, 2009. [1](#), [2](#), [6](#)
- [8] A. Brunton, M. Wand, S. Wuhrer, H.-P. Seidel, and T. Weinkauff. A low-dimensional representation for robust partial isometric correspondences computation. *Graphical Models*, 76(2):70–85, 2014. [2](#)
- [9] R. E. Burkard, E. Cela, P. M. Pardalos, and L. S. Pitsoulis. The quadratic assignment problem. In *Handbook of combinatorial optimization*, pages 1713–1809. Springer, 1998. [2](#)
- [10] Y. Choukroun, A. Shtern, A. M. Bronstein, and R. Kimmel. Hamiltonian operator for spectral shape analysis. *TVCG*, 2018. [3](#)
- [11] M. Chu and G. Golub. *Inverse eigenvalue problems: theory, algorithms, and applications*. Oxford University Press, 2005. [2](#)
- [12] P. G. Ciarlet. *The finite element method for elliptic problems*, volume 40. Siam, 2002. [8](#)
- [13] L. Cosmo, M. Panine, A. Rampini, M. Ovsjanikov, M. M. Bronstein, and E. Rodolà. Isospectralization, or how to hear shape, style, and correspondence. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#), [3](#), [6](#), [7](#), [8](#)
- [14] L. Cosmo, E. Rodolà, M. Bronstein, A. Torsello, D. Cremers, and Y. Sahillioglu. SHREC’16: Partial matching of deformable shapes. In G. A. Ferreira A., Giorgi D., editor, *Eurographics Workshop on 3D Object Retrieval, EG 3DOR*, pages 61–67. Eurographics Association, 2016. [5](#), [6](#)
- [15] L. Cosmo, E. Rodolà, J. Masci, A. Torsello, and M. M. Bronstein. Matching deformable objects in clutter. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 1–10. IEEE, 2016. [2](#)
- [16] M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 209–216. ACM Press/Addison-Wesley Publishing Co., 1997. [6](#)
- [17] C. Gordon, D. L. Webb, and S. Wolpert. One cannot hear the shape of a drum. *Bulletin of the American Mathematical Society*, 27(1):134–138, 1992. [6](#)
- [18] R. Gregor, I. Sipiran, G. Papaioannou, T. Schreck, A. Andreadis, and P. Mavridis. Towards automated 3d reconstruction of defective cultural heritage objects. In *Proceedings of the Eurographics Workshop on Graphics and Cultural Heritage*, pages 135–144. Eurographics Association, 2014. [2](#)
- [19] D. J. Griffiths. *Introduction to Quantum Mechanics*. Prentice Hall, 1994. [4](#)
- [20] S. M. Heilman and R. S. Strichartz. Localized eigenfunctions: Here you see them, there you don’t. *Notices of the AMS*, 57(5):624–629, 2010. [4](#)
- [21] M. Kac. Can one hear the shape of a drum? *The American Mathematical Monthly*, 73(4):1–23, 1966. [2](#)
- [22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [8](#)
- [23] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Fully spectral partial shape matching. *Computer Graphics Forum*, 36(2):247–258, 2017. [6](#)
- [24] O. Litany, E. Rodolà, A. M. Bronstein, M. M. Bronstein, and D. Cremers. Non-rigid puzzles. *Computer Graphics Forum*, 35(5):135–143, 2016. [2](#), [6](#)
- [25] S. Melzi, E. Rodolà, U. Castellani, and M. M. Bronstein. Localized manifold harmonics for spectral shape analysis. *Computer Graphics Forum*, 37(6):20–34, 2018. [3](#)
- [26] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III*, pages 35–57. Springer, 2003. [3](#)
- [27] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5115–5124, 2017. [1](#)
- [28] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30, 2012. [1](#), [2](#)
- [29] Q.-H. Pham, M.-K. Tran, W. Li, S. Xiang, H. Zhou, W. Nie, A. Liu, Y. Su, M.-T. Tran, N.-M. Bui, et al. SHREC’18: Rgb-d object-to-cad retrieval. In *Proc. 3DOR*, 2018. [2](#)
- [30] J. Pokrass, A. M. Bronstein, and M. M. Bronstein. Partial shape matching without point-wise correspondence. *Numerical Mathematics: Theory, Methods and Applications*, 6(1):223–244, 2013. [2](#)
- [31] I. Pratikakis, M. A. Savelonas, F. Arnaoutoglou, G. Ioannakis, A. Koutsoudis, T. Theoharis, M.-T. Tran, V.-T. Nguyen, V.-K. Pham, H.-D. Nguyen, et al. Partial shape queries for 3d object retrieval. In *Proc. 3DOR*, pages 79–88. Eurographics Association, 2016. [2](#)
- [32] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace–beltrami spectra as ‘shape-dna’ of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006. [3](#)

- [33] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. *Computer Graphics Forum*, 36(1):222–236, 2017. [2](#), [6](#), [7](#)
- [34] E. Rodolà, L. Cosmo, O. Litany, M. Bronstein, A. Bronstein, et al. SHREC’17: Deformable shape retrieval with missing parts. In P. I. Dupont F., Ovsjanikov M., editor, *Eurographics Workshop on 3D Object Retrieval, EG 3DOR*, pages 85–94. Eurographics Association, 2017. [2](#)
- [35] E. Rodolà, Z. Löhner, A. M. Bronstein, M. M. Bronstein, and J. Solomon. Functional maps representation on product manifolds. *Computer Graphics Forum*, 38(1):678–689, 2019. [3](#)
- [36] Z. Shentu, M. Al Hasan, C. Bystroff, and M. J. Zaki. Context shapes: Efficient complementary shape matching for protein–protein docking. *Proteins: Structure, Function, and Bioinformatics*, 70(3):1056–1073, 2008. [2](#)
- [37] I. Sipiran, R. Meruane, B. Bustos, T. Schreck, H. Johan, B. Li, and Y. Lu. SHREC’13 track: large-scale partial shape retrieval using simulated range images. In *Proceedings of the Sixth Eurographics Workshop on 3D Object Retrieval*, pages 81–88. Eurographics Association, 2013. [2](#)
- [38] R. Toldo, U. Castellani, and A. Fusiello. A bag of words approach for 3d object categorization. In *Proc. MIRAGE*, pages 116–127. Springer, 2009. [1](#), [6](#)
- [39] F. Tombari, S. Salti, and L. Di Stefano. Unique signatures of histograms for local surface description. In *Proc. ECCV*, pages 356–369. Springer, 2010. [6](#)