Int J Comput Vis
DOI 10.1007/s11263-016-0883-8

Spectral Generalized Multi-dimensional Scaling

Yonathan Aflalo[1] · **Anastasia Dubrovina**[1] · **Ron Kimmel**[1]

Received: 28 February 2015 / Accepted: 19 January 2016

Abstract Multidimensional scaling (MDS) is a family of methods that embed a given set of points into a simple, usually flat, domain. The points are assumed to be sampled from some metric space, and the mapping attempts to preserve the distances between each pair of points in the set. Distances in the target space can be computed analytically in this setting. Generalized MDS is an extension that allows mapping one metric space into another, that is, MDS into target spaces in which distances are evaluated numerically rather than analytically. Here, we propose an efficient approach for computing such mappings between surfaces based on their natural spectral decomposition, where the surfaces are treated as sampled metric-spaces. The resulting *spectral-GMDS* procedure enables efficient embedding by incorporating smoothness of the metric structure into the problem, thereby substantially reducing the complexity involved in its solution while practically overcoming its non-convex nature. The method is compared to existing techniques that compute dense correspondence between shapes. Numerical experiments of the proposed method demonstrate its efficiency and accuracy compared to state-of-the-art approaches especially when isometry invariance is a dominant property.

Keywords Spectral domain · GMDS · Shape matching

Communicated by B. C. Vemuri.

✉ Yonathan Aflalo
yaflalo@cs.technion.ac.il

Anastasia Dubrovina
nastyad@cs.technion.ac.il

Ron Kimmel
ron@cs.technion.ac.il

[1] Technion Institute of Technology, Haifa, Israel

1 Introduction

Matching non-rigid or deformable shapes is a challenging problem involving a large number of degrees of freedom. While matching rigid objects one needs to search for isometries in a three dimensional Euclidean space, a problem that can be described by six parameters. Matching solvers for rigid surfaces in $\mathbb{R}^3$ are known as iterative closest point algorithms or ICP (Chen and Medioni 1991; Besl and McKay 1992). Non-rigid matching usually involves much more dimensions that can add up to the number of points of the sampled surfaces that one wishes to match. When ignoring the continuity of matching one surface to another, the problem can be viewed as a combinatorial one, for which the computational complexity is exponential. The problem in this setting is NP hard, the hardest to solve in terms of computational complexity. The question we address is how to efficiently solve this notoriously hard problem.

Various attempts to define robust and invariant meaningful measures by which articulated objects and deformable shapes could be identified were made. Adopting tools from metric geometry, the Gromov–Hausdorff distance (Gromov 1981; Burago et al. 2001), and its variants were suggested as candidates for measuring the discrepancy between two deformable shapes (Memoli and Sapiro 2005; Memoli 2007). The Gromov–Hausdorff distance between two surfaces S and Q, or $d_{\mathrm{GH}}(S, Q)$ in short, is the maximal distortion introduced when bijectively embedding S into Q and vice-versa. Motivated by early attempts of finding a common parametrization for surfaces (Schwartz et al. 1989) the idea of treating surfaces as metric spaces that can be embedded into simple spaces was thought of. There, the metric of each surface is first embedded into a small dimensional Euclidean space, say $\mathbb{R}^3$, by a procedure known as multidimensional scaling (MDS) (Borg and Groenen 1997). The flat mappings

Published online: 29 March 2016

or *canonical forms* in the Euclidean space are then treated as rigid surfaces and matched, for example, by ICP. Though the idea is appealing, the embedding error of mapping a non-flat manifold into a flat finite dimensional domain can be substantial with little hope for convergence. The question is how to avoid intermediate simple spaces while still being able to computationally handle the seemingly complicated task of matching non-rigid surfaces. Towards that end, Memoli and Sapiro (2005) provided the support that sampling surfaces could be tolerated within the Gromov–Hausdorff framework. In other words, the sampling error is linear as a function of the distance between the sampled points, and could thus be bounded when comparing sampled surfaces. Equipped with that encouraging result, Bronstein et al. (2008) exploited the fact that d_{GH} could be formulated as three coupled generalized multidimensional scaling (GMDS) problems for which they introduced a numerical solver.

In retrospective, the Hausdorff measure optimized for by the celebrated iterative closet point (ICP) procedure (Chen and Medioni 1991; Besl and McKay 1992) can be interpreted as a Gromov–Hausdorff distance, where distances are computed in the embedding $\mathbb{R}^3$ Euclidean space. Other simple intermediate embedding spaces for matching non-rigid shapes were advocated. The eigenspace of the Laplace–Baltrami operator (LBO) was suggested in various flavors, for example by Mateus et al. (2008), and by Rustamov et al. (2013), as potential Euclidean target space, see also Bérard et al. (1994), Coifman and Lafon (2006), Lévy (2006). Lipman et al. (2009, 2011) embedded shapes conformally into disks between which the correspondence boils down again to a six parameters Möbius transform, see also Gu et al. (2004), Jin et al. (2004), Zeng et al. (2012). In that case, metric embedding errors are replaced by numerical ones, as important features with effective Gaussian curvature often scale down substantially and can practically vanish when sub-sampled. Kim et al. (2011) suggested a refinement procedure, while using conformal mappings that perform well only locally. They softly tailored a handful of such locally good maps, using a procedure they coined as *blending*.

Heuristics that reduce the complexity of the dense matching problem and detect some initial state at a significant basin of attraction for convex solvers to refine were often employed by the above approaches. Such heuristics use feature point detectors and descriptors. Some examples include the heat kernel signature (HKS) (Sun et al. 2009; Gębal et al. 2009), global point signature (GPS) (Rustamov et al. 2013), wave kernel signature (WKS) (Aubry et al. 2011), and scale-space representation (Zaharescu et al. 2009). Matching the metric spaces with either geodesic (Memoli and Sapiro 2005) or diffusion (Bérard et al. 1994; Coifman and Lafon 2006) distances, could then be treated as a regularization or refinement term. It produced dense correspondence from the sparse one provided by matching the feature points (Dubrovina and Kimmel 2010). Higher order structures were suggested, for example, in Zeng et al. (2010). Dense matching was further accelerated by hierarchical solvers like Raviv et al. (2012), Sahillioğlu and Yemez (2011). Still, the complexity of searching over the space of all possible point-to-point correspondences was determined by the number points one wishes to match.

Ovsjanikov et al. (2012) illuminated the fact that given two functional spaces, and given the correspondence between these two spaces, there is a linear relation between the functional representation of a function in one space (shape) and its corresponding functional representation in the second space (shape). This linear relation is due to the given correspondence and can be viewed as a matrix translating the decomposition coefficients of a function in one metric space to its set of corresponding coefficients in the other. When the functional spaces are the eigenfunctions of the surface LBO (Lévy 2006), the right matching matrix for isometric surfaces would be nothing but the identity. Ovsjanikov et al. (2012) named these linear connections between functional spaces as *functional maps* and used them to find dense correspondence between shapes. Under the assumption of *smooth* function representation, for which the Laplace–Beltrami provides a natural basis, as was proven in Aflalo and Kimmel (2013), Aflalo et al. (2015a), only a small number of leading eigenfunctions may be considered. Thus, the combinatorial problem of correspondence detection can be casted as low dimensional functional map identification. As always, in Ovsjanikov et al. (2012), Pokrass et al. (2013), a number of matching regions or feature points was required for computing the correspondence using functional maps. To the best of our knowledge, the most accurate shape matching results, obtained using functional representations, were realized in Shtern and Kimmel (2014, 2015), after this paper has been submitted. When restricting the deformations to almost isometries, the proposed method obtains similar correspondence rates.

When matching non-isometric shapes, the corresponding Laplace–Beltrami eigenspaces are incompatible. This effect is substantial, for instance, in the case of various human body shapes in the SCAPE dataset (Anguelov et al. 2004). Pokrass et al. (2013) subsequently formulated the non-rigid isometric matching problem as permuted sparse coding. There, the dense correspondence is extracted through coupling the functional map representation with that of matching corresponding regions. The computation is performed by alternating minimization over the unknown functional map, while penalizing non-diagonal solutions, and a permutation matrix, representing the correspondences.

In this paper, we argue that the L_2 version of the Gromov–Hausdorff framework for matching deformable shapes, namely, the GMDS (Bronstein et al. 2006), can be naturally casted into the spectral domain. The present spec-

tral formulation is denoted as *spectral GMDS*, or S-GMDS in short. We utilize the following important observations:

- The point-to-point correspondence between two shapes induces a map between the natural eigen-spaces of the shapes. Furthermore, one can use truncated eigen-spaces to faithfully approximate this correspondence.
- Distances measured on shapes are smooth functions, and as such are well suited for compact spectral representation, see (Aflalo and Kimmel 2013; Aflalo et al. 2015a), thereby allowing us to translate the Gromov–Hausdorff framework into the spectral domain.

In a nutshell, we show that treating the shape matching problem as a mapping between metric spaces can lead to an efficient and accurate solver. Specifically, both metric spaces and the mapping are dealt with in the dual spectral domain—a representation space which is justified by its provable theoretical optimality (Aflalo et al. 2015a). Furthermore, we show that when isometry invariance is approximately maintained through poses, the suggested procedure outperforms state-of-the-art dense correspondence solvers in terms of complexity and accuracy, while avoiding the need for supporting features.

2 Generalized Multidimensional Scaling

Consider the shape correspondence problem that involves in searching for the best point to point assignment of two given shapes, S and Q. The *Generalized Multi-Dimensional Scaling* (Bronstein et al. 2008) is a procedure that computes the map that best preserves the inter-geodesic distances while embedding one surface into another. Formally, if $d_S(s, s')$ and $d_Q(q, q')$ represent the inter-geodesic distances between $s, s' \in S$ and $q, q' \in Q$, respectively, then, the generalized multidimensional scaling problem is defined as finding a mapping $\psi : S \to Q$ that minimizes

$$\underset{\psi}{\operatorname{argmin}} \int_{S\times S} \left(d_S(s, s') - d_Q(\psi(s), \psi(s'))\right)^2 da_s da_{s'}. \quad (1)$$

Here, $s, s' \in S$, and da_s is an area element about $s \in S$.

The GMDS is related to the Gromov–Hausdorff distance between metric spaces defined as

$$d_{\mathrm{GH}}(S, Q) = \frac{1}{2} \min_{C} \max_{\substack{(s,q)\in C \\ (s',q')\in C}} \left|d_S(s, s') - d_Q(q, q')\right| \quad (2)$$

where

$$\forall s \in S \;\; \exists q \in Q \text{ s.t. } (s, q) \in C$$

and

$$\forall q \in Q \;\; \exists s \in S \text{ s.t. } (s, q) \in C. \quad (3)$$

In this case, C represents the set of corresponding points. The set C could be defined through an indicator function $p(s, q)$ such that $p(s, q) = 1$ if $(s, q) \in C$ and $p(s, q) = 0$ for $(s, q) \notin C$. Several variations were proposed to reduce the complexity of the problem (Bronstein et al. 2008; Lipman and Daubechies 2011; Pokrass et al. 2013). In the GMDS, the L_∞ Hausdorff distance was replaced by an L_2 norm.

We next symmetrize the GMDS problem through the function $p(s, q)$ that indicates whether two points $s \in S$ and $q \in Q$ correspond to one another. Instead of $p : S \times Q \to \{0, 1\}$, a continuous weak form for $p : S \times Q \to \mathbb{R}^+$ is employed, such that $\int_S p(s, q)da_s = 1$ and $\int_Q p(s, q)da_q = 1$. Our new $p(s, q)$ defines a fuzzy correspondence between the surfaces. To define the matching problem, let us denote the correspondence between the two shapes by a pair of mappings $\phi : Q \to S$ and $\psi : S \to Q$, such that $\psi = \phi^{-1}$. That is, for any $s \in S, q \in Q$, their corresponding points are $\psi(s)$ and $\phi(q)$, respectively, and it holds that $s = \phi(\psi(s))$.

Let us assume that for any pair of corresponding points $q_0 \in Q$ and $s_0 \in S$, the function $p(s_0, q_0)$ defined above is approximately a Dirac delta function, that is $p(s_0, q_0) \approx \delta(\phi(q_0) - s_0)$, and, similarly, $p(s_0, q_0) \approx \delta(q_0 - \psi(s_0))$. Here, the delta function $\delta(s)$ is defined in a classical sense, such that $\int_S \delta(s)da_s = 1$ and $\int_S f(s)\delta(s - s_0)ds = f(s_0)$, for any continuous function $f : S \to \mathbb{R}$. Under this assumption, given a mapping $\phi(q)$ of a point $q \in Q$, the distance measured between $\phi(q)$ and some $s \in S$ is given by

$$d_S(\phi(q), s) = \int_S d_S(s', s)p(q, s')da_{s'}. \quad (4)$$

Similarly, the distance measured between q and $\psi(s)$ on Q is given by

$$d_Q(q, \psi(s)) = \int_Q d_Q(q, q')p(q', s)da_{q'}. \quad (5)$$

To measure the quality of the mapping $p(s, q)$ we compare the distances $d_S(\phi(q), s)$ and $d_Q(q, \psi(s))$, for all $s \in S, q \in Q$. Thus, when writing the GMDS problem in its weak correspondence form using (4) and (5), we search for $p : S \times Q \to \mathbb{R}^+$ which minimizes

$$\int_{S,Q} \left(d_S(\phi(q), s) - d_Q(q, \psi(s))\right)^2 da_s da_q$$
$$= \int_{S,Q} \left(\int_S d_S(s', s) p(q, s') da_{s'} \right.$$
$$\left. - \int_Q d_Q(q, q') p(q', s) da_{q'} \right)^2 da_s da_q. \quad (6)$$

In practice, we detect correspondences between sampled triangulated surfaces, for which we can re-write our optimization problem in matrix notation, so that it reads

$$\min_{\mathbf{P}} \|\mathbf{P}\mathbf{A}_S\mathbf{D}_S - \mathbf{D}_Q\mathbf{A}_Q\mathbf{P}\|_{S,Q}$$
s.t.
$$\mathbf{P}\mathbf{A}_S\mathbb{1} = \mathbb{1},$$
$$\mathbf{P}^T\mathbf{A}_Q\mathbb{1} = \mathbb{1}, \quad (7)$$

where $\|\mathbf{F}\|_{S,Q} = \text{trace}\left(\mathbf{F}^T\mathbf{A}_Q\mathbf{F}\mathbf{A}_S\right)$ represents the L_2 norm of the function $F : S \times Q \to \mathbb{R}$ discretised by the matrix $\mathbf{F}$. Here, $\mathbf{P}$, the discretization of $p(s, q)$, is the doubly stochastic weighted approximate-permutation matrix we are looking for. $\mathbf{A}_S$ and $\mathbf{A}_Q$ are diagonal area elements matrices, where $(\mathbf{A}_S)_{ii} \approx da_{s_i}$. In a triangulated surface, an area element da_{s_i} about a specific vertex $s_i \in S$ can be approximated by the area of the Voronoi cell about that vertex, as specified for example in Pinkall and Polthier (1993). Next, $\mathbf{D}_S$ and $\mathbf{D}_Q$ are the symmetric inter-geodesics matrices, such that $(\mathbf{D}_S)_{ij} = d_S(s_i, s_j)$, that is, the geodesic distance between points $s_i \in S$ and $s_j \in S$. The doubly stochastic requirement is captured by the first two constraints in which $\mathbb{1}$ is a vector of ones. As it was shown in Aflalo et al. (2015b), the solution of the matching problem is not influenced by omitting a non-negativity constraint $\mathbf{P} \geq \mathbf{0}$; instead, it is sufficient to relax $\mathbf{P}$ to the affine space defined by the two constrains in (7).

3 Spectral Formulation

In order to reduce the dimensionality of the problem, we propose to use a spectral representation of the permutation matrix $\mathbf{P}$. Unlike the functional maps approach (Ovsjanikov et al. 2012), we do not search for the linear map between spectral representations of corresponding features, but rather express the correspondence matrix $\mathbf{P}$ in the dual spectral domains.

Let $\mathbf{\Phi}_S$ be the matrix with columns given by the eigenvectors $\{\phi_i^S\}$ of the Laplace Beltrami operator (LBO) of S, and $\mathbf{\Lambda}_S$ its associated eigenvalues diagonal matrix, such that $\mathbf{W}_S\mathbf{\Phi}_S = \mathbf{A}_S\mathbf{\Phi}_S\mathbf{\Lambda}_S$. Here, w.l.o.g. we used the LBO discretization of Pinkall and Polthier (1993) by which a triangulated surface Beltrami operator is defined by $\mathbf{L} = \mathbf{A}^{-1}\mathbf{W}$, where again $(\mathbf{A})_{ii} \approx da_{s_i}$, and $\mathbf{W}$ represents the well established cotangent weights coefficients organized in a sparse laplacian matrix. Any other Laplace Beltrami operator could be handled in a similar manner. In the continuous setting, the fuzzy correspondence $p(s, q)$ is a mapping from $S \times Q$ to $\mathbb{R}$. For any given q, we can express $p(s, q)$ using the spectral domain of S by

$$p(s, q) = \sum_i \left\langle p(s, q), \phi_i^S(s) \right\rangle_S \phi_i^S(s) = \sum_i \alpha_i^S(q)\phi_i^S(s). \quad (8)$$

Now, for each i, the coefficient $\alpha_i^S(q)$ is a mapping (a function) $\alpha_i^S : Q \to \mathbb{R}$ that can be expressed in the spectral domain of Q as

$$\alpha_i^S(q) = \sum_j \left\langle \alpha_i^S(q), \phi_j^Q(q) \right\rangle_Q \phi_j^Q(q) = \sum_j \alpha_{ij}\phi_j^Q(q).$$

Plugging the last expression into Eq. (8) we have

$$p(s, q) = \sum_i \sum_j \alpha_{ij}\phi_j^Q(q)\phi_i^S(s),$$

or in matrix form

$$\mathbf{P} = \mathbf{\Phi}_Q\boldsymbol{\alpha}\mathbf{\Phi}_S^T, \quad (9)$$

which is the spectral representation of $\mathbf{P}$ with respect to $\mathbf{\Phi}_S$ and $\mathbf{\Phi}_Q$, that is captured by a matrix $\boldsymbol{\alpha}$.

A specific interesting mapping would be one where $Q = S$. In that case we write $\mathbf{P} : S \times S \to \mathbb{R}$ and express $\mathbf{P} = \mathbf{\Phi}_S\boldsymbol{\alpha}\mathbf{\Phi}_S^T$. It is obvious that for $\boldsymbol{\alpha} = \mathbf{I}$, the trivial identity mapping should hold. Now, what would happen if we consider just a few leading eigenvectors of $\mathbf{\Phi}_S$? The truncation effect on mapping surface points to their original locations using a different number of eigenvectors is captured in Fig. 1. The effect could be thought of as a low pass filter, affecting the maximum points of filtered delta functions, representing the original point locations. It appears from our experiments that 100 eigenfunctions are sufficient for accurate localization of the identity mapping. Figure 2 illustrates the accuracy of the mapping obtained for different numbers of eigenvectors. This accuracy is measured by geodesic distances between the original point locations and the locations of the maxima of the corresponding smoothed delta functions, weighted by the square root of the area of S, $\sqrt{\mathbf{A}_S}$.

The double stochastic conditions

$$\mathbf{P}\mathbf{A}_S\mathbb{1} = \mathbb{1} \text{ and } \mathbf{P}^T\mathbf{A}_Q\mathbb{1} = \mathbb{1},$$

can be rewritten as

$$\mathbf{\Phi}_Q\boldsymbol{\alpha}\mathbf{\Phi}_S^T\mathbf{A}_S\mathbb{1} = \mathbb{1} \text{ and } \mathbf{\Phi}_S\boldsymbol{\alpha}^T\mathbf{\Phi}_Q^T\mathbf{A}_Q\mathbb{1} = \mathbb{1}.$$

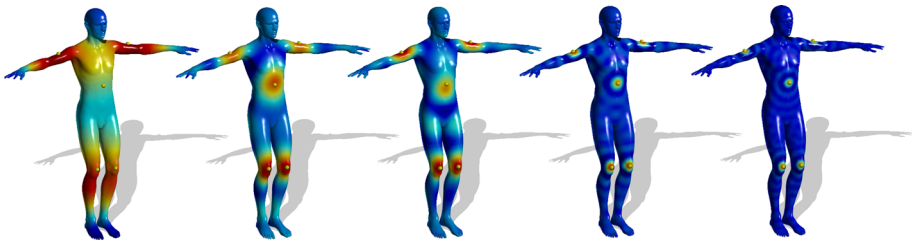

Fig. 1 Mapping 5 surface-points (indicated by *yellow spheres*) to their own location, using (from *left to right*) 10, 50, 100, 500 and 1000 eigenvectors of the Laplace–Beltrami operator, respectively (Color figure online)

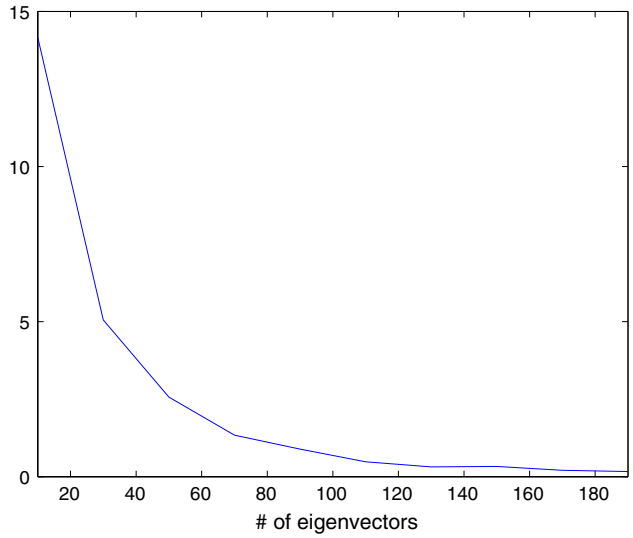

Fig. 2 Geodesic distance error between surface points and their mapping to themselves using 10–1000 eigenfunctions, averaged over 50 points randomly sampled from S

Setting $\boldsymbol{\eta}_S = \boldsymbol{\Phi}_S^T \mathbf{A}_S \mathbb{1}$ and similarly $\boldsymbol{\eta}_Q = \boldsymbol{\Phi}_Q^T \mathbf{A}_Q \mathbb{1}$, the double stochastic conditions can be compactly written as

$$\boldsymbol{\alpha}\boldsymbol{\eta}_S = \boldsymbol{\eta}_Q \text{ and } \boldsymbol{\alpha}^T \boldsymbol{\eta}_Q = \boldsymbol{\eta}_S. \tag{10}$$

Putting all ingredients together, we consider the spectral representation of Eq. (7) by which $\mathbf{P} = \boldsymbol{\Phi}_Q \boldsymbol{\alpha} \boldsymbol{\Phi}_S^T$, and $\boldsymbol{\alpha}\boldsymbol{\eta}_S = \boldsymbol{\eta}_Q$ and $\boldsymbol{\alpha}^T \boldsymbol{\eta}_Q = \boldsymbol{\eta}_S$. In order to make the spectral story complete, we still have to handle the inter-geodesic distance maps and represent them in the spectral domain.

4 Correspondence in Spectral Domain

The correspondence $\mathbf{P}$ may be written in the spectral bases of S and Q, up to area normalization. Thus, following the analysis in Sect. 3, we may write

$$\mathbf{P} = \boldsymbol{\Phi}_Q \boldsymbol{\alpha} \boldsymbol{\Phi}_S^T. \tag{11}$$

One of the important consequences of using the spectral representation of the correspondence is a reduction of the size of the problem. We started by searching for a point-wise matching between the vertices of S and those of Q, with $\mathbf{P}$ of size $|S| \times |Q|$, where $|S|$ is the number of points in S (vertices in case of a triangulated surface). Now, working in the spectral domains we consider the matrix $\boldsymbol{\alpha}$ relating the bases $\boldsymbol{\Phi}_S$ and $\boldsymbol{\Phi}_Q$. The size of $\boldsymbol{\alpha}$ relates to the size of the truncated spectral domains. In our experiments, for a given surface sampled at $n = 10{,}000$ vertices, it was enough to consider $m = 100$ eignefunctions for a faithful reconstruction of $\mathbf{P}$. Note, that we do not require $|S| = |Q|$, and two differently sampled surfaces and any high order approximation of the maximal point would allow us to find its position inside a triangle approximating the surface, rather than restricting the mapping to an existing vertex.

Next, similar to the spectral expression of $\mathbf{P}$, we could express $\mathbf{D} : S \times S \to \mathbb{R}$ in the natural eigenspace of S, namely $\boldsymbol{\Phi}_S$. The inter-geodesic distance matrix $\mathbf{D}_S$ could then be written as

$$\mathbf{D}_S = \boldsymbol{\Phi}_S \boldsymbol{\beta}_S \boldsymbol{\Phi}_S^T, \tag{12}$$

where

$$\boldsymbol{\beta}_S = \boldsymbol{\Phi}_S^T \mathbf{A}_S \mathbf{D}_S \mathbf{A}_S \boldsymbol{\Phi}_S,$$

or in its continuous setting

$$\boldsymbol{\beta}_{ij} = \int_{S \times S} d_S(s, s') \phi_i(s) \phi_j(s') da_s da_{s'}.$$

Truncating the eigenspace in the above expression to $m \ll n$ eigenfunctions instead of n, would lead to an approximation of the distance function $\mathbf{D}_S$ that we denote by $\tilde{\mathbf{D}}_S$.

In fact, in order to compute an approximation to $\boldsymbol{\beta}_S$, there is no need to compute all inter-geodesic distances. By substantially sub-sampling the matrix $\mathbf{D}_S$ we could solve for $\boldsymbol{\beta}_S$ and keep the interpolated geodesic distances matrix in its implicit form. Here, following the ideas put forward in Aflalo and Kimmel (2013) we use the biharmonic equation for the interpolation. It eliminates the need to compute the $O(n^2)$ components of $\mathbf{D}_S$ to almost linear space (and time) complexity in n. The motivation for that efficient construction is that the global structure of $\mathbf{D}_S$ is captured by the geodesic distances from the sampled points to the rest of the domain, while the local smooth structure is naturally interpolated by the leading eigenfunctions of $\boldsymbol{\Phi}_S$, which can be proven to

be optimally tailored for the representation task of $\mathbf{D}_S$, see Aflalo et al. (2015a).

We are now ready to reformulate Problem (7) in the spectral domain, that now reads,

$$\begin{aligned}&\max_{\mathbf{P}} \|\mathbf{P}\mathbf{A}_S\mathbf{D}_S - \mathbf{D}_Q\mathbf{A}_Q\mathbf{P}\|_{S,Q} \\ &\quad\approx \max_{\mathbf{P}} \|\mathbf{P}\mathbf{A}_S\tilde{\mathbf{D}}_S - \tilde{\mathbf{D}}_Q\mathbf{A}_Q\mathbf{P}\|_{S,Q}\end{aligned} \tag{13}$$

Morevover, we have,

$$\mathbf{P}\mathbf{A}_S\tilde{\mathbf{D}}_S = \boldsymbol{\Phi}_Q\boldsymbol{\alpha}\boldsymbol{\Phi}_S^T\mathbf{A}_S\boldsymbol{\Phi}_S\boldsymbol{\beta}_S\boldsymbol{\Phi}_S^T = \boldsymbol{\Phi}_Q\boldsymbol{\alpha}\boldsymbol{\beta}_S\boldsymbol{\Phi}_S^T$$

and, similarly,

$$\tilde{\mathbf{D}}_Q\mathbf{A}_Q\mathbf{P} = \boldsymbol{\Phi}_Q\boldsymbol{\beta}_Q\boldsymbol{\alpha}\boldsymbol{\Phi}_S^T.$$

It is also straightforward to see that

$$\begin{aligned}\|\boldsymbol{\Phi}_Q\mathbf{F}\boldsymbol{\Phi}_S^T\|_{SQ} &= \text{trace}\left((\boldsymbol{\Phi}_Q\mathbf{F}\boldsymbol{\Phi}_S^T)^T\mathbf{A}_Q(\boldsymbol{\Phi}_Q\mathbf{F}\boldsymbol{\Phi}_S^T)\mathbf{A}_S\right) \\ &= \text{trace}\left((\boldsymbol{\Phi}_S\mathbf{F}^T\boldsymbol{\Phi}_Q^T\mathbf{A}_Q(\boldsymbol{\Phi}_Q\mathbf{F}\boldsymbol{\Phi}_S^T)\mathbf{A}_S\right) \\ &= \text{trace}\left((\mathbf{F}^T(\boldsymbol{\Phi}_Q^T\mathbf{A}_Q\boldsymbol{\Phi}_Q)\mathbf{F}(\boldsymbol{\Phi}_S^T\mathbf{A}_S\boldsymbol{\Phi}_S)\right) \\ &= \text{trace}\left(\mathbf{F}^T\mathbf{F}\right) \\ &= \|\mathbf{F}\|_F^2.\end{aligned}$$

Then, we have

$$\|\mathbf{P}\mathbf{A}_S\tilde{\mathbf{D}}_S - \tilde{\mathbf{D}}_Q\mathbf{A}_Q\mathbf{P}\|_{S,Q} = \|\boldsymbol{\alpha}\boldsymbol{\beta}_S - \boldsymbol{\beta}_Q\boldsymbol{\alpha}\|_F^2.$$

Finally, plugging all ingredients together we are lead to an optimization problem

$$\begin{aligned}&\min_{\boldsymbol{\alpha}} \ \|\boldsymbol{\alpha}\boldsymbol{\beta}_S - \boldsymbol{\beta}_Q\boldsymbol{\alpha}\|_F^2 \\ &\text{s.t.} \\ &\qquad \boldsymbol{\alpha}\boldsymbol{\eta}_S = \boldsymbol{\eta}_Q \quad \text{and} \quad \boldsymbol{\alpha}^T\boldsymbol{\eta}_Q = \boldsymbol{\eta}_S.\end{aligned} \tag{14}$$

The first eigenvector of the Laplace–Beltrami operator is the constant vector $\phi_1 = \tau(1\,1\,1\,1\,1\cdots 1)^T$, where τ is a scalar constant. The orthonormality of the basis $\boldsymbol{\Phi}$ by which $\boldsymbol{\Phi}^T\mathbf{A}\boldsymbol{\Phi} = I$ allows us to write $\phi_1^T\mathbf{A}\phi_1 = 1$, from which τ can be extracted to be $\tau = \left(\sum_i \mathbf{A}_{ii}\right)^{-1/2}$. By the same orthonormality property, for any eigenvector ϕ_j, $j > 1$, we have $\phi_1^T\mathbf{A}\phi_j = 0$. Thus,

$$\boldsymbol{\eta} = \boldsymbol{\Phi}^T\mathbf{A}\mathbb{1} = \boldsymbol{\Phi}^T\mathbf{A}\,\tau^{-1}\phi_1 = \tau^{-1}\begin{pmatrix}1\\0\\\vdots\\0\end{pmatrix}.$$

When the shapes are almost isometric and can be assumed to have similar area, the double stochastic constraints can be directly enforced by

$$\boldsymbol{\alpha}\begin{pmatrix}1\\0\\\vdots\\0\end{pmatrix} = \begin{pmatrix}1\\0\\\vdots\\0\end{pmatrix} \text{ and } \boldsymbol{\alpha}^T\begin{pmatrix}1\\0\\\vdots\\0\end{pmatrix} = \begin{pmatrix}1\\0\\\vdots\\0\end{pmatrix},$$

or equivalently

$$\boldsymbol{\alpha} = \begin{pmatrix}1 & 0 & \cdots & 0\\ 0 & \alpha_{22} & \cdots & \alpha_{2m}\\ \vdots & \vdots & \vdots & \vdots\\ 0 & \alpha_{m2} & \cdots & \alpha_{mm}\end{pmatrix}.$$

Finally, in order to solve Problem (14) numerically, we use the PBM toolbox by M. Zibulevsky (Ben-Tal and Zibulevsky 1997).

5 Additional Constraints

In order to guarantee convergence to significant solutions one usually adds constraints to the optimization problem (14). Let us consider some popular alternatives. First, region correspondences could be applied by pre-segmentation of the shapes using tools like the MSER (Ovsjanikov et al. 2012; Pokrass et al. 2013). Another option is to start with sparse point to point correspondences, for example, by considering the local maxima of the geodesic spectrum (Baloch et al. 2005) and finding correspondences at few special points using descriptors such as WKS (Aubry et al. 2011) and HKS (Sun et al. 2009). In this paper, we follow the latter approach. Given a shape S, at each point $s \in S$, we compute the integral of the geodesic distances between s and the rest of the points in S. That is,

$$g(s) = \int_S D(s,s')da(s').$$

This quantity can be efficiently evaluated using the spectral approximation of the geodesic distances described in Sect. 4. That is,

$$g(s) = \mathbf{D}_S\mathbf{A}_S\mathbb{1} \approx \boldsymbol{\Phi}_S\boldsymbol{\beta}_S\boldsymbol{\Phi}_S^T\mathbf{A}_S\mathbb{1}.$$

Identifying local maxima of $g(s)$ provides a unique set of feature-points; usually up to 5 points are enough for each surface. Then, exhaustive search for correspondences could be applied between the two sets.

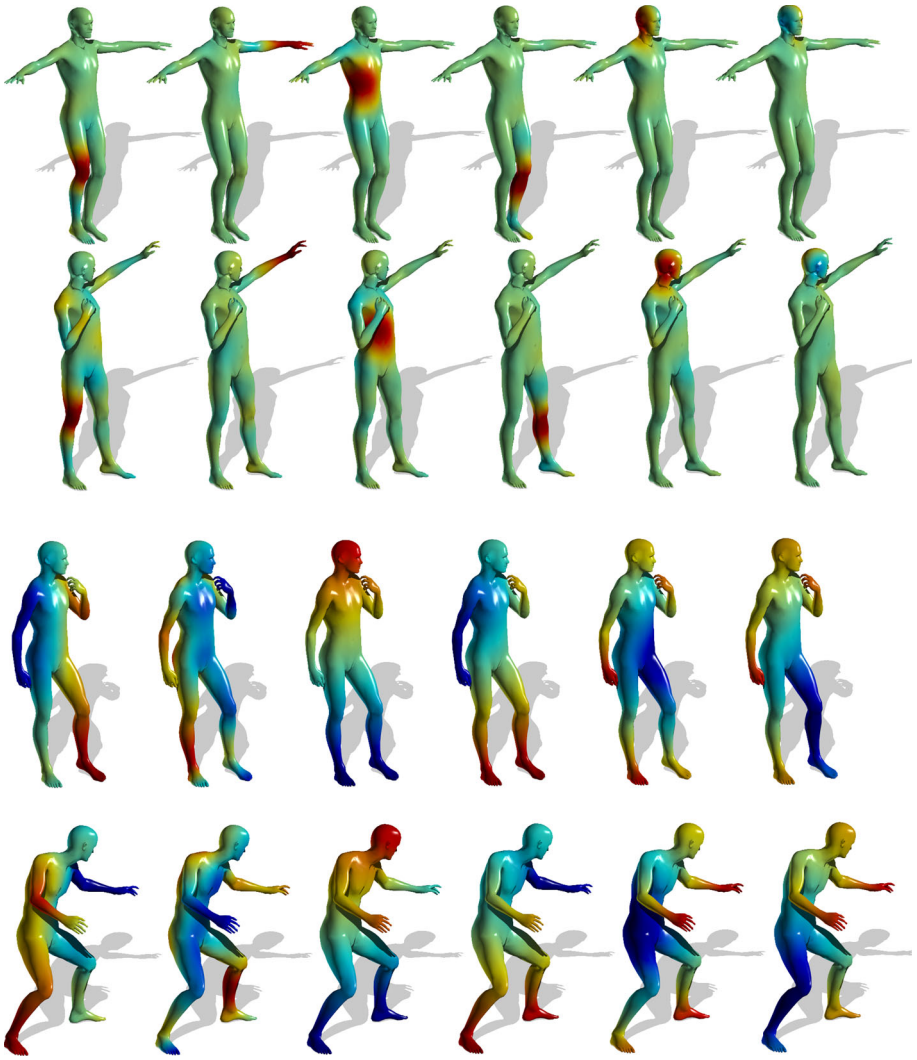

Fig. 3 Mapping functions between two almost isometric shapes via SGMDS (Color figure online)

6 Experimental Results

Several experiments were performed in order to evaluate the accuracy and efficiency of the proposed method. We used two publicly available datasets—TOSCA (Bronstein et al. 2008) and SCAPE (Anguelov et al. 2004). The TOSCA dataset contains 90 densely sampled synthetic human and animal surfaces, divided into several classes with given point-to-point correspondences between the shapes within each class. The SCAPE dataset contains scans of real human bodies in different poses.

In all of our experiments, we used pre-computed geodesic distances between a small subset of surface points. The geodesic distances were calculated using the fast marching method between 5 % of the surface points, sampled using the farthest point sampling method (Hochbaum and Shmoys 1985; Gonzalez 1985). To minimize the objective function in Eq. (14) we used the PBM toolbox by M. Zibulevsky (Ben-Tal and Zibulevsky 1997). All the experiments were executed on a 2.7 GHz Intel Core i7 machine with 16GB RAM.

In our first experiment, we selected almost isometric surfaces within the same class from the TOSCA dataset, and computed correspondences between them using the proposed Spectral-GMDS algorithm. We visualize the quality of the mapping by transferring a couple of functions defined on one shape to the other, as shown in Figs. 3 and 4. In Fig. 5 we visualize point-to-point correspondences between several almost isometric poses of a horse, obtained using the S-GMDS. Note, that for intrinsically symmetric shapes, like the ones used in our experiments, there can be more than one optimal mapping. In our examples, there exist two possible correspondences between the surface to itself, and thus, between each object and its appearance in other poses. The objective function in Eq. (14) does not distinguish between such symmetries. In such cases, the algorithm is expected to detect one of the possible mappings, and may therefore

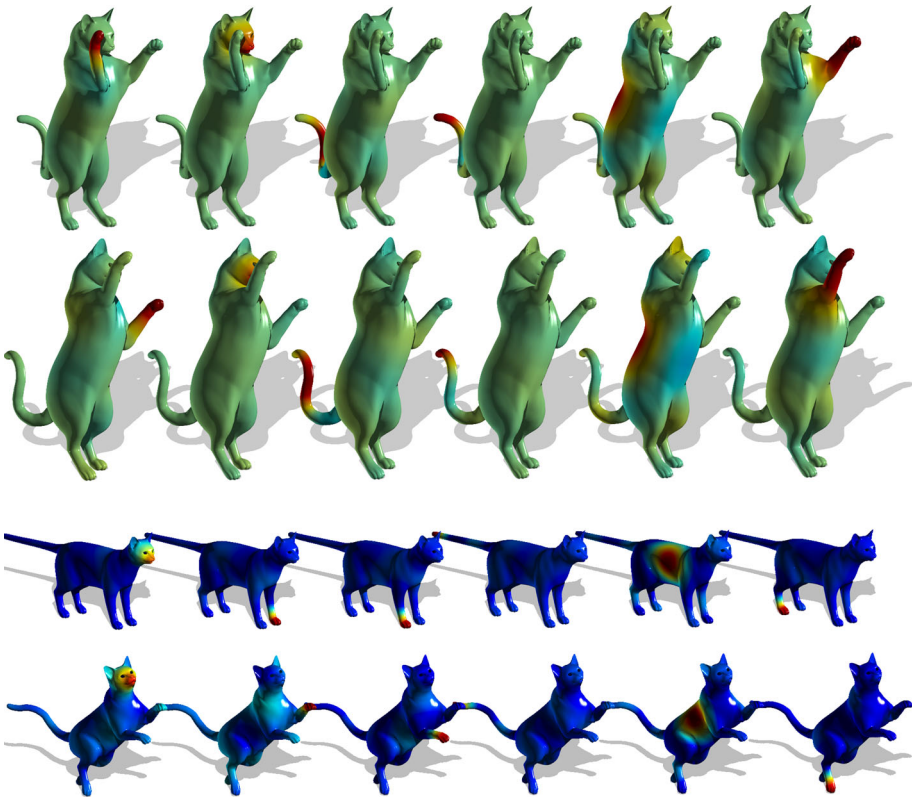

Fig. 4 Mapping functions between two almost isometric shapes via SGMDS (Color figure online)

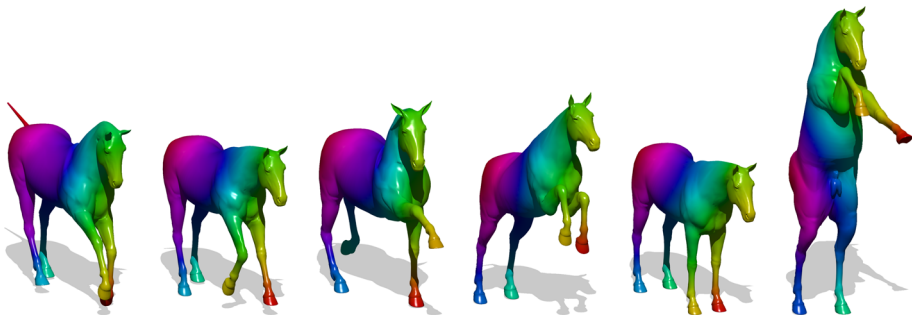

Fig. 5 Dense point-to-point correspondence between six almost isometric shapes of a horse from the TOSCA dataset (Color figure online)

produce symmetric flips, as shown in Fig. 3 (bottom), Fig. 4 (top), and Fig. 6 (bottom).

Figure 7 compares the accuracy of the proposed method to other methods using the evaluation procedure proposed in Kim et al. (2011). In these experiment we used 6 first eigenvectors of the Laplace–Beltrami operator. The evaluation protocol was applied to both TOSCA (Bronstein et al. 2008) and SCAPE (Anguelov et al. 2004) datasets. For the other methods, we used the information provided in Kim et al. (2011), and in Ovsjanikov et al. (2012); Pokrass et al. (2013). Figure 6 demonstrates the robustness of the proposed approach to typical types of noise.

In the benchmark protocol proposed by Kim et al. (Kim et al. 2011), the ground-truth correspondence between shapes is assumed to be given. Then, a script, provided by the authors, computes the geodesic departure of each point, mapped by the evaluated method, from what the authors refer to as true location. The distortion curves describe the percentage of surface points falling within a relative geodesic distance from what are assumed to be their true locations. For each shape, the geodesic distance is normalized with respect to the shape's squared root of the area.

Notice that "true location" here is actually a subjective measure. In fact, measuring the geodesic distortion of the given correspondences demonstrates a substantial discrepancy between corresponding pairs of points on most surface pairs from the given datasets. The distortion curves incorporate an intrinsic ambiguity of about 5–25 % due to the lack of absolute isometry within seemingly identical objects at different poses. In a favorable scenario, given a

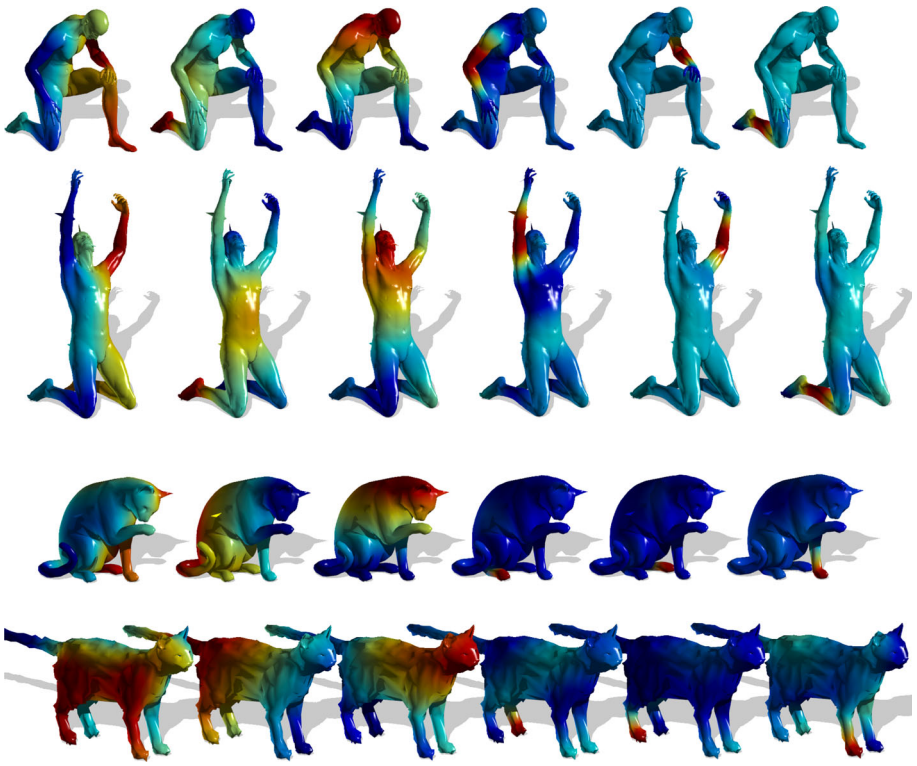

Fig. 6 Mapping functions between two, almost isometric, noisy shapes via SGMDS (Color figure online)

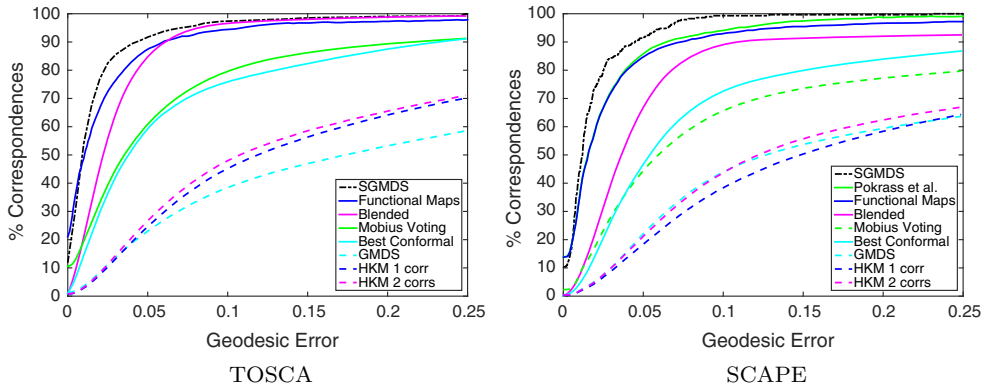

Fig. 7 Quantitative evaluation of the SGMDS applied to the shapes from the TOSCA and SCAPE datasets, using the protocol from Kim et al. (2011) (Color figure online)

pair of shapes for which the corresponding geodesic distortion is relatively small, the SGMDS provides superior results compared to existing methods, as demonstrated in Fig. 8.

Figure 9 illustrates the accuracy of the mapping found by the proposed SGMDS method. A set of random points on the left shape are mapped to the right one. Each shape was then colored according to the Voronoi regions of the points. Note that the Voronoi Diagram was generated separately for each shape after mapping of the points. Given the substentially different poses, the accuracy of the mapping compared to the existing alternatives is relatively high. Next, Fig. 10 presents a quantitative evaluation of the SGMDS using a different number of eigenvectors for each experiment. Note that in this experiment, increasing the number of the eigenvectors beyond 100 does not provide a significant improvement in the matching accuracy.

7 Conclusions

Spectral generalized multidimensional scaling method (S-GMDS) was proposed and proven to be an accurate model and efficient tool for matching non-rigid shapes. It accounts

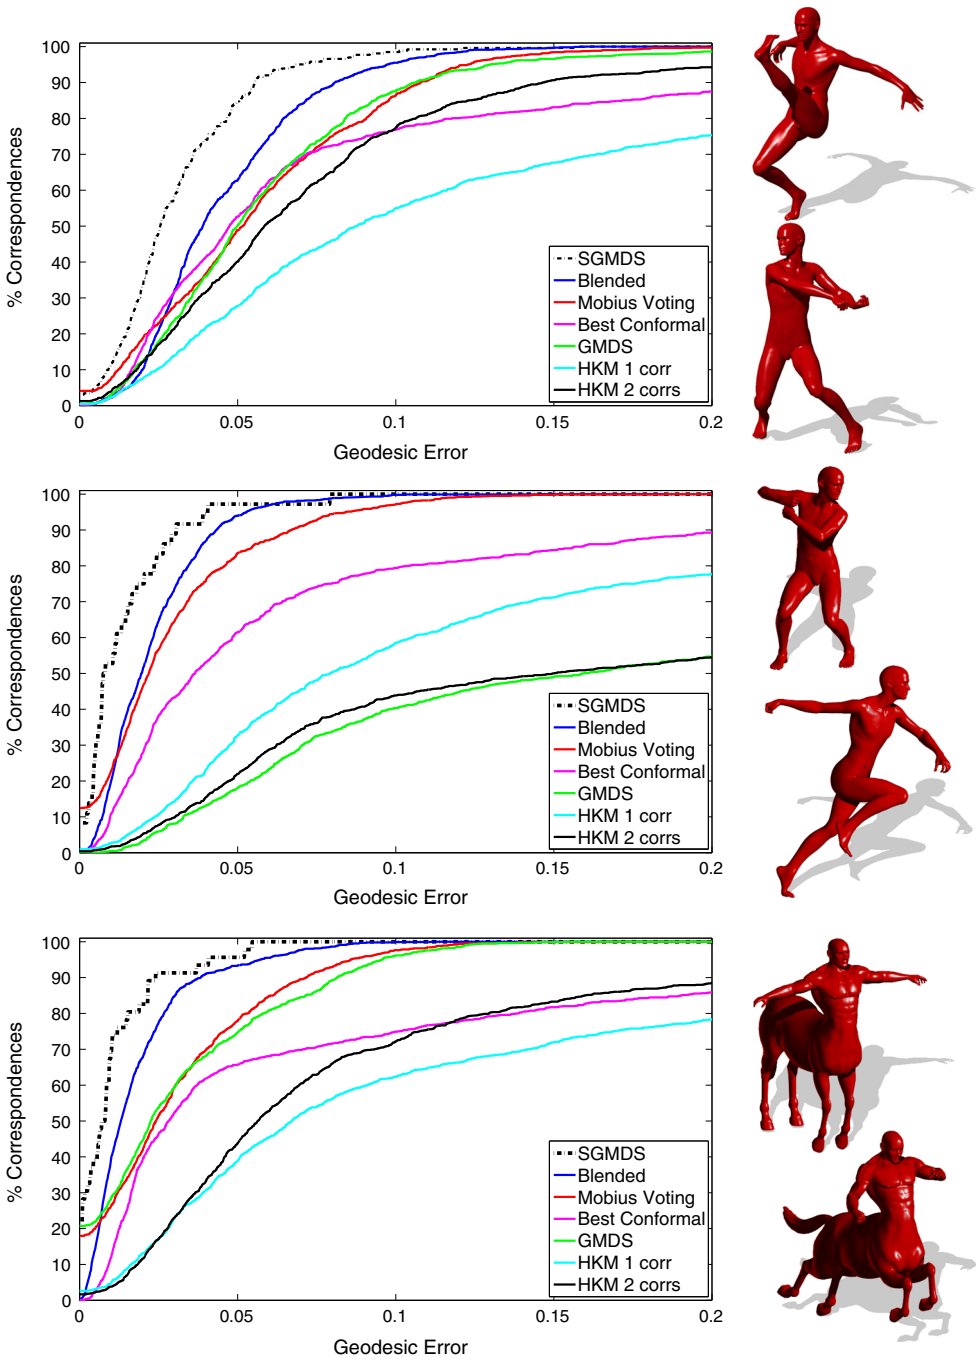

Fig. 8 Performance evaluation of the S-GMDS compared to other methods applied to pairs of "David" and "Centaur" shapes from the TOSCA dataset which are relatively isometric. See *shapes* on the *right*. The comparison protocol is adopted from Kim et al. (2011) (Color figure online)

for almost isometric deformations of surfaces with respect to the regular metric. Being able to consider large as well as small distances when comparing two surfaces, with a natural regularization of the matching, reduces the need for the support of heuristics or initializations. By incorporating the smoothness of the inter-geodesic distances matrix for subsampling and then implicitly interpolating in that domain, we treat the shape matching problem holistically rather than as an interpolation between multiple matched features such as points, regions, or localized functions.

Here, we used a regular metric in which geodesic distances on the surface determine the isometric quantity we try to preserve and whose distortions we use as a discrepancy measure. In our future research, we will try to axiomatically tackle the problem of analyzing objects between which local scale can be a notable factor, yet, the conceptual meaning of such local structures with different scale is preserved. For example, utilizing a scale invariant metric (Aflalo et al. 2013) could allow us to handle shapes with similar structures that vary in their relative size, like a small face with

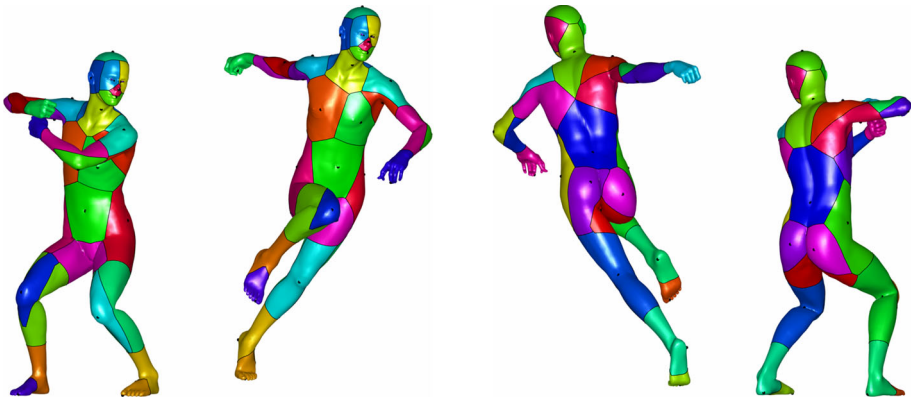

Fig. 9 SGMDS mapping: Visualization of a Voronoi diagram on one shape and its corresponding map on a different pose of the articulate object as computed by the propsoed method (Color figure online)

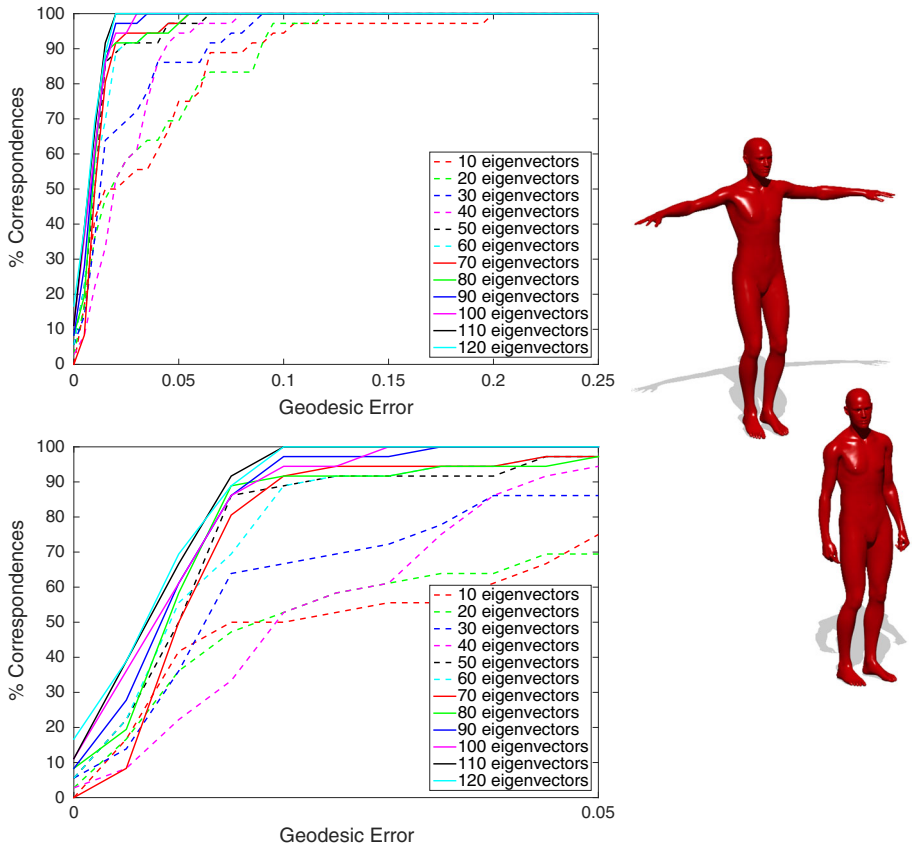

Fig. 10 *Top* quantitative evaluation presented as rate distortion graphs of the SGMDS applied to two human shapes from the TOSCA dataset (shown in *red*), using a different number of eigenvectors in each experiment. *Bottom* Zoom in into the rate distortion graph for geodesic errors in the range of [0, 0.05]. The comparison protocol is adopted from Kim et al. (2011) (Color figure online)

large ears that we try to match to a large face with smaller ears.

Moreover, when topological noise is present, a robust matching that could gracefully handle cuts and holes, can be achieved by using diffusion or commute time distances (Coifman and Lafon 2006; Bronstein and Bronstein 2011; Bronstein et al. 2010), instead of the geodesc distance, in the proposed framework. In fact, when dealing with diffusion distances, it is natural to approximate the distance matrices in the spectral domain.

Acknowledgments The authors would like to thank Alon Shtern and Matan Sela for stimulating discussions throughout this research. This work has been supported by Grant agreement No. 267414 of the European Communitys FP7-ERC program.

Compliance with ethical standards

Conflicts of interest The authors declare that they have no conflict of interest.

References

Aflalo, Y., & Kimmel, R. (2013). Spectral multidimensional scaling. *Proceedings of the National Academy of Sciences*, *110*(45), 18,052–18,057.

Aflalo, Y., Kimmel, R., & Raviv, D. (2013). Scale invariant geometry for nonrigid shapes. *SIAM Journal on Imaging Sciences*, *6*(3), 1579–1597.

Aflalo, Y., Brezis, H., & Kimmel, R. (2015a). On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences*, *8*(2), 1141–1160.

Aflalo, Y., Bronstein, A., & Kimmel, R. (2015b). On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, *112*(10), 2942–2947. doi:10.1073/pnas.1401651112.

Anguelov, D., Srinivasan, P., Pang, H. C., Koller, D., Thrun, S., & Davis, J. (2004). The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. *Advances in Neural Information Processing Systems*, *17*, 33–40.

Aubry, M., Schlickewei, U., & Cremers, D. (2011). The wave kernel signature: A quantum mechanical approach to shape analysis. In: *2011 IEEE international conference on computer vision workshops (ICCV workshops)*, IEEE, pp 1626–1633.

Baloch, S., Krim, H., Kogan, I., & Zenkov, D. (2005). Rotation invariant topology coding of 2d and 3d objects using morse theory. In *IEEE international conference on image processing, 2005. ICIP 2005* (Vol. 3, pp. III-796–III-799).

Ben-Tal, A., & Zibulevsky, M. (1997). Penalty/barrier multiplier methods for convex programming problems. *SIAM Journal on Optimization*, *7*(2), 347–366.

Bérard, P., Besson, G., & Gallot, S. (1994). Embedding riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, *4*(4), 373–398.

Besl, P. J., & McKay, N. D. (1992). Method for registration of 3-d shapes. In *Robotics-DL tentative. International Society for Optics and Photonics*, pp. 586–606.

Borg, I., & Groenen, P. (1997). *Modern multidimensional scaling: Theory and applications*. New York: Springer.

Bronstein, A., Bronstein, M., & Kimmel, R. (2008). *Numerical geometry of non-rigid shapes*. New York: Springer.

Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2006). Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences of the USA*, *103*(5), 1168–1172.

Bronstein, A. M., Bronstein, M. M., Kimmel, R., Mahmoudi, M., & Sapiro, G. (2010). A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision*, *89*(2–3), 266–286. doi:10.1007/s11263-009-0301-6.

Bronstein, M., & Bronstein, A. M. (2011). Shape recognition with spectral distances. *IEEE Transaction on Pattern Analysis and Machine Intelligence (PAMI)*, *33*(5), 1065–1071.

Burago, D., Burago, Y., & Ivanov, S. (2001). *A course in metric geometry*. Providence: American Mathematical Society.

Chen, Y., & Medioni, G. (1991). Object modeling by registration of multiple range images. In *Proceedings. IEEE international conference on robotics and automation*, pp. 2724–2729.

Coifman, R. R., & Lafon, S. (2006). Diffusion maps. *Applied and Computational Harmonic Analysis*, *21*(1), 5–30.

Dubrovina, A., & Kimmel, R. (2010). Matching shapes by eigendecomposition of the laplace_belrami operator. In *Proceedings of the symposium on 3D data processing visualization and transmission (3DPVT)*.

Gębal, K., Bærentzen, J. A., Aanæs, H., & Larsen, R. (2009). Shape analysis using the auto diffusion function. *Computer Graphics Forum, Wiley Online Library*, *28*, 1405–1413.

Gonzalez, T. F. (1985). Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, *38*, 293–306.

Gromov, M. (1981). Structures metriques pour les varietes riemannienes. *Textes Mathematiques*, no. 1.

Gu, X., Wang, Y., Chan, T. F., Thompson, P. M., & Yau, S. T. (2004). Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, *23*(8), 949–958.

Hochbaum, D., & Shmoys, D. (1985). A best possible heuristic for the k-center problem. *Mathematics of Operations Research*, *10*, 180–184.

Jin, M., Wang, Y., Yau, S. T., & Gu, X. (2004). Optimal global conformal surface parameterization. *IEEE Visualization*, pp. 267–274.

Kim, V. G., Lipman, Y., & Funkhouser, T. (2011). Blended intrinsic maps. In *ACM Transactions on Graphics (TOG), ACM* (Vol. 30, p. 79).

Lévy, B. (2006). Laplace-Beltrami eigenfunctions towards an algorithm that "understands" geometry. In *IEEE International Conference on Shape Modeling and Applications, 2006. SMI 2006*, pp. 13–13.

Lipman, Y., & Daubechies, I. (2011). Surface comparison with mass transportation. *Advances in Mathematics*, *227*(3)

Lipman, Y., & Funkhouser, T. (2009). Möbius voting for surface correspondence. *ACM Transactions on Graphics (Proc SIGGRAPH)*, *28*(3), 72.

Mateus, D., Horaud, R., Knossow, D., Cuzzolin, F., & Boyer, E. (2008). Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *IEEE Conference on Computer Vision and Pattern Recognition*. CVPR 2008. pp. 1–8.

Memoli, F. (2007). On the use of Gromov-Hausdorff distances for shape comparison. In M. Botsch, R. Pajarola, B. Chen, & M. Zwicker (Eds.), *Symposium on point based graphics* (pp. 81–90). Prague: Czech Republic, Eurographics Association.

Memoli, F., & Sapiro, G. (2005). A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, *5*(3), 313–347.

Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., & Guibas, L. (2012). Functional maps: A flexible representation of maps between shapes. *ACM Transaction on Graph*, *31*(4), 30:1–30:11.

Pinkall, U., & Polthier, K. (1993). Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, *2*(1), 15–36.

Pokrass, J., Bronstein, A. M., Bronstein, M. M., Sprechmann, P., & Sapiro, G. (2013). Sparse modeling of intrinsic correspondences. *Computer Graphics Forum (EUROGRAPHICS)*, *32*, 459–468.

Raviv, D., Dubrovina, A., & Kimmel, R. (2012). Hierarchical matching of non-rigid shapes. In *Scale space and variational methods in computer vision*, Springer, Berlin, pp. 604–615.

Rustamov, R., Ovsjanikov, M., Azencot, O., Ben-Chen, M., Chazal, F., & Guibas, L. (2013). Map-based exploration of intrinsic shape differences and variability. In *SIGGRAPH, ACM*.

Sahillioğlu, Y., & Yemez, Y. (2011). Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum, Wiley Online Library*, *30*, 1461–1470.

Schwartz, E. L., Shaw, A., & Wolfson, E. (1989). A numerical solution to the generalized mapmaker's problem: Flattening nonconvex polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *11*(9), 1005–1008.

Shtern, A., & Kimmel, R. (2014). Iterative closest spectral kernel maps. In *2nd International Conference on 3D Vision (3DV)* (Vol. 1, pp. 499–505).

Shtern, A., & Kimmel, R. (2015). Spectral gradient fields embedding for nonrigid shape matching. *Computer Vision and Image Understanding*, *140*, 21–29.

Sun, J., Ovsjanikov, M., & Guibas, L. (2009). A concise and provably informative multi-scale signature based on heat diffusion. In *Proceedings of the symposium on geometry processing*, Eurographics Association, Aire-la-Ville, Switzerland, SGP'09, pp. 1383–1392.

Zaharescu, A., Boyer, E., Varanasi, K., & Horaud, R. (2009). Surface feature detection and description with applications to mesh matching. In *IEEE conference on computer vision and pattern recognition, 2009*. CVPR 2009, pp. 373–380.

Zeng, W., Lui, L. M., Luo, F., Chan, T. F. C., Yau, S. T., & Gu, D. X. (2012). Computing quasiconformal maps using an auxiliary metric and discrete curvature flow. *Numerische Mathematik*, *121*(4), 671–703.

Zeng, Y., Wang, C., Wang, Y., Gu, X., Samaras, D., & Paragios, N. (2010). Dense non-rigid surface registration using high-order graph matching. In *IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 382–389.