SIAM J. IMAGING SCIENCES
Vol. 8, No. 2, pp. 1141–1160

On the Optimality of Shape and Data Representation in the Spectral Domain*

Yonathan Aflalo†, Haim Brezis‡, and Ron Kimmel†

Abstract. A proof of the optimality of the eigenfunctions of the Laplace–Beltrami operator (LBO) in representing smooth functions on surfaces is provided and adapted to the field of applied shape and data analysis. It is based on the Courant–Fischer min-max principle adapted to our case. The theorem we present supports the new trend in geometry processing of treating geometric structures by using their projection onto the leading eigenfunctions of the decomposition of the LBO. Utilization of this result can be used for constructing numerically efficient algorithms to process shapes in their spectrum. We review a couple of applications as possible practical usage cases of the proposed optimality criteria. We refer to a scale invariant metric, which is also invariant to bending of the manifold. This novel pseudometric allows constructing an LBO by which a scale invariant eigenspace on the surface is defined. We demonstrate the efficiency of an intermediate metric, defined as an interpolation between the scale invariant and the regular one, in representing geometric structures while capturing both coarse and fine details. Next, we review a numerical acceleration technique for classical scaling, a member of a family of flattening methods known as multidimensional scaling (MDS). There, the optimality is exploited to efficiently approximate all geodesic distances between pairs of points on a given surface and thereby match and compare between almost isometric surfaces. Finally, we revisit the classical principal component analysis (PCA) definition by coupling its variational form with a Dirichlet energy on the data manifold. By pairing the PCA with the LBO we can efficiently handle cases that go beyond the scope defined by the observation set that is handled by regular PCA.

Key words. Laplace–Beltrami, shape analysis, principal component analysis

AMS subject classification. 43A77

DOI. 10.1137/140977680

1. Introduction.

The field of shape analysis has been evolving rapidly in recent decades. The constant increase in computing power has allowed image and shape understanding algorithms to efficiently handle difficult problems that could not have been practically addressed in the past. A large set of theoretical tools from metric geometry, differential geometry, and spectral analysis has been imported and translated into action within the shape and image understanding arena. Among the myriad operators recently explored, the *Laplace–Beltrami operator* (LBO) is ubiquitous. The LBO is an extension of the Laplacian to nonflat multidimensional manifolds. Its properties have been well studied in differential geometry, and it

*Received by the editors July 15, 2014; accepted for publication (in revised form) March 16, 2015; published electronically May 21, 2015. This work was supported by the European Community's FP7- ERC program, grant agreement 267414.

http://www.siam.org/journals/siims/8-2/97768.html

†Computer Science Department, Technion, Israel Institute of Technology, Haifa 32000, Israel (yaflalo@cs.technion.ac.il, ron@cs.technion.ac.il).

‡Department of Mathematics, Rutgers University, Piscataway, NJ 08854, and Department of Mathematics, Technion, Israel Institute of Technology, Haifa 32000, Israel (brezis@math.rutgers.edu). This author's research was partially supported by NSF grant DMS-1207793 and by grant 238702 of the European Commission (ITN, project FIRST).

was used extensively in computer graphics. It is used to define the *heat equation*, which models the conduction of heat in solids, and is fundamental in describing basic physical phenomena. In its more general setting, the LBO admits an eigendecomposition that yields a *spectral domain* that can be viewed as a generalization of the Fourier analysis to any Riemannian manifold. The LBO invariance to isometric transformations allowed the theories developed by physicists and mathematicians to be useful for modern shape analysis. Here, we justify the selection of the leading eigenfunctions in the spectral domain as an optimal subspace for representing smooth functions on a given manifold. It is used for solving and accelerating existing solvers of various problems in data representation, information processing, and shape analysis. As one example, in section 4 we pose the dilemma of metric selection for shape representation while interpolating between a scale invariant metric and the regular one. Next, in section 5 it is shown how the recently introduced spectral classical scaling can benefit from the efficacy property of the suggested subspace. Finally, in section 6 we revisit the definition of the celebrated *principal component analysis* (PCA) by regularizing its variational form with an additional Dirichlet energy. The idea is to balance between two optimal subspaces, one for the data points themselves, captured by the PCA, and one optimally encapsulating the relation between the data points as defined by decomposition of the LBO.

2. Notation and motivation. Consider a parametrized surface $S : \Omega \subset \mathbb{R}^2 \to \mathbb{R}^3$ (with or without boundary) and a metric (g_{ij}) that defines the affine differential relation of navigating with coordinates $\{\sigma_1, \sigma_2\}$ in Ω to a distance measured on S. That is, an arc length on S expressed by σ_1 and σ_2 would read as $ds^2 = g_{11}d\sigma_1^2 + 2g_{12}d\sigma_1 d\sigma_2 + g_{22}d\sigma_2^2$. The LBO acting on the scalar function $f : S \to \mathbb{R}$ is defined as

$$-\Delta_g f = \frac{1}{\sqrt{g}} \sum_{ij} \partial_i \left(\sqrt{g} g^{ij} \partial_j f \right), \tag{2.1}$$

where g is the determinant of the metric matrix, and $(g^{ij}) = (g_{i,j})^{-1}$ is the inverse metric, while ∂_i is a derivative with respect to the ith coordinate σ_i. The LBO operator is symmetric and admits a spectral decomposition (λ_i, ϕ_i), with $\lambda_1 \leq \lambda_2 \leq \cdots$, such that

$$\begin{aligned} \Delta_g \phi_i &= \lambda_i \phi_i, \\ \langle \phi_i, \phi_j \rangle &= \delta_{ij}, \end{aligned} \tag{2.2}$$

where $\langle u, v \rangle = \int_S uv\sqrt{g}dx$, and $\|u\|_2^2 = \int_S |u|^2 \sqrt{g}dx$. In case S has a boundary, we add Neumann boundary condition

$$\frac{\partial \phi_i}{\partial \nu} = 0 \quad \text{on } \partial S. \tag{2.3}$$

Being defined by the metric rather than the explicit embedding makes the LBO and its spectral decomposition invariant to isometries and thus a popular operator for shape processing and analysis. For example, the eigenfunctions and eigenvalues can be used to efficiently approximate diffusion distances and commute time distances [31, 7, 13, 14, 12] that were defined as computational alternatives to geodesic distances and were shown to be robust to topology changes and global scale transformations. At another end, Lévy [24] proposed manipulating

the geometry of shapes by operating in their spectral domain, while Karni and Gotsman [21] chose the eigenfunctions as a natural basis for approximating the coordinates of a given shape. Feature point detectors and descriptors of surfaces were also extracted from the same spectral domain. Such measures include the heat kernel signature (HKS) [37, 17], the global point signature (GPS) [34], the wave kernel signature (WKS) [4], and the scale-space representation [41]. It was exploited for dimensionality reduction in shape matching and recognition; see, for example, [33].

Given two surfaces S and Q and a bijective mapping between them, $\rho : S \to Q$, Ovsjanikov et al. [27] emphasized the fact that the relation between the spectral decomposition of a scalar function $f : S \to \mathbb{R}$ and its representative on Q, that is, $f \circ \rho^{-1} : Q \to \mathbb{R}$, is linear. In other words, the geometry of the mapping is captured by ρ, allowing the coefficients of the decompositions to be related in a simple linear manner. The basis extracted from the LBO was chosen in this context because of its intuitive efficiency in representing functions on manifolds and thus far justified heuristically. The linear relation between the spectral decomposition coefficients of the same function on two surfaces, when the mapping between manifolds is provided, was exploited by Pokrass et al. [30] to find the correspondence between two almost isometric shapes. They assumed that the matrix that links between the LBO eigenfunctions of two almost isometric shapes should have dominant coefficients along its diagonal, a property that was first exploited in [23].

One could use the relation between the eigenstructures of two surfaces to approximate non-scalar and nonlocal properties of the manifolds [1]. Examples of such functions are geodesic distances [22, 39, 35, 26, 38], which serve as an input for the *multidimensional scaling* (MDS) [8], the *generalized multidimensional scaling* (GMDS) [11], and the Gromov–Hausdorff distance [25, 10]. Using the optimality of representing surfaces and functions on surfaces with a truncated basis, geodesic distances can now be efficiently computed and matched in the spectral domain.

In classical signal processing, a low pass filter version of a given signal in a Fourier sense is traditionally associated with the Karhunen–Loève transform, resorting to covariance functions and stochastic processes. The same nondeterministic viewpoint motivated the geometry processing community. For example, in [6] the optimality of the LBO eigenfunctions for mesh compression was explored from a statistical perspective, relating between the LBO and the covariance matrix of the distribution of mesh geometries. In a different domain, the eigenvectors of LBOs approximated from sampled data were often used in manifold learning. In [5] the LBO decomposition of sampled data manifolds was shown to well represent the geometric structure of the data, while in [42], a way to estimate the number of required eigenfunctions to represent a sampled manifold was presented. See also [36, 18, 15] for related efforts.

Among the myriad reasons that motivated the selection of the spectral domain for shape analysis the following are often emphasized:

- The spectral domain is isometric invariant.
- Countless signal processing tools that exploit the Fourier basis are available. Some can be generalized to shapes for representation, processing, analysis, and synthesis.
- Most interesting functions defined on surfaces are smooth and can thus be well approximated by their projection onto a small number of eigenfunctions.
- For most articulated objects in nature the problem of finding correspondences between

the shapes of the same object in different poses appears to have simple formulations in the spectral domain.

Still, a rigorous justification for the selection of the basis defined by the LBO was missing in the geometry processing arena. Along the same lines, combining the eigenstructure of the LBO with classical data representation and analysis procedures that operate in other domains like the PCA [20], MDS [8], and GMDS [11] was yet to come. Here, we review recent improvements of tools that make use of the decomposition of the LBO and introduce new procedures that utilize the proven optimality of this space. We provide a theoretical justification for using the LBO eigendecomposition in many shape analysis methods. With this property in mind, we demonstrate that it is possible to migrate algorithms to the spectral domain while establishing a substantial reduction in complexity.

3. Optimality of the LBO eigenspace. In this section we provide a theoretical justification for the choice of the LBO eigenfunctions by proving that the resulting spectral decomposition is optimal in approximating functions with L^2 bounded gradient magnitudes. Let S be a given Riemannian manifold with a metric (g_{ij}) and an induced LBO, Δ_g, with associated spectral basis ϕ_i, where $\Delta_g \phi_i = \lambda_i \phi_i$. It is shown, for example in [2], that for any $f : S \to \mathbb{R}$, the representation error

$$\|r_n\|_2^2 \equiv \left\| f - \sum_{i=1}^{n} \langle f, \phi_i \rangle \phi_i \right\|_2^2 \leq \frac{\|\nabla_g f\|_2^2}{\lambda_{n+1}}. \tag{3.1}$$

Our next result asserts that the eigenfunctions of the LBO are optimal with respect to estimate error (3.1).

Theorem 3.1. *Let $0 \leq \alpha < 1$. There is no integer n and no sequence $\{\psi_i\}_{i=1}^n$ of linearly independent functions in L^2 such that*

$$\left\| f - \sum_{i=1}^{n} \langle f, \psi_i \rangle \psi_i \right\|_2^2 \leq \frac{\alpha \|\nabla_g f\|_2^2}{\lambda_{n+1}} \qquad \forall f. \tag{3.2}$$

Proof. Recall the Courant–Fischer min-max principle; see [9, Problems 37 and 49] and [40]. We have for every $n \geq 0$,

$$\lambda_{n+1} = \max_{\substack{\Lambda \\ \operatorname{codim} \Lambda = n}} \min_{\substack{f \in \Lambda \\ f \neq 0}} \left\{ \frac{\|\nabla_g f\|_2^2}{\|f\|_2^2} \right\}. \tag{3.3}$$

That is, the min is taken over a linear subspace $\Lambda \subset H^1(S)$ (where $H^1(S)$ is the Sobolev space $\{f \in L^2, \nabla_g f \in L^2\}$) of co-dimension n and the max is taken over all such subspaces.

Set $\Lambda_0 = \{f \in H^1(S);\ \langle f, \psi_i \rangle = 0,\ i = 1, 2, \ldots, n\}$, so that Λ_0 is a subspace of co-dimension n. By (3.2) we have that for all $f \neq 0,\ f \in \Lambda_0$,

$$\frac{\|\nabla_g f\|_2^2}{\|f\|_2^2} \geq \frac{\lambda_{n+1}}{\alpha},$$

and thus

$$X_0 = \min_{\substack{f\in\Lambda_0 \\ f\neq 0}} \left\{ \frac{\|\nabla_g f\|_2^2}{\|f\|_2^2} \right\} \geq \frac{\lambda_{n+1}}{\alpha}. \tag{3.4}$$

On the other hand, by (3.3),

$$\lambda_{n+1} \geq X_0. \tag{3.5}$$

Combining (3.4) and (3.5) yields $\alpha \geq 1$. ∎

For the convenience of the reader we present in Appendix A a direct proof of a special case of the above result which does not make use of the Courant–Fischer min-max principle. The above theorem proves the optimality of the eigenfunctions of the LBO in representing H^1 functions on manifolds. In the following sections we apply the optimality property for solving various shape analysis problems.

4. Scale invariant geometry. Almost isometric transformations are probably the most common ones for surfaces and volumes in nature. Still, in some scenarios, relations between surfaces should be described by slightly more involved deformation models. Though a small child and an adult are obviously not isometric, and the same goes for a whale and a dolphin, the main characteristics are morphometrically similar for mammals at large. In order to extend the scope of matching and comparing shapes, a semilocal scale invariant geometry was introduced in [3]. There, it was used to define a new LBO by which one can construct an eigenspace which is invariant to semilocal and obviously global scale transformations.

Let (g_{ij}) be the regular metric defined on the manifold. In [3], the *scale invariant pseudometric* $(\tilde{g}_{ij})$ is defined as

$$\tilde{g}_{ij} = |K| g_{ij},$$

where K is the Gaussian curvature at each point on the manifold. One could show that this metric is scale invariant and that the same goes for the LBO that it induces, namely $\Delta_{\tilde{g}} f = -\frac{1}{\sqrt{\tilde{g}}} \sum_{ij} \partial_i \left(\sqrt{\tilde{g}} \tilde{g}^{ij} \partial_j f \right)$. A discretization of this operator and experimental results that outperformed state-of-the-art algorithms for shape matching, when scaling is involved, were presented in [3] . Specifically, the scale invariant geometry allows finding correspondence between two shapes related by semilocal scale transformation.

Next, one could think of searching for an optimal representation space for shapes by interpolating between the scale invariant metric and the regular one. We define the interpolated pseudometric to be

$$\hat{g}_{ij} = |K|^{\alpha} g_{ij},$$

where $(\hat{g}_{ij})$ represents the new pseudometric, K is the Gaussian curvature, and $\alpha \in [0, 1]$ is the metric interpolation scalar that we use to control the representation error. In our setting, $\hat{g}$ depends on α and represents the regular metric when $\alpha = 0$ or represents the scale invariant one when $\alpha = 1$.

Figure 1 depicts the effect of representing a triangulated shape's coordinates projected to the first 300 eigenfunctions of the LBO with a regular metric (left), the scale invariant one (right), and the interpolated pseudometric with $\alpha = 0.4$ (middle). The idea is to use

only a part of eigenfunctions to approximate smooth functions on the manifold, treating the coordinates as such. While the regular natural basis captures the global structure of the surface, as expected, the scale invariant one concentrates on the fine features with effective curvature. The interpolated one is a good compromise between the global structure and the fine details. Note again that each shape at the bottom row reflects the optimal reconstruction with respect to a different metric.

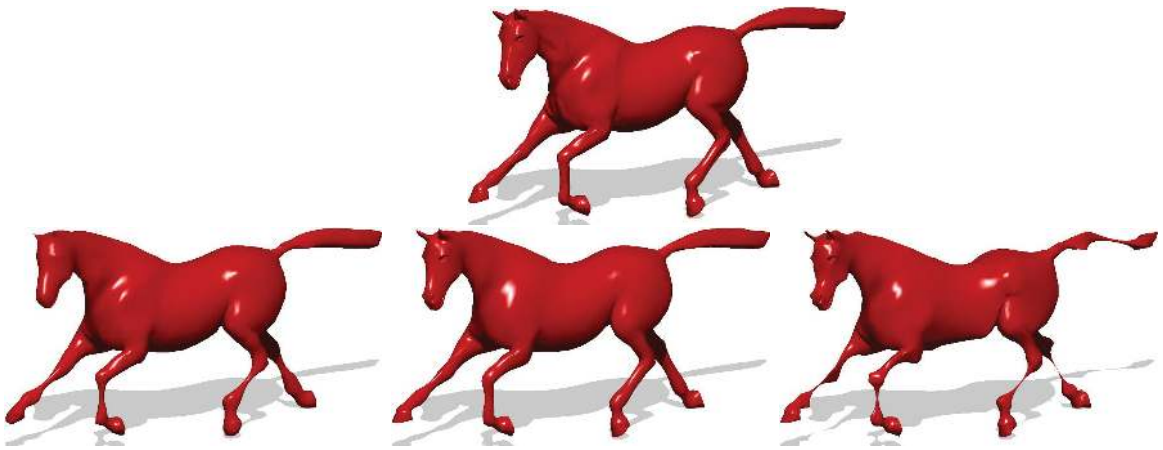

Figure 1. *Top: A given horse model. Bottom: The horse coordinates projected to the first* 300 *LBO eigenfunctions using a regular metric (left), an intermediate one (middle), and a scale invariant one (right).*

We emphasize again that once a proper metric is defined, the Laplace–Beltrami eigenspace is the best possible space for approximation of smooth functions. Next, we exploit this property to reformulate classical shape analysis algorithms such as MDS in the spectral domain.

5. Spectral classical scaling. MDS [8] is a family of data analysis methods that is widely used in machine learning and shape analysis. Given an $m \times m$ pairwise distance matrix $\mathbf{D}$, the MDS method finds an embedding of points of $\mathbb{R}^k$, given by an $m \times k$ matrix $\mathbf{X}$, such that the pairwise euclidean distances between every two points, each defined by a row of $\mathbf{X}$, is as close as possible to their corresponding input pair given by the right entry in $\mathbf{D}$. The classical MDS algorithm minimizes the following functional:

$$\mathbf{X} = \underset{\mathbf{X}\in\mathbb{R}^{m\times k}}{\operatorname{argmin}} \left\| \mathbf{X}\mathbf{X}^T + \frac{1}{2}\mathbf{J}\mathbf{D}_2\mathbf{J} \right\|_F, \tag{5.1}$$

where $\mathbf{D}_2$ is a matrix such that $(\mathbf{D}_2)_{ij} = \mathbf{D}_{ij}^2$, and $\mathbf{J}$ is a centering matrix defined by $\mathbf{J} = \mathbf{I} - \frac{1}{m}\mathbb{1}\mathbb{1}^T$, where $\mathbf{I}$ is the identity matrix, $\mathbb{1}$ is a vector of ones, and $\|\cdot\|_F$ is the Frobenius norm. The solution can be obtained by a singular value decomposition (SVD) of the matrix $\mathbf{J}\mathbf{D}_2\mathbf{J}$. This method was found to be useful when comparing isometric shapes using their inter-geodesic distances [16, 10] and texture mapping in computer graphics [43]. The computation of geodesic distances as well as the SVD of an $m \times m$ matrix can be expensive in terms of memory and computational time. High resolution shapes with more than 10^4 vertices are difficult to handle with this method.

In order to reduce the complexity of the problem, it was proposed in [2] to compute geodesic distances between a small set of sample points and then interpolate the rest of the distances by minimizing a biharmonic energy in the spectral domain. We find a spectral representation of the matrix $\mathbf{D}_2 = \mathbf{\Phi}\boldsymbol{\alpha}\mathbf{\Phi}^T$, where $\mathbf{\Phi}$ represents the matrix that discretizes the spectral domain. We then embed our problem into the eigenspace of the LBO, defining $\mathbf{X} = \mathbf{\Phi}\boldsymbol{\beta}$, where $\boldsymbol{\beta}$ is an $n \times k$ matrix, and $n \ll m$ in order to reduce the overall complexity. Note that here $\mathbf{\Phi}$ is truncated to an $m \times n$ matrix, so that only the first n eigenvectors are used. $\mathbf{X}$ is obtained by minimizing

$$\min_{\boldsymbol{\beta}} \left\| \mathbf{\Phi}\boldsymbol{\beta}\boldsymbol{\beta}^T\mathbf{\Phi}^T + \frac{1}{2}\mathbf{J}\mathbf{\Phi}\boldsymbol{\alpha}\mathbf{\Phi}^T\mathbf{J} \right\|_F . \tag{5.2}$$

Experimental results of shape canonization comparing shapes flattened with spectral classical scaling to regular classical scaling results were presented in [2]. The spectral approach outperformed the classical scaling in terms of time and space complexities, which led to overall better accuracy for the spectral version; see Figure 2. In the next section we introduce a novel design of functional spaces that benefit from both the LBO and the classical PCA, while extending the scope of each of these measures.

Figure 2. *MDS flat embedding: Result of regular MDS (middle) and the spectral MDS (right) of the given horse surface (left). Here, the target embedding space is $k = 3$, the number of surface points is $m = 8000$, and the number of eigenfunctions used for the analysis is $n = 150$.*

6. Regularized PCA. When dealing with an articulated object, we would like to have a model that is robust to possible realizations or poses of the object in the three-dimensional world. To that end, we revisit the definition of the classical PCA and regularize its structure exploiting the intrinsic geometry of the given manifold. The resulting basis defined on the manifold allows us to synthesize and analyze surface intrinsic as well as extrinsic properties, such as a smooth texture painted on the surface and the surface coordinates, as well as features and signatures that could be useful for alignment and recognition. Here again, we use the fact that the eigenfunctions provided by decomposition of the LBO depict the optimal basis for smooth functions on a given manifold. Combining this property with a basis optimized for the minimal projection error on a set of given functions defined on the manifold, we build a continuous intrinsically regularized version of the classical PCA. We provide quantitative and qualitative empirical evidence supporting the efficiency of the proposed novel construction. Note that a related regularized PCA promoting a different geometric model recently appeared in [19]; see also [32] for a nongeometric smooth PCA construction.

In the field of machine learning, PCA [28] is a basis designed to capture a different kind of structure. It provides an orthonormal basis that minimizes, in a Frobenius L_2 sense, the projection error of a given set of features. The PCA basis is extracted one element at a time in a greedy manner such that each new direction is orthonormal to the previous one and sequentially minimizes the error of representation by projection. Each basis element is also referred to as a moment of the given dataset.

When projecting functions defined on the shape into the LBO eigenfunctions while limiting the number of basis functions, the resulting representation error can be significant. Yet, resorting to the eigenfunctions provided by a PCA of the underlying manifold coordinates, considered at various poses, would make little sense if our goal were to represent functions defined on the manifold itself. The elements provided by a PCA basis capture the nature of continuous functions for given sets of corresponding points, or corresponding surfaces, in the embedding space defined by a specific coordinate system. At the other end, the LBO basis optimally represents continuous smooth functions defined on the intrinsic geometry of the shape manifold, that is, the embedded space.

In this section we marry the LBO eigenfunctions with a regular PCA. The result is a simple eigenvalue problem that naturally balances between surface representation in the embedding space and smooth functions defined on the surface itself. We demonstrate our construction with several numerical examples justifying the proposed PCA regularization by the LBO model that is defined through a Dirichlet energy.

The regularized PCA provides an efficient representation for a family of continuous functions defined for a given set of isometrically similar manifolds. It is demonstrated to efficiently represent out-of-sample functions using just a few projections onto the proposed basis. An example of such functions are the surface coordinates as used in our numerical experiments.

Given the set of functions $\{f_i\}_{i=1}^k$, where $f_i : S \to \mathbb{R}$, the principal components of the set are defined by an orthonormal basis $\{\psi_i\}$ optimized with respect to the truncation error

$$\sum_{i=1}^{k} \left\| f_i - \sum_{j=1}^{n} \langle f_i, \psi_j \rangle \psi_j \right\|_2^2 \tag{6.1}$$

for any given n or in its geometric flavor,

$$\sum_{i=1}^{k} \left\| f_i - \sum_{j=1}^{n} \langle f_i, \psi_j \rangle \psi_j \right\|_g^2. \tag{6.2}$$

Note that now the functions we are trying to model are assumed to belong to a given set, and no assumption about their smoothness with respect to the underlying manifold is assumed.

Next, assume that we have information about the manifold given by Δ_g and that the functions we deal with have bounded gradient $\|\nabla_g f\|_2^2$ on the manifold. The functional that would produce the eigenfunctions of the LBO as the minimization of the truncation error defined in (3.1) is the Dirichlet energy

$$\sum_{j=1}^{n} \|\nabla_g \psi_j\|_g^2. \tag{6.3}$$

In fact, feeding in the eigenfunctions of the LBO, namely $\{\phi_i\}$, into the above measure we obtain

$$\begin{aligned}\sum_{j=1}^{n} \|\nabla_g \phi_j\|_g^2 &= \sum_{j=1}^{n} \langle \Delta_g \phi_j, \phi_j \rangle_g \\ &= \sum_{j=1}^{n} \langle \lambda_j \phi_j, \phi_j \rangle_g \\ &= \sum_{j=1}^{n} \lambda_j. \end{aligned} \tag{6.4}$$

Here, the only knowledge about the functions we would like to represent by $\{\psi_i\}_{i=1}^n$ is the fact that their gradient magnitude on S is bounded.

In a discrete setting, PCA [20] is the process of extracting a low rank orthonormal approximate representation from a given set of vectors $\{x_i\}_{i=1}^k$, sometimes referred to as data points, that in our setting approximate the continuous functions $\{f_i\}_{i=1}^k$. Given a set of k vectors $x_i \in \mathbb{R}^m$, the PCA algorithm finds an orthonormal basis of $n \leq k$ vectors defined by $\{P_j\}_{i=j}^n$ by minimizing

$$\begin{aligned}&\min_{\mathbf{P}} \sum_{i=1}^{k} \|\mathbf{P}\mathbf{P}^T x_i - x_i\|_2^2 \\ &\text{such that (s.t.)} \\ &\mathbf{P}^T\mathbf{P} = \mathbf{I}_n. \end{aligned} \tag{6.5}$$

Here, the quantity $\mathbf{P}\mathbf{P}^T x_i$ is the projection of the vector x_i onto the orthonormal basis $\mathbf{P}$. Thus, the term $\|\mathbf{P}\mathbf{P}^T x_i - x_i\|_2^2$ represents the error of projecting x_i onto $\mathbf{P}$. It is straightforward to show that

$$\begin{aligned}\|\mathbf{P}\mathbf{P}^T x_i - x_i\|_2^2 &= \left(\mathbf{P}\mathbf{P}^T x_i - x_i\right)^T \left(\mathbf{P}\mathbf{P}^T x_i - x_i\right) \\ &= x_i^T \mathbf{P} \underbrace{\mathbf{P}^T\mathbf{P}}_{=\mathbf{I}_n} \mathbf{P}^T x_i - 2x_i^T \mathbf{P}\mathbf{P}^T x_i + x_i^T x_i \\ &= x_i^T x_i - x_i^T \mathbf{P}\mathbf{P}^T x_i \\ &= x_i^T x_i - \operatorname{trace}(x_i^T \mathbf{P}\mathbf{P}^T x_i) \\ &= x_i^T x_i - \operatorname{trace}(\mathbf{P}\mathbf{P}^T x_i x_i^T). \end{aligned} \tag{6.6}$$

Thus, the minimization problem is equivalent to

$$\begin{aligned}&\max_{\mathbf{P}} \sum_{i=1}^{k} \operatorname{trace}(\mathbf{P}\mathbf{P}^T x_i x_i^T) \\ &\text{s.t.} \\ &\mathbf{P}^T\mathbf{P} = \mathbf{I}_n. \end{aligned} \tag{6.7}$$

Denoting

$$\sum_{i=1}^{k} x_i x_i^T = \mathbf{X}\mathbf{X}^T, \tag{6.8}$$

where $\mathbf{X}$ is the matrix whose columns account for the x_i's, our problem can be written as

$$\begin{aligned} &\max_{\mathbf{P}} \operatorname{trace}(\mathbf{P}\mathbf{P}^T\mathbf{X}\mathbf{X}^T) \\ &\text{s.t.} \\ &\mathbf{P}^T\mathbf{P} = \mathbf{I}_n. \end{aligned} \tag{6.9}$$

The solution of this problem is obtained by an SVD of the matrix $\mathbf{X}$.

In a geometric setting, the PCA problem reads as

$$\begin{aligned} &\min_{\mathbf{P}} \sum_{i=1}^{k} \|\mathbf{P}\mathbf{P}^T\mathbf{A}x_i - x_i\|_g^2 \\ &\text{s.t.} \\ &\mathbf{P}^T\mathbf{A}\mathbf{P} = \mathbf{I}_n, \end{aligned} \tag{6.10}$$

where $\mathbf{A}$ is a discretization matrix of $\sqrt{g}$, the square root of the determinant of the metric (g_{ij}), which quantifies the volume element as defined by the given parametrization. For a triangulated surface, $\mathbf{A}$ is given by a diagonal matrix whose A_{ii} element is the sum of areas of all triangles that share the surface vertex i.

As for the discrete Dirichlet energy, one can show that for one vector P_j, which will later define the jth column of $\mathbf{P}$, we have

$$\|\nabla_g P_j\|_g^2 = (\mathbf{L}P_j)^T\mathbf{A}P_j, \tag{6.11}$$

where $\mathbf{L} = \mathbf{A}^{-1}\mathbf{W}$ is a discretization matrix of the LBO; see Appendix B.

The term $(\mathbf{L}P_j)^T\mathbf{A}P_j$ can be written as

$$(\mathbf{L}P_j)^T\mathbf{A}P_j = P_j^T\mathbf{W}^T\mathbf{A}^{-T}\mathbf{A}P_j = P_j^T\mathbf{W}P_j = \operatorname{trace}\left(\mathbf{W}P_jP_j^T\right). \tag{6.12}$$

Using this notation, the formulation of the discrete geometric PCA can be approximated by

$$\begin{aligned} &\max_{\mathbf{P}} \operatorname{trace}\left(\mathbf{P}\mathbf{P}^T\mathbf{A}\mathbf{X}\mathbf{X}^T\mathbf{A}\right) \\ &\text{s.t.} \\ &\mathbf{P}^T\mathbf{A}\mathbf{P} = \mathbf{I}_n, \end{aligned} \tag{6.13}$$

while the optimally smooth basis is obtained by

$$\begin{aligned} &\min_{\mathbf{P}} \operatorname{trace}\left(\mathbf{P}\mathbf{P}^T\mathbf{W}\right) \\ &\text{s.t.} \\ &\mathbf{P}^T\mathbf{A}\mathbf{P} = \mathbf{I}_n. \end{aligned} \tag{6.14}$$

Next, we would like to combine both measures, namely the energy defined by the data projection error (PCA) and the Dirichlet energy (LBO) of the representation space. A straightforward addition of the energies would lead to a numerical challenge; see, for example, [19].

Here, we propose a different strategy. We use the fact that the solution of (6.14) is equivalent to that of

$$\begin{aligned} &\max_{\mathbf{P}} \operatorname{trace}\left(\mathbf{P}\mathbf{P}^T\tilde{\mathbf{W}}^{-1}\right) \\ &\text{s.t.} \\ &\mathbf{P}^T\mathbf{A}\mathbf{P} = \mathbf{I}_n, \end{aligned} \tag{6.15}$$

where $\tilde{\mathbf{W}}^{-1}$ represents the pseudoinverse matrix of $\mathbf{W}$.

Exploiting this property, the regularized-PCA problem can be defined as

$$\begin{aligned} &\max_{\mathbf{P}} \operatorname{trace}\left(\mathbf{P}\mathbf{P}^T\left(\mathbf{A}\mathbf{X}\mathbf{X}^T\mathbf{A} + \mu\tilde{\mathbf{W}}^{-1}\right)\right) \\ &\text{s.t.} \\ &\mathbf{P}^T\mathbf{A}\mathbf{P} = \mathbf{I}_n. \end{aligned} \tag{6.16}$$

In this setting, the matrix $\mathbf{M} = \mathbf{A}\mathbf{X}\mathbf{X}^T\mathbf{A} + \mu\tilde{\mathbf{W}}^{-1}$ is symmetric and positive definite.

Classical algorithms performing a numerical eigendecomposition of a given matrix $\mathbf{M}$ usually produce the eigenvectors one by one sorted by the magnitudes of their corresponding eigenvalues. Several numerical optimization libraries such as ARPACK, LAPACK, and EIGS implement the Arnoldi iteration algorithm to extract the eigenvector associated with the eigenvalue of largest magnitude. The core operation in these methods is a multiplication of the matrix we would like to decompose by a vector. Therefore, these methods are particularly efficient when the input is a sparse matrix. Here, we show how to adapt such a method to our problem. We virtually deal with two matrices, the low rank $\mathbf{A}\mathbf{X}\mathbf{X}^T\mathbf{A}$ and the sparse one $\mathbf{W}$. Decomposition of their combination as given in (6.16) using off the shelf solvers would be inefficient. In order to resolve these issues, we recall that several implementations of Arnoldi iterations that find the eigenvectors of a given matrix $\mathbf{M}$ do not require the explicit matrix but rather a function whose input is a vector y and whose output is the product $\mathbf{M}y$.

Performing an economy-sized SVD of the matrix $\mathbf{X}$ (whose size is $m \times k$) such that $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^T$, where $\mathbf{U}$ is an $m \times m$ orthonormal matrix, $\mathbf{D}$ is an $m \times k$ diagonal matrix, and $\mathbf{V}$ is an $k \times k$ orthonormal matrix, we have

$$\mathbf{A}\mathbf{X}\mathbf{X}^T\mathbf{A} = \mathbf{A}\mathbf{U}\mathbf{D}^2\mathbf{U}^T\mathbf{A} = \tilde{\mathbf{U}}\mathbf{D}^2\tilde{\mathbf{U}}^T, \tag{6.17}$$

where $\tilde{\mathbf{U}} = \mathbf{A}\mathbf{U}$.

Next, given the vector y, the matrix vector multiplication $(\tilde{\mathbf{U}}\mathbf{D}^2\tilde{\mathbf{U}}^T + \tilde{\mathbf{W}}^{-1})y$ can equivalently be written as $\tilde{\mathbf{U}}(\mathbf{D}^2(\tilde{\mathbf{U}}^T y)) + \tilde{\mathbf{W}}^{-1}y$, for which the numerical multiplication procedure reads as follows.

Algorithm 1. Compute $\tilde{\mathbf{U}}(\mathbf{D}^2(\tilde{\mathbf{U}}^T y)) + \tilde{\mathbf{W}}^{-1}y$.

input y
$v \leftarrow \tilde{\mathbf{U}}^T y$
$v \leftarrow \mathbf{D}^2 v$
$v \leftarrow \tilde{\mathbf{U}} v$
Compute u by solving $y = \mathbf{W}u$ in a least squares sense.
return $(v + u)$

The above computation is extremely efficient. First, we calculate $\tilde{\mathbf{U}}^T y$, where here $\tilde{\mathbf{U}}$ is truncated to an $m \times k$ matrix (k is equal to the number of PCA input functions, $k \ll m$) that yields an $k \times 1$ vector. This operation has a complexity of $O(km)$. Then, we multiply the result by $\mathbf{D}^2$, which is effectively a $k \times k$ matrix, an operation that takes $O(k^2)$. Finally, the result is multiplied by $\tilde{\mathbf{U}}$ in $O(km)$. The overall computational complexity is $O(km)$, and the same holds for space complexity. The second phase consists of solving a sparse system (with $O(m)$ effective elements), which can be efficiently executed with classical algorithms such as the EIGS function in MATLAB. The proposed method efficiently converges and can provide one hundred eigenvectors in a couple of seconds for a triangulated surface with about 10^4 vertices.

6.1. Experimental results. In our first example, we consider a set of smooth functions defined on a smooth manifold. This set is composed of several *heat kernel signatures* (HKSs) evaluated at various times [37]. In this setting, time refers to an intrinsic scale at which the signature is computed, while the signature itself could be thought of as a smoothed version of the Gaussian curvature. Noise was then added to these functions, which were used to construct the PCA basis and our regularized-PCA basis. Next, we selected other smooth functions that belong to the same family, namely signatures taken at new time scales, and computed the representation error using the regular PCA and the proposed regularized PCA. The representation error as a function of the signal to noise ratio (SNR) is shown in Figure 3. Notice the significant advantage of using the regularized PCA compared to its nonregularized version.

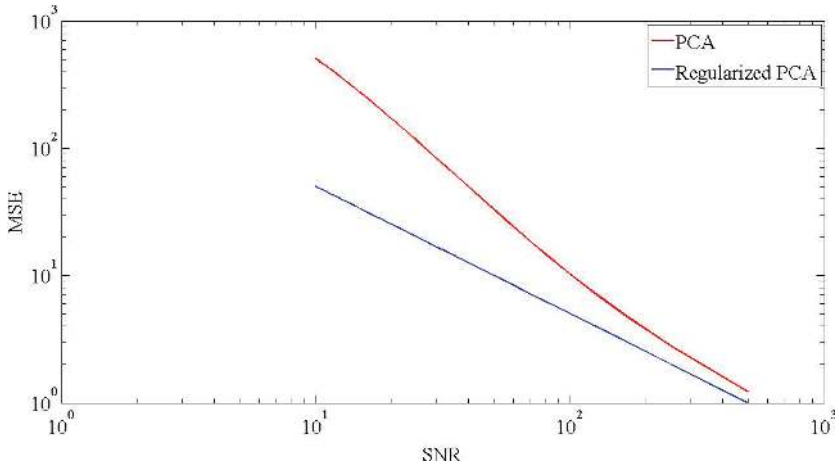

Figure 3. *Mean square error of functions represented by PCA and its regularized version.*

Next, we consider five different poses of a cat which are approximately isometries. We computed an orthonormal basis that minimizes the reconstruction error of the coordinates extracted from two shapes of the cat presented at the top. These six coordinate vectors provide us with six functions defined on the surface. We then computed the regularized version that accounts for efficiently representing continuous smooth functions on the manifold. Figure 4 exemplifies the superiority of the regularized PCA in this extreme case. In another example

the benefits of this hybrid model in representing out of sample information are demonstrated in Figure 5.

In the next example, shown in Figure 6, the 300 coordinate functions of 100 postures of a hand, two of which are presented in the box in the first row, were utilized to construct a 100-sized PCA basis. This basis was then used to reconstruct the open hand posed at the top right, the result of which is shown on the bottom left. The first 100 eigenfunctions of the LBO reconstruction of the same hand, shown in the middle of the second row, demonstrate the power of the LBO in reconstructing the smooth structure of the shape rather than its details. The reconstruction with the 100 eigenfunctions obtained for the regularized PCA provides an almost perfect reconstruction, as depicted on the bottom right.

In this section we demonstrated that while smooth functions are indeed important as far as data about the manifold goes, in some instances, additional information about the manifold could and should be utilized. At the other end of the spectrum, PCA procedures allow us to greedily construct an efficient basis for a set of given functions representing such additional information, such as the coordinates of shapes in different poses that are treated as functions on the surface. These, seemingly opposite, basis generation strategies were unified into a coherent basis construction framework that enjoys the benefits of both. The LBO was shown to be a natural regularizer for the PCA, while the PCA of a given set of known functions was shown to steer the basis provided by the LBO into one that better represents functions far and beyond the scope of the available training functions captured by a regular PCA. The proposed model allowed us to design an alternative basis which is efficient to construct, related to the spectral domain but whose properties can be tuned to fit specific information about the data.

7. Conclusion.

A theoretical support for the selection of the leading eigenfunctions of the LBO of a given shape as a natural domain for morphometric study of surfaces was provided. The optimality result motivates the design of efficient shape matching and analysis algorithms. It enabled us to find the most efficient representations of smooth functions on surfaces in terms of both accuracy and complexity when projected onto the leading eigenfunctions of the LBO. Our optimality criterion is obviously defined with respect to a given metric. In shape representation, the choice of an appropriate metric is probably as important as the selection of the most efficient subspace. This was demonstrated in approximating the fine details and the general structure of a shape of a horse in section 4 using a regular metric, a scale invariant metric, and a metric interpolating between the two. Spectral classical scaling and its generalized version benefit from the presented optimality result, as does the novel regularized PCA presented in the previous section. In both cases it was demonstrated that the decomposition of the LBO provides a natural subspace within which to operate.

The provably optimal representation space allows us to construct efficient tools for computational morphometry—the numerical study of shapes. Revisiting the optimality criteria obviously leads to alternative domains and hopefully better analysis tools that we plan to explore in the future.

Appendix A. Alternative proof.

Theorem A.1. *Given a Riemannian manifold S with a metric (g_{ij}), the induced LBO, Δ_g, and its spectral basis ϕ_i, where $\Delta_g \phi_i = \lambda_i \phi_i$, and a real scalar value $0 \leq \alpha < 1$, there is no*

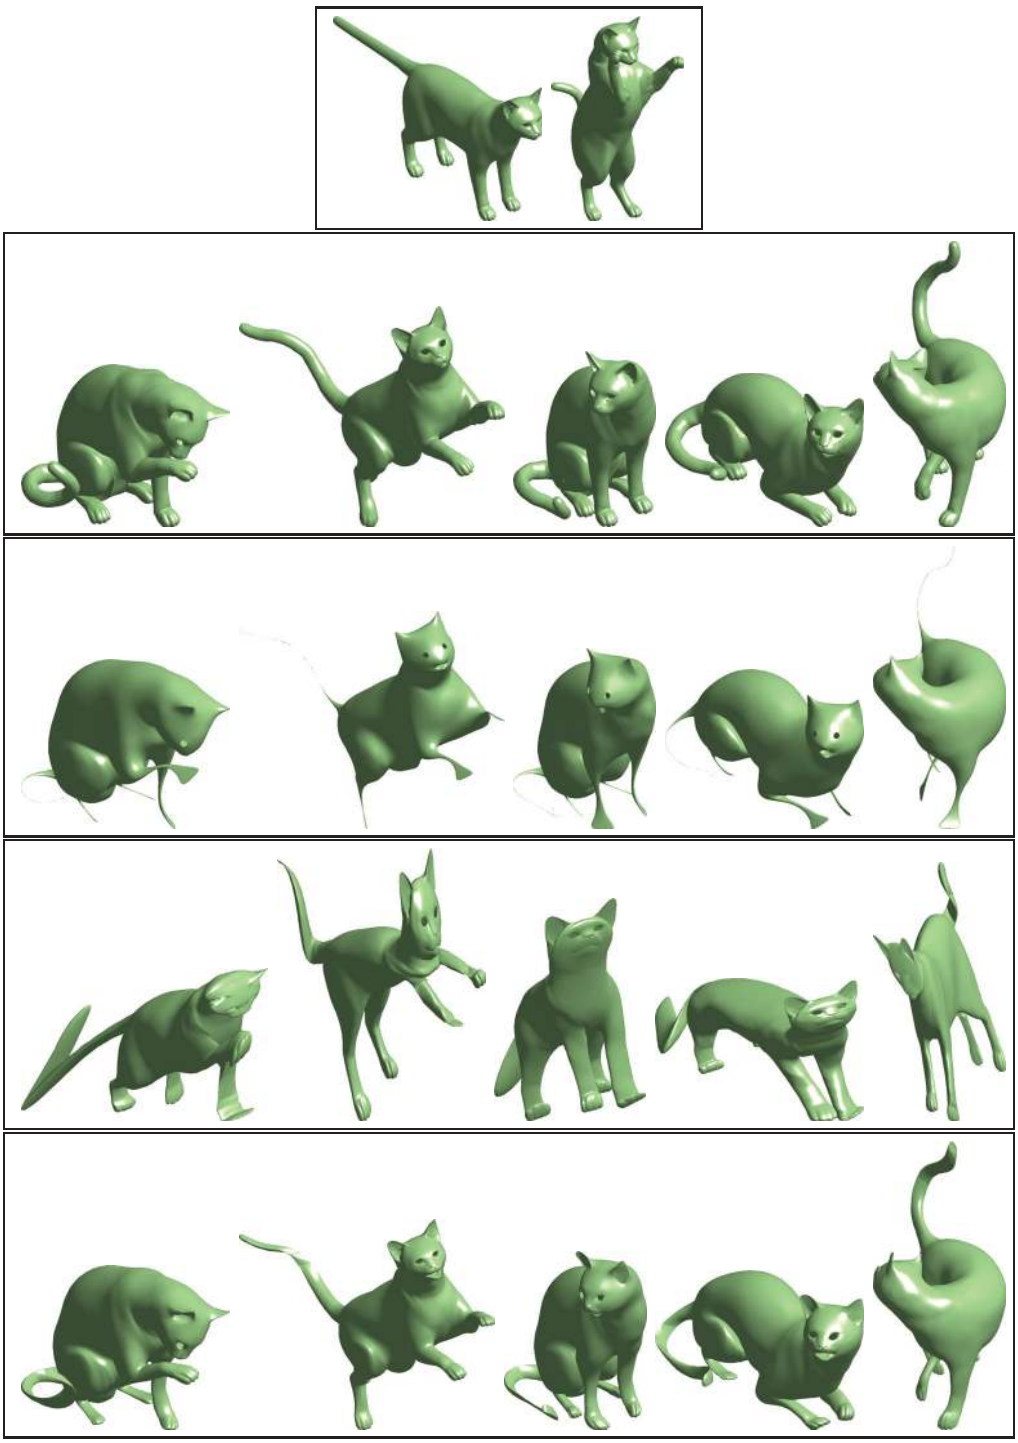

Figure 4. *Reconstruction of the original cats (second row) by projecting the coordinates to the first leading* 100 *vectors of the Laplace–Beltrami eigenbasis (third row) , to the PCA trained with two shapes (fourth row), and (bottom) to the regularized-PCA basis trained with the two cats (at the top).*

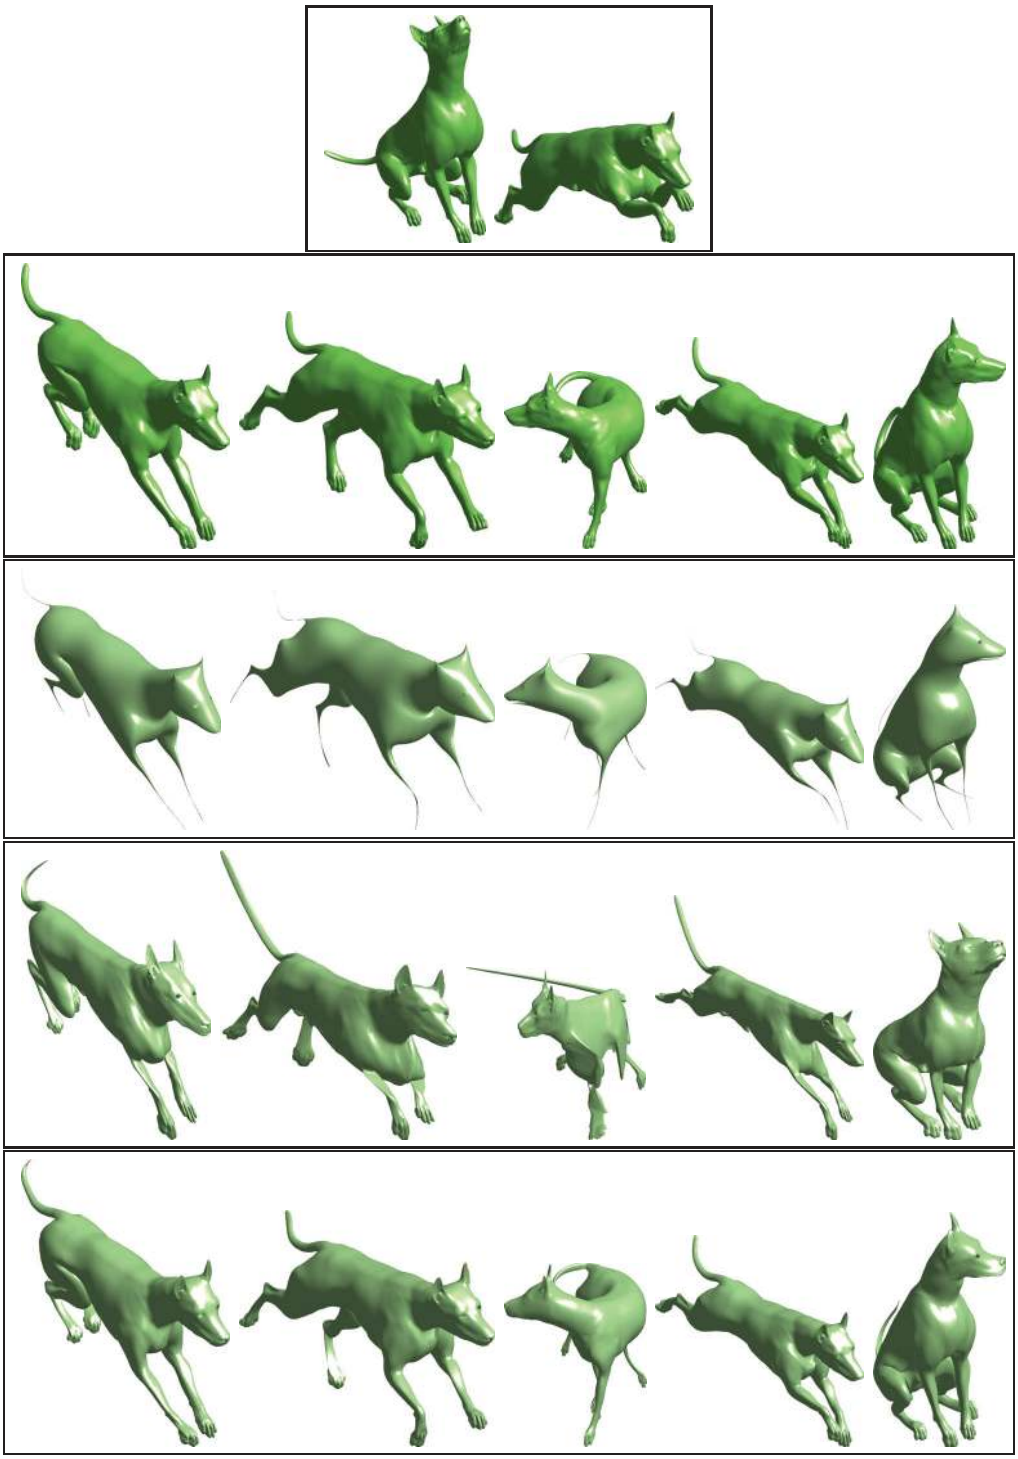

Figure 5. *Reconstruction of the dog shapes in the second row by projecting their coordinates to the first* 100 *eigenfunctions of the Laplace–Beltrami eigenbasis (third row), to the PCA basis trained with the two shapes from the first row (fourth row), and to the* 100 *basis functions of the regularized-PCA basis trained with the two dogs from the first row (bottom/fifth row).*

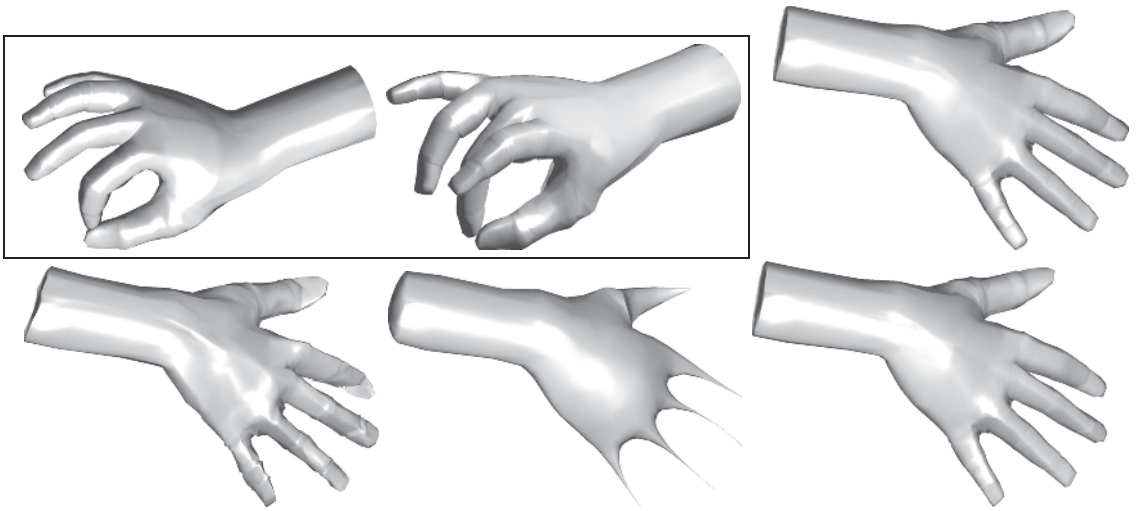

Figure 6. *Top left box: Two out of the* 100 *hand poses whose coordinates were used to construct the PCA. Top right: The hand we wish to reconstruct. Bottom left: PCA reconstruction. Bottom middle: LBO reconstruction. Bottom right: Regularized-PCA reconstruction result.*

orthonormal basis of functions $\{\psi_i\}_{i=1}^{\infty}$ *and an integer* n *such that*

$$\left\| f - \sum_{i=1}^{n} \langle f, \psi_i \rangle \psi_i \right\|_2^2 \leq \alpha \frac{\|\nabla_g f\|_2^2}{\lambda_{n+1}} \qquad \forall f. \tag{A.1}$$

To prove the optimality of the LBO eigenbasis, let us first prove the following lemma.

Lemma A.2. *Given an orthonormal basis* $B = \{b_1, b_2, \ldots\}$ *of a Hilbert space* $\mathcal{V}$ *and an orthogonal projection operator* P *of* $\mathcal{V}$ *such that*

$$\|Pv\| \leq k\|v\| \quad \forall v \in \operatorname{span}\{b_i, \quad 1 \leq i \leq n\}, \tag{A.2}$$

where $0 < k < 1$*, then*

$$\dim(\ker(P)) \geq n. \tag{A.3}$$

Proof. Let us denote

$$\begin{aligned} \mathcal{B}_1 &= \{b_i, \|Pb_i\| < 1\}, \\ \mathcal{B}_2 &= \{b_i, \|Pb_i\| = 1\}, \\ \mathcal{P}_1 &= \ker(P), \\ \mathcal{P}_2 &= \operatorname{im}(P). \end{aligned} \tag{A.4}$$

Because the operator P is orthogonal and the basis B is orthonormal, we have

$$\mathcal{V} = \mathcal{B}_1 \bigoplus \mathcal{B}_2 = \mathcal{P}_1 \bigoplus \mathcal{P}_2 \tag{A.5}$$

and

$$\begin{aligned} \mathcal{B}_1^{\perp} &= \mathcal{B}_2, \\ \mathcal{P}_1^{\perp} &= \mathcal{P}_2. \end{aligned} \tag{A.6}$$

By definition, we have that

$$\mathcal{B}_2 \subset \mathcal{P}_2. \tag{A.7}$$

Then,

$$\mathcal{B}_2^{\perp} \supset \mathcal{P}_2^{\perp}, \tag{A.8}$$

and since $\mathcal{B}_2^{\perp} = \mathcal{B}_1$ and $\mathcal{P}_2^{\perp} = \mathcal{P}_1$, we have

$$\mathcal{P}_1 \subset \mathcal{B}_1. \tag{A.9}$$

Now, assume that

$$\dim(\mathcal{P}_1) = \dim(\ker(P)) < n \leq \dim \mathcal{B}_1. \tag{A.10}$$

Then, $\mathcal{P}_1 \neq \mathcal{B}_1$, and we can find a vector $u \in \mathcal{P}_1^{\perp}$ such that $\|u\| = 1$ and $u \in \mathcal{B}_1$. Since $\mathcal{P}_1^{\perp} = \ker(P)^{\perp} = \mathcal{P}_2$, it follows that

$$\|Pu\| = 1. \tag{A.11}$$

But, this contradicts the fact that $u \in \mathcal{B}_1$ because $u \in \mathcal{B}_1$ implies

$$\|Pu\| < 1. \tag{A.12}$$

Then,

$$\dim(\ker(P)) \geq n. \tag{A.13}$$ ■

Equipped with this result we are now ready to prove Theorem A.1.

Proof. Assume that there exists such a basis, $\{\psi_i\}$. Then, the representation of a function f in the eigenbasis of the LBO can be written as

$$f = \sum_{i=1}^{\infty} \langle f, \phi_i \rangle \phi_i = \sum_{i=1}^{\infty} \beta_i \phi_i. \tag{A.14}$$

We straightforwardly have

$$\|\nabla_g f\|_2^2 = \sum_{i=1}^{\infty} \lambda_i \beta_i^2, \tag{A.15}$$

and it follows that

$$\alpha \frac{\|\nabla_g f\|_2^2}{\lambda_{n+1}} = \sum_{i=1}^{\infty} \underbrace{\frac{\alpha \lambda_i}{\lambda_{n+1}}}_{\tilde{\lambda}_i} \beta_i^2 = \sum_{i=1}^{\infty} \tilde{\lambda}_i \beta_i^2. \tag{A.16}$$

Moreover,

$$\left\| f - \sum_{i=1}^{n} \langle f, \psi_i \rangle \psi_i \right\|_2^2 \leq \alpha \frac{\|\nabla_g f\|_2^2}{\lambda_{n+1}} \leq \sum_{i=1}^{\infty} \tilde{\lambda}_i \beta_i^2. \tag{A.17}$$

Then, replacing f with $\sum_{j=1}^{n+1} \beta_j \phi_j$, we have

$$\left\| \sum_{j=1}^{n+1} \beta_j \phi_j - \sum_{i=1}^{n} \left\langle \sum_{j=1}^{n+1} \beta_j \phi_j, \psi_i \right\rangle \psi_i \right\|_2^2 \leq \sum_{j=1}^{n+1} \beta_j^2 \tilde{\lambda}_j \leq \left(\max_{j=1}^{n+1} \tilde{\lambda}_j \right) \left(\sum_{j=1}^{n+1} \beta_j^2 \right), \tag{A.18}$$

and since $\tilde{\lambda}_i < 1$ for all i, $1 \leq i \leq n+1$, we can state that there is a set of $n+1$ orthonormal vectors ϕ_i belonging to an orthonormal basis whose projection error (over the space spanned by ψ_i) is smaller than one. According to the previous lemma, the original assumption leads to a contradiction because the dimension of the kernel of the projection on the space spanned by ψ_i, $1 \leq i \leq n$, is n. ■

Appendix B. Cotangent weight approximation for Δ_g. The cotangent weight discretization [29] is widely used in the field of shape analysis as an approximation for an LBO operator $\mathbf{L}$ on a triangulated surface. The discrete LBO is defined as

$$\mathbf{L} = \mathbf{A}^{-1}\mathbf{W}, \tag{B.1}$$

where $\mathbf{W}$ is the famous cotangent weight matrix

$$\mathbf{W}_{\xi\eta} = \begin{cases} \sum\limits_{(\xi,\tau)\in E} (\cot \gamma_{\xi\tau} + \cot \beta_{\xi\tau}) & \text{if } \xi = \eta, \\ -(\cot \gamma_{\xi\eta} + \cot \beta_{\xi\eta}) & \text{if } \xi \neq \eta,\ (\xi,\eta) \in E, \end{cases} \tag{B.2}$$

where $\gamma_{\xi\eta}$ and $\beta_{\xi\eta}$ are defined as shown in Figure 7, and E represents the set of edges in the triangulated surface.

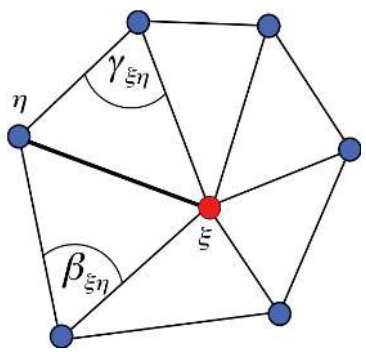

Figure 7. *Cotangent weight construction of the LBO about the vertex ξ.*

Acknowledgments. YA and RK thank Anastasia Dubrovina, Alon Shtern, Aaron Wetzler, and Michael Zibulevsky for intriguing discussions.

REFERENCES

[1] Y. Aflalo, A. Dubrovina, and R. Kimmel, *Spectral Generalized Multi-Dimensional Scaling*, preprint, arXiv:1311.2187, 2013.

[2] Y. Aflalo and R. Kimmel, *Spectral multidimensional scaling*, Proc. Natl. Acad. Sci. USA, 110 (2013), pp. 18052–18057.

[3] Y. Aflalo, R. Kimmel, and D. Raviv, *Scale invariant geometry for nonrigid shapes*, SIAM J. Imaging Sci., 6 (2013), pp. 1579–1597.

[4] M. Aubry, U. Schlickewei, and D. Cremers, *The wave kernel signature: A quantum mechanical approach to shape analysis*, in Proceedings of the IEEE International Conference on Computer Vision Workshops, 2011, pp. 1626–1633.

[5] M. Belkin and P. Niyogi, *Convergence of Laplacian eigenmaps*, in Advances in Neural Information Processing Systems, Vol. 19, MIT Press, Cambridge, MA, 2007, pp. 129–136.

[6] M. Ben-Chen and C. Gotsman, *On the optimality of spectral compression of mesh data*, ACM Trans. Graphics, 24 (2005), pp. 60–80.

[7] P. Bérard, G. Besson, and S. Gallot, *Embedding Riemannian manifolds by their heat kernel*, Geom. Funct. Anal., 4 (1994), pp. 373–398.

[8] I. Borg and P. Groenen, *Modern Multidimensional Scaling: Theory and Applications*, Springer, New York, 1997.

[9] H. Brezis, *Functional Analysis, Sobolev Spaces and Partial Differential Equations*, Universitext, Springer, New York, 2010.

[10] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, *Efficient computation of isometry-invariant distances between surfaces*, SIAM J. Sci. Comput., 28 (2006), pp. 1812–1836.

[11] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, *Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching*, Proc. Natl. Acad. Sci. USA, 103 (2006), pp. 1168–1172.

[12] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro, *A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching*, Int. J. Comput. Vis., 89 (2010), pp. 266–286.

[13] R. R. Coifman and S. Lafon, *Diffusion maps*, Appl. Comput. Harmon. Anal., 21 (2006), pp. 5–30.

[14] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, *Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps*, Proc. Natl. Acad. Sci. USA, 102 (2005), p. 7426–7431.

[15] R. R. Coifman and M. Maggioni, *Diffusion wavelets*, Appl. Comput. Harmon. Anal., 21 (2006), pp. 53–94.

[16] A. Elad and R. Kimmel, *On bending invariant signatures for surfaces*, IEEE Trans. Pattern Anal. Mach. Intell., 25 (2003), pp. 1285–1295.

[17] K. Gębal, J. Andreas Bærentzen, H. Aanæs, and R. Larsen, *Shape analysis using the auto diffusion function*, in Computer Graphics Forum, Vol. 28, Wiley Online Library, 2009, pp. 1405–1413.

[18] D. K. Hammond, P. Vandergheynst, and R. Gribonval, *Wavelets on graphs via spectral graph theory*, Appl. Comput. Harmon. Anal., 30 (2011), pp. 129–150.

[19] B. Jiang, C. Ding, B. Luo, and J. Tang, *Graph-Laplacian PCA: Closed-form solution and robustness*, in Proceedings of the 26th Annual IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 3492–3498.

[20] I. T. Jolliffe, *Principal Component Analysis*, 2nd ed., Springer, New York, 2002.

[21] Z. Karni and C. Gotsman, *Spectral compression of mesh geometry*, in Proceedings of the 27th Annual ACM Conference on Computer Graphics and Interactive Techniques, 2000, pp. 279–286.

[22] R. Kimmel and J. A. Sethian, *Fast marching methods on triangulated domains*, Proc. Natl. Acad. Sci. USA, 95 (1998), pp. 8341–8435.

[23] A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, K. Glashoff, and R. Kimmel, *Coupled quasi-harmonic bases*, Comput. Graph. Forum, 32 (2013), pp. 439–448.

[24] B. Lévy, *Laplace-Beltrami eigenfunctions towards an algorithm that "understands" geometry*, in Proceedings of the IEEE International Conference on Shape Modeling and Applications, 2006, p. 13.

[25] F. Memoli and G. Sapiro, *A theoretical and computational framework for isometry invariant recognition*

of point cloud data, Found. Comput. Math., 5 (2005), pp. 313–347.
[26] J. S. B. Mitchell, D. M. Mount, and C. H. Papadimitriou, *The discrete geodesic problem*, SIAM J. Comput., 16 (1987), pp. 647–668.
[27] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas, *Functional maps: A flexible representation of maps between shapes*, ACM Trans. Graphics, 31 (2012), 30.
[28] K. Pearson, *On lines and planes of closest fit to systems of points in space*, Phil. Mag., 2 (1901), pp. 559–572.
[29] U. Pinkall and K. Polthier, *Computing discrete minimal surfaces and their conjugates*, Experim. Math., 2 (1993), pp. 15–36.
[30] J. Pokrass, A. M. Bronstein, M. M. Bronstein, P. Sprechmann, and G. Sapiro, *Sparse modeling of intrinsic correspondences*, Comput. Graph. Forum, 32 (2013), pp. 459–468.
[31] H. Qiu and E. R. Hancock, *Clustering and embedding using commute times*, IEEE Trans. Pattern Anal. Mach. Intell., 29 (2007), pp. 1873–1890.
[32] J. O. Ramsay and B. W. Silverman, *Functional Data Analysis*, 2nd ed., Springer, New York, 2005.
[33] G. Rong, Y. Cao, and X. Guo, *Spectral mesh deformation*, Vis. Comput., 24 (2008), pp. 787–796.
[34] R. Rustamov, M. Ovsjanikov, O. Azencot, M. Ben-Chen, F. Chazal, and L. Guibas, *Map-based exploration of intrinsic shape differences and variability*, ACM Trans. Graphics, 32 (2013), 72.
[35] J. A. Sethian, *A fast marching level set method for monotonically advancing fronts*, Proc. Nat. Acad. Sci. U.S.A., 93 (1996), pp. 1591–1595.
[36] J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*, Cambridge University Press, New York, 2004.
[37] J. Sun, M. Ovsjanikov, and L. Guibas, *A concise and provably informative multi-scale signature based on heat diffusion*, in Proceedings of the Symposium on Geometry Processing (SGP '09), Aire-la-Ville, Switzerland, 2009, pp. 1383–1392.
[38] V. Surazhsky, T. Surazhsky, D. Kirsanov, S. J. Gortler, and H. Hoppe, *Fast exact and approximate geodesics on meshes*, ACM Trans. Graphics, 24 (2005), pp. 553–560.
[39] J. N. Tsitsiklis, *Efficient algorithms for globally optimal trajectories*, IEEE Trans. Automat. Control, 40 (1995), pp. 1528–1538.
[40] H. F. Weinberger, *Variational Methods for Eigenvalue Approximation*, SIAM, Philadelphia, 1974.
[41] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud, *Surface feature detection and description with applications to mesh matching*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 373–380.
[42] X. Zhou and N. Srebro, *Error analysis of Laplacian eigenmaps for semi-supervised learning*, in Proceedings of the International Conference on Artificial Intelligence and Statistics, 2011, pp. 901–908.
[43] G. Zigelman, R. Kimmel, and N. Kiryati, *Texture mapping using surface flattening via multidimensional scaling*, IEEE Trans. Visual. Comput. Graph., 8 (2002), pp. 198–207.