

A Fast Modal Space Transform for Robust Nonrigid Shape Retrieval

Jianbo Ye*

Yizhou Yu†

Abstract

Nonrigid or deformable 3D objects are common in many application domains. Retrieval of such objects in large databases based on shape similarity is still a challenging problem. In this paper, we take advantages of functional operators as characterizations of shape deformation, and further propose a framework to design novel shape signatures for encoding nonrigid geometries. Our approach constructs a context-aware integral kernel operator on a manifold, then applies modal analysis to map this operator into a low-frequency functional representation, called *fast functional transform*, and finally computes its spectrum as the shape signature. In a nutshell, our method is fast, isometry-invariant, discriminative, smooth and numerically stable with respect to multiple types of perturbations. Experimental results demonstrate that our new shape signature for nonrigid objects can outperform all methods participating in the nonrigid track of the SHREC'11 contest. It is also the second best performing method in the real human model track of SHREC'14.

1 Introduction

Content-based 3D object retrieval facilitates the search for desired objects within a large 3D object repository. It has become increasingly popular due to the rapid development of 3D scanning technologies and the emergence of large 3D object databases. Content-based object retrieval is useful in many application domains, including CAD/CAM, medicine, molecular biology, 3D computer games and virtual worlds.

Since many 3D object models, such as avatars, creatures and biomedical objects, can take various types of deformations, it is much desired for an object retrieval technique to be able to recognize deformed versions of an object. Nonetheless, nonrigid object retrieval is a very challenging task because a deformed object may not be visually similar to the original one any more. In the following, we summarize a few criteria for measuring the performance of nonrigid object retrieval techniques.

- **Isometry Invariance.** There exist a large variety of nonrigid deformations. In theory, one can always work out a

deformation that transforms one object into another completely different object. Therefore, it is very important to define a subclass of deformations that usually do not alter the identity of an object, and then develop nonrigid object retrieval methods that are invariant or at least insensitive to this subclass of deformations. In the literature, isometric deformations are commonly adopted for this purpose.

- **Discrimination Power.** In most retrieval techniques, an object is represented with a relatively short shape descriptor or signature. It is important for the shape descriptor or signature to encode the intrinsic characteristics of an object while discarding unimportant details so that it can distinguish two different objects even when they are visually very similar, and accurately measure the degree of dissimilarity between them. Such a capability makes it possible for a retrieval technique to return a ranked list of objects with decreasing similarity with a query object.
- **Efficiency.** Computation of shape signatures or descriptors contains offline and online stages. Many methods often require significant amount of time to process a single mesh in offline stage, which prohibits their use in large database (say if a PC takes ten minutes to process a mesh, it would take tens of years to process an entire database in size of a million). In online stage, comparison of shape is also often time-critical, such that signature-based approaches are still favored over shape matching approaches in a typical retrieval system.
- **Smoothness and Stability.** Although isometric modeling provides a state-of-the-art theoretical base for nonrigid object retrieval, numerical stability of isometry-invariant shape descriptors or signatures has not been sufficiently explored. It is much desired that a shape descriptor only changes slightly when a small amount of nonisometric deformation is introduced to the original object. Likewise, it is also desired that the shape descriptor is stable when other types of defects, such as noise and holes, are introduced to the original object. Stable shape descriptors and signatures give rise to robust retrieval results. Mathematical definitions with respect to the smoothness of shape perturbations and the stability of descriptors are still missing.

In addition to the above four criteria, there exist other considerations, including *extensibility*, *applicability* and *complexity of implementation* (e.g. parameter free).

*J. Ye is with College of Information Science and Technology, The Pennsylvania State University.

†Y. Yu is with College of Computer Science, Zhejiang University.

1.1 Related Work

Among extensive work on 3D object retrieval, most techniques are devoted to rigid objects and are based on extrinsic geometry such as Euclidean distance, curvatures and snapshots of 3D canonical views. Nevertheless, nonrigid models have gained increasing popularity. Their extrinsic geometry often varies under nonrigid deformations. Isometric shape deformation was initially addressed in [13], where researchers began to consider bending invariant or insensitive 3D shape recognition. Recently, more extensive classes of invariance have been studied, such as elastic deformations [19] and affine transformations [33].

To our knowledge, two major classes of approaches are proposed for nonrigid shape retrieval during the last decade. The first class includes all local feature based approaches (such as earlier work [40, 17, 48, 29]). Inspired by the success of the SIFT feature descriptor in image retrieval, researchers proposed local descriptors, such as meshSIFT [29], Mesh-HOG [48], covariance descriptor [43], and descriptors based on the spectral graph wavelet transform [23], for representing features on mesh surfaces. ShapeGoogle [8] emphasizes the robustness of association and classification, especially for objects with missing parts and topological noises. It integrates the local Heat Kernel Signature (HKS) [40] with the bag-of-word framework [21]. By extracting isometrically invariant dense point descriptors and quantizing them into binary codes, shapes are registered for efficient indexing, comparison and association. A more recent method [15] also endeavors to characterize 3D surfaces by measuring the likelihood of points in the 3D space being local centers of radial symmetry at selected scales. It could be useful in domain-specific 3D object retrieval, where the scale is known.

The second class emphasizes coarser-scale or global shape of a 3D nonrigid model. The skeleton-based method (e.g. [2]) encodes the geometric and topological information in the form of a skeleton graph and uses graph matching to retrieve similar skeletons. An approach for matching 3D objects in the presence of nonrigid transformations and partially similar models is presented in [42], which uses 3D curves extracted around feature points to represent surfaces. Another method based on distributions of diffusion distances has been proposed in [9] for matching nonrigid shapes. Techniques, such as histograms, D2 distributions [36] and spatial pyramids [22], are also popular in designing descriptors in respect of coarse-scale shapes. Some of them are “context-aware”: the term “context-aware” refers to dependencies between geometric characteristics not within each other’s immediate neighborhoods, distinguishing itself from “bag-of-words” approaches that focus on local geometric neighborhoods. Laga et al. [20] model the context of a shape part as a set of walks in the graph of specific length originating from the part’s node. In particular, they construct a context-aware similarity measure by relating their model of contexts to semantic correspondences. In our paper, we merely

consider the context as a purely geometric property, aka the context of a point given a shape domain, and therefore, do not intend to recognize functionalities of shape parts. As other global shape descriptors, we do not generate an exact signature for contexts, yet our proposed signature is sensitive to any changes of such geometric contexts.

Among those methods concerning global shapes, spectrum-based techniques became popular in recent years. Shape-DNA [35] proposes to use the spectrum of the Laplace-Beltrami operator as an isometry-invariant shape descriptor. Another similar method, SD-GDM [38] proposes to compute a singular value decomposition (spectrum) for the geodesic distance matrix, which seems to outperform shape-DNA [24]. However, compared to shape-DNA, matrix assembly in SD-GDM requires all-pairs geodesic distances, which are computationally prohibitive to obtain even with the latest developments in fast geodesic distance computation [47], and a uniform mesh decimation is typically employed in practice to speed up this process. Utilizing spectral analysis is often effective for extracting isometric information. For example, spectral multidimensional scaling [1] has used the spectral projection of geodesic distances for better MDS embedding of data points. An efficient method for computing a robust spectrum-based shape descriptor insensitive to noises and small topological changes is presented in [46], which achieves efficiency and robustness by performing a modal space transform. This paper is an extended version of [46].

Local features, such as SIFT [28] and its variants, have been successfully adopted in image retrieval. However, in the context of 3D nonrigid shape retrieval, local feature based methods have often been outperformed by shape descriptors emphasizing coarser-scale shapes. The reason for this phenomenon is two-fold. First, pixel values captured by a camera are related to high-frequency surface properties, such as textures, of real scenes and objects. Given such high-frequency signals, it is possible for a local feature descriptor to encode sufficiently discriminative visual information for recognition or retrieval tasks. Second, 3D scanning techniques are not as mature as digital photography. Even when a 3D surface does have high-frequency details, such as pores and wrinkles on skin, it is unlikely for them to be accurately captured by a 3D scanner. Very often, such high-frequency details are buried in noises. Therefore, high-frequency geometry over a 3D surface tend to be inaccurate and unreliable (see [31] for related discussions).

Content-based multimedia retrieval is known as a problem/scenario dependent task. If one is interested in retrieving objects that share similar parts with the query object, methods based on point signatures should be more useful. However, if one is interested in finding globally similar nonrigid objects with potentially different poses and deformations, a global shape signature can more effectively prune the list of candidates. Even when a state-of-the-art point signature, such as the one in [23], is used in global shape retrieval, a spatial

aggregation and matching scheme, such as histogram matching and spatial pyramid matching, still needs to be employed. Inevitably, such aggregation partially loses relative spatial information among point features. On the other hand, a global signature is capable of encoding the complete geometric shape up to an invariant class, and thus is a more powerful global shape representation. Therefore, global and local shape signatures are complementary to each other.

1.2 Overview and Contributions

In this paper, we analyze functional operators over modal space and further introduce spectrum-based shape signatures to encode the shape of a nonrigid model. The basic idea underlying our shape signature is to compute the spectrum of the newly proposed functional operator, which is constructed from an intrinsic, context-aware integral kernel operator by projecting it to the linear space spanned by low-frequency modes defined over an object surface, while the integral kernel operator is itself based on a modal based pairwise distance. The resulting transformation matrix can be analytically written in a succinct form. Our method is isometry-invariant and stable with respect to noise, holes and nonisometric deformations. The implementation of our shape signature is based on linear FEM. The signature itself can be efficiently computed for meshes in a wide range of resolutions and topologies.

The main contribution of this paper is the introduction of a substantially improved global shape signature, which has been experimentally validated to meet most of the aforementioned criteria. Our experiments demonstrate that our new shape signature for nonrigid objects can outperform all methods participating in the nonrigid track of the SHREC 2011 contest [24], including the best performing method. In particular, our method achieves obviously higher precision than all other methods when recall is above 50%. Comparisons were performed on two representative datasets, the dataset used in the nonrigid track of SHREC'11 and a more challenging one created by ourselves. We have also tested our method on the more recent SHREC 2014 dataset for the retrieval of nonrigid human models [31], which have two tracks (synthetic and real). Our method is among the best performing methods in the real human model track.

The rest of the paper is organized as follows. In Section 2, we discuss the fundamental mathematics behind our signature design. In Section 3, we discuss our numerical implementation. In Section 4, we present experimental results to validate our shape signature. Section 5 concludes our paper.

2 Our approach

2.1 Laplace-Beltrami Operator

Shape-DNA [35] exploits eigenvalues to achieve an impressive shape retrieval performance, while there are a number of other

methods, such as [36, 40], utilizing eigenvectors. All these methods compute the spectrum of the Laplace-Beltrami operator Δ_M , where M is the underlying manifold embedded in the 3D Euclidean space as a surface, by solving the following eigenvalue problem,

$$-\Delta_M u = \lambda u. \quad (1)$$

The Laplace-Beltrami operator is a generalized operator for functions defined on Riemannian manifolds. By solving the above equation, we obtain a complete set of modal bases $\{\phi_i\}_{i=0}^{\infty}$ over a manifold, where each ϕ_i corresponds to a normalized eigenvector with eigenvalue λ_i in an ascending order.

2.2 Functional Operator

The eigenvectors of Laplace-Beltrami operator intrinsically span a low-frequency functional space Φ which is useful in modal analysis. For example, in [30], the authors construct intrinsic flexible maps between two shapes by solving for a linear transform matrix between their modal spaces $\Phi_1 \mapsto \Phi_2$ subject to certain constraints.

We instead focus on intrinsic functional transforms i.e. maps between Φ and itself. Of course, Laplace-Beltrami operator is itself a functional transform that map $\phi_i \mapsto -\lambda_i \phi_i$, which is isometry-invariant. We in this section introduce a new set of functional transforms whose spectra are more robust than the spectrum of the Laplace-Beltrami operator.

Given a symmetric function, $k(\cdot, \cdot)$, let us consider the following integral kernel operator,

$$\mathcal{K}f(y) = \int_M k(x, y)f(x)dx. \quad (2)$$

If $k(\cdot, \cdot)$ is isometry-invariant, the spectrum of \mathcal{K} is also isometry-invariant. For example, in the non-rigid track of the 2011 3D shape retrieval contest (SHREC'11) [24], the retrieval performance of SD-GDM [38], a method based on geodesic distance matrices (GDM), is ranked first. It outperforms shapeDNA. GDMs are in fact a class of integral kernel operators. If we set $k(x, y)$ as the geodesic distance between x and y , the spectrum of \mathcal{K} is the mathematically precise form of SD-GDM and is provable isometry-invariant. More importantly, we show in Appendix A that an integral kernel operator defines a family of "smooth" deformations, and the total deviation of the spectral signature of the integral kernel operator is finitely bounded under ϵ -deformation no matter how many dimensions the signature has. Such a property is important for spectral signatures, as their performance should improve rather than degrade with an increasing number of eigenvalues. A theoretical justification for the use of integral kernel operators is as follows. If the intended invariant class (typically larger than the isometric class) of shape perturbations is the family of deformations characterized by an integral kernel

operator \mathcal{K} whose perturbed kernel $\mathcal{K}_\epsilon^{-1}$ admits certain spatial smoothness condition, the distance measure based on its spectral signature with a varying number of dimensions stably changes with respect to the amount of perturbation. That is to say, in nonrigid shape retrieval, integral kernel operators can typically tolerate a larger class of smooth deformations than isometric deformations.

The major contribution of our approach is instead of naively computing a transform matrix in the complexity of the original geometry (i.e. pair-wise values $k(x_i, x_j)$ for all $x_i, x_j \in M$), we restrict the kernel to modal space Φ (i.e. applying the kernel to functions in Φ , the space spanned by the lower eigenvectors, and then projecting the solution back into Φ). We can write the transform matrix as

$$\tilde{K} = \Phi^T \mathcal{K} \Phi, \quad (3)$$

where $\Phi = [\phi_0, \phi_1, \dots, \phi_m, \dots]$.

It is worth noting that this restriction is not a simple efficient approximation, the functional transforms before and after restriction are different in the numerical aspects. Precisely speaking, the original functional space before restriction is a vector space endowed with an inner product; while the functional space Φ in restriction is an infinite dimensional separable Hilbert space (and is numerically approximated by truncated lower eigenvectors, which converge in a weak sense) spanned by the modes. These two vector spaces are endowed with the same inner product structure in the limit. However their numerical convergence properties are different. Table 1 summarizes their differences.

There exists a theoretical issue when we work with integral kernel operators numerically without modal restriction: numerical convergence is only guaranteed in the weak sense, while the underlying representation is point-wise. Thus the numerically estimated eigenvector of an operator may converge, but it is not guaranteed for the estimation to converge point-wise to the true eigenvector as the resolution of the mesh goes to the infinity. On the other hand, by restricting the functions of interest in the modal space, we make them have required smoothness that implicitly guarantee pointwise convergence. In our experiments (see comparison of BiHDM and R-BiHDM in Section 4.2), it has been observed that spectra computed by modal space restriction can better tolerate manifold deformations, and outperform the spectrum of a pairwise kernel in retrieval tasks.

2.3 Distance Map using Modes

Computing the geodesic distance matrix is very computationally expensive for large meshes. We choose to compute (squared) biharmonic distance [26] instead because it exhibits multiple nice properties while being more efficient to compute,

and can be restricted in modal space analytically in a succinct way.

In the continuous case, the (squared) biharmonic distance is defined as follows,

$$d^2(x, y) = \sum_{i=1}^{\infty} \frac{(\phi_i(x) - \phi_i(y))^2}{\lambda_i^2}, \quad (4)$$

where ϕ_i and λ_i are the eigenfunctions and eigenvalues (resp.) of the semi-positive definite Laplace-Beltrami operator, $-\Delta\phi_i(x) = \lambda_i\phi_i(x)$, where $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$ and $\int_M |\phi_i|^2 = 1$. The distance is a metric, and is smooth, locally isotropic, globally “shape aware”, isometry invariant, insensitive to noise and small topology changes, parameter-free, and practical to compute on a discrete mesh. In [26] these two types of distances have been extensively compared in detail. Biharmonic distance provides a nice trade-off between being nearly geodesic for small distances and global shape-awareness for large distances.

2.4 Functional Biharmonic Distance Map

Combining previous two section together, i.e. let $k(x, y) = d^2(x, y)$, we formulate \tilde{K} explicitly as follows.

$$\begin{aligned} \mathcal{K}\phi_0(y) &= \sum_{i=1}^{\infty} \int_M \frac{(\phi_i(x) - \phi_i(y))^2}{\lambda_i^2} \phi_0(x) dx \\ &= \frac{1}{\sqrt{A}} \sum_{i=1}^{\infty} \frac{1}{\lambda_i^2} + \sqrt{A} \sum_{i=1}^{\infty} \frac{\phi_i^2(y)}{\lambda_i^2} \end{aligned}$$

where A is the total area of M , and note $\phi_0 = 1/\sqrt{A}$. Let $\langle \cdot, \cdot \rangle$ be the standard inner product of L^2 functions, we have

$$\begin{aligned} a_0 &= \langle \phi_0, \mathcal{K}\phi_0 \rangle = \sum_{i=1}^{\infty} \frac{2}{\lambda_i^2}, \\ a_j &= \langle \phi_j, \mathcal{K}\phi_0 \rangle = \sqrt{A} \int_M \sum_{i=1}^{\infty} \frac{\phi_i^2}{\lambda_i^2} \phi_j \quad j > 0. \end{aligned} \quad (5)$$

We also have

$$\begin{aligned} \mathcal{K}\phi_j(y) &= \sum_{i=1}^{\infty} \int_M \frac{(\phi_i(x) - \phi_i(y))^2}{\lambda_i^2} \phi_j(x) dx \\ &= \int_M \sum_{i=1}^{\infty} \frac{\phi_i^2}{\lambda_i^2} \phi_j - \frac{2\phi_j(y)}{\lambda_j^2} \quad j > 0, \end{aligned}$$

where $\langle \phi_0, \mathcal{K}\phi_j \rangle = a_j$ and $\langle \phi_i, \mathcal{K}\phi_j \rangle = -\frac{2}{\lambda_j^2} \delta_{ij}$. Thus we

have obtained the projected matrix \tilde{K} , called reduced biharmonic distance matrix (R-BiHDM),

$$\tilde{K} = \begin{bmatrix} a_0 & a_1 & a_2 & \dots \\ a_1 & -2/\lambda_1^2 & & \\ a_2 & & -2/\lambda_2^2 & \\ \vdots & & & \ddots \end{bmatrix}. \quad (6)$$

¹See Appendix A for the definition of perturbed kernel.

	space	inner product	numerics	representation	numerical convergence
before	$L_2(M)$	$\langle f, g \rangle = \int_M fg$	finite element	pointwise	weak convergence in Sobolev space $W^{k,2}$
after	ℓ^2	$\langle f, g \rangle_{\ell^2} = \sum_{i=1}^{\infty} \langle f, \phi_i \rangle \langle \phi_i, g \rangle$	truncated modes	coefficient series	weak convergence in L^2

Table 1: Differences between integral kernel operators before and after modal space restriction

The above matrix is infinite. We first define $\tilde{K}_{m,n}$ be an $(n+1) \times (n+1)$ matrix

$$\tilde{K}_{m,n} = \begin{bmatrix} a_0^{(m)} & a_1^{(m)} & a_2^{(m)} & \dots & a_n^{(m)} \\ a_1^{(m)} & \frac{-2}{(\lambda_1^{(m)})^2} & 0 & & \\ a_2^{(m)} & & \frac{-2}{(\lambda_2^{(m)})^2} & 0 & \\ \vdots & & & \ddots & \vdots \\ a_n^{(m)} & 0 & 0 & \dots & \frac{-2}{(\lambda_n^{(m)})^2} \end{bmatrix}, \quad (7)$$

which is formed by the following two steps: i) take the first $n+1$ rows and first $n+1$ columns of \tilde{K} ; ii) set $\lambda_i^{(m)} := \lambda_i$ for $i \leq m$ and $\lambda_i^{(m)} := \infty$ if $i > m$ in Eq. (7), thus when calculating each $a_j^{(m)}$ in this truncated matrix $\tilde{K}_{m,n}$, every infinite summation in (5) is approximated by the first m terms. The spectrum of $\tilde{K}_{m,n}$ can converge in two ways:

$$\text{BiHDM: } \lim_{m \rightarrow \infty} \text{spectrum}(\lim_{n \rightarrow \infty} \tilde{K}_{m,n}), \quad (8)$$

and

$$\text{R-BiHDM: } \lim_{m \rightarrow \infty} \text{spectrum}(\tilde{K}_{m,m}). \quad (9)$$

The first scheme corresponds to the case where no modal space restriction is used. That is, $n \rightarrow \infty$ already before the spectrum is estimated. The second scheme is our method based on modal space restriction. The spectra under these two schemes are different asymptotically. For a finite m , BiHDM computes the spectrum of $\tilde{K}_{m,\infty}$, which still includes the high frequency components in the first row and the first column, i.e. $a_j^{(m)} \neq 0$ for $j > m$. On the other hand, R-BiHDM computes the spectrum of $\tilde{K}_{m,m}$ without any high frequency components. In general, the scheme with high frequency components removed is not guaranteed to outperform the one without. But we remark that our proposed signature has a close connection with the traditional Fourier descriptor [49], where coefficients corresponding to low-frequency components of a contour signal are used for constructing informative shape signatures. Since high frequency components are more likely to be contaminated with noises, which are typically high frequency signals themselves, high frequency components are less reliable in characterizing geometric shapes. Our approach extends a similar spirit to integral kernel operators on a manifold: the ‘‘low-frequency representation’’ of an integral kernel operator is more useful and reliable in characterizing its shape domains.

We call $\tilde{K}_m := \tilde{K}_{m,m}$ empirical R-BiHDM. As $m \rightarrow \infty$, the largest tens of eigenvalues of \tilde{K}_m enjoy a fast convergence

rate. Figure 7 shows the maximum error of the first 30 eigenvalues versus m , the number of eigenpairs of the Laplace-Beltrami operator. It has been observed that the asymptotic eigenvalues converge linearly.

Theorem 2.1 (Spectral Convergence of empirical R-BiHDM). *If M is a two-dimensional Riemannian manifold and is compact, then $\sum_{i=1}^{\infty} \frac{\phi_i^2}{\lambda_i^2}$ is bounded both pointwise and in the form of the L_2 norm. Furthermore, the leading eigenvalues of \tilde{K}_m converge to that of \tilde{K} as $m \rightarrow \infty$.*

Proof. See Appendix B for a proof with a more general condition. \square

Note that $\text{tr}(\tilde{K}) = 0$. We denote all eigenvalues of \tilde{K} in a magnitude descending order as $\{\mu_j\}_{j=0}^L$. We have observed that $\mu_0 > 0$ and $\mu_j < 0 \quad \forall j > 0$. (Such matrix has a single positive eigenvalue, and the rest are negative. See [3] and references therein) Hence a scale invariant spectrum can be defined as

$$\bar{\mu}_j = |\mu_j / \mu_0|. \quad (10)$$

Our shape signature is defined as a vector $S = [\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_L]^T$, which in theory is also isometric invariant. In practice, we select L ranging from 10 \sim 30, and $m > \max\{60, 2L\}$.

To compare two shape signatures, S^p and S^q , one can reference the dissimilarity measures in [38]. In particular, let us mention two useful ones here, mean normalized Manhattan distance,

$$D_1 = \sum_{j=1}^L \left| \frac{S_j^p - S_j^q}{S_j^p + S_j^q} \right|,$$

which is used in SD-GDM, and normalized Euclidean distance

$$D_2^2 = \sum_{j=1}^L \frac{(S_j^p - S_j^q)^2}{S_j^p S_j^q}, \quad (11)$$

which performs slightly better for our approach by experiments.

Computational Efficiency. A clear advantage of working with functionally transformed biharmonic distance is significantly reduced size and complexity of the kernel matrix. The computation of the original kernel matrix (either geodesics or biharmonic distance) has a much higher complexity, $O(Cn^2)$, than that of the reduced one (6), whose complexity is $O(nm)$.

Here n is the number of sample points on the surface, $m \ll n$, and C is the time for computing a distance between two sample points, e.g. $C = O(m)$ for biharmonic distance. Also, computing the eigenvalues of the original kernel matrix could be even more expensive because it is usually a large-scale $n \times n$ dense matrix, while computing the eigenvalues of the reduced one takes much less time.

3 Implementation

To compute the eigenspace of the Laplace-Beltrami operator on a manifold, we use finite elements [50] to formulate the eigenfunction space as in [35]. Although solving PDEs with FEM is sampling invariant, mesh quality during discretization is an important factor affecting numerical accuracy. Fortunately, there is already considerable amount of work [12, 5] in mesh generation, repairing and quality improvement. Our method only assumes that the mesh is properly refined to at least a few thousand triangles and has no (near) degenerated faces. In our implementation, we use Neumann boundary condition whenever a mesh has boundaries.

Following the standard setup of FEM, we need to solve the following sparse symmetric generalized eigenvalue problem that has efficient solvers, such as IRAM [4] and Krylov-Schur [39],

$$L\mathbf{u} = \lambda D\mathbf{u},$$

where L is the stiff matrix and D is the mass matrix. They are defined as follows. $L_{ij} = \int_M \nabla \varphi_i \cdot \nabla \varphi_j$, and $D_{ij} = \int_M \varphi_i \varphi_j$, where L_{ij} and D_{ij} are analytically calculated from the geometry of the elements, and φ_i is chosen to be a piecewise polynomial in each element. See appendix C for a detailed implementation of linear and cubic triangular FEMs.

Once L and D have been assembled, the inner products between two functions f and g , discretely represented by finite elements over a mesh domain, are given as

$$\int_M fg = \sum_{i,j=1}^N D_{ij} f_i g_j, \quad \int_M \nabla f \cdot \nabla g = \sum_{i,j=1}^N L_{ij} f_i g_j,$$

$$\int_M f^2 g = \sum_{i,j=1}^N f_i^2 D_{ij} g_j \text{ or } \sum_{i,j=1}^N f_i D_{ij} f_j g_j.$$

Once eigenfunctions of the Laplace-Beltrami operator have been computed, we are able to use the above equations to compute a_i in (6).

4 Experimental Results

4.1 Efficiency

Efficiency is usually important for shape retrieval techniques to be practical on large datasets. Computational intensity is a

major limitation of methods based on optimization, geodesic computation, and per-node based quantization. In contrast, our approach can always compute shape signatures in an efficient manner. Because our approach is directly based on computing lower eigenvalues/eigenvectors of Laplace-Beltrami operator and requires very little extra cost of assembly R-BiHDM (see eq. (6)) and solving for its eigenvalues. The eigensolver of Laplace-Beltrami operator contain a sparse direct preconditioner in complexity $O(n^2)$ and a dimensional free iterative eigensolver in complexity $O(n)$.

Model	#vert.	Time	Model	#vert.	Time
R-BiHDM: linear FEM			SD-GDM		
gmm_prisms	969	0.5s	gmm_prisms	969	14.8s
abstract	4096	1.7s	ant_dec	2502	2m28s
nonrigid ant	9501	4s	abstract	4096	12m13s
human meta	13336	5.8s	meshSIFT		
helicopter	22664	9.8s	ant_dec	2502	1m31s
bimba_cvd	74764	41s	abstract	4096	2m5s
desktop	106961	1m20s	nonrigid ant	9501	13m5s

Table 2: Timings for constructing shape signatures or descriptors. Our shape signature can be computed much more efficiently than most successful methods in the literature. For a mesh with 10k vertices, we can compute its signature within seconds.

On an Intel Core2 Duo CPU E8400@3.00GHz, running times required for the methods used in the subsequent section 4.2 are reported in Table 2. Such running times are based on our implementation of R-BiHDM and SD-GDM [38]², and the original authors' implementation of meshSIFT [29]³.

It is seen from the timing table that, computations of SD-GDM and meshSIFT are expensive even for a mesh with only thousands vertices and grow super linearly, while our method is much faster at the same resolutions (see timing of model "abstract"). Some of more recent leading performance methods, such as shapeGoogle [8], Supervised DL [27] and ISPM [22, 23], also suffer from the prohibitive long running time [31, 27].

4.2 Signature Based Retrieval

We have tested the nonrigid shape retrieval performance of our method on two representative datasets. One is provided by the nonrigid track of SHREC'11 [24] (Figure 1), which has 600 watertight triangle meshes that were derived from 30 original models. The second dataset mixes the nonrigid dataset from SHREC'11 with a dataset custom built by ourselves. The SHREC'11 nonrigid track dataset serves as a background dataset. The custom built dataset contains 200

²Geodesic distance is computed using the fast marching Matlab toolbox, <http://www.mathworks.com/matlabcentral/fileexchange/6110>

³This code can be downloaded at <https://mirc.uzleuven.be/MedicalImageComputing/downloads/meshSIFT.php>




Figure 1: 30 classes of nonrigid models in the nonrigid track of SHREC'11

deformed meshes derived from 4 biped models, 3 dinosaur models and 3 quadruped models (Figure 2(a)) and 40 background meshes (other biped, dinosaur and quadruped models, see Figure 2 (b)). The second dataset was designed to be more challenging and practical than the first one because it contains multiple similar meshes from each category of models, such as bipeds, dinosaurs and quadrupeds. Distinguishing similar models from the same category requires a retrieval technique to be more discriminative and stable. In the second dataset, we have performed retrieval tests on the 200 deformed meshes in a repository with a total of 840 (200+40+600) meshes.⁴


We applied the same evaluation methodology of the SHREC'11 contest to evaluate our method. It is based on the Precision-Recall curve and five quantitative measures: Nearest Neighbor (NN), First Tier (FT), Second Tier (ST), E-measure (E), and Discounted Cumulative Gain (DCG). We refer to [37] for detailed definitions. In our method, we use Normalized Euclidean distance to measure similarity among R-BiHDM signatures (see eq. (10)).

We have compared the retrieval performance of our method with that of Shape-DNA (OrigM-n12-norm1) [35], meshSIFT [29], and SD-GDM [38]. These are the best state-of-art performing methods in the nonrigid track of the SHREC'11 contest.⁵ The parameters in these methods were set empirically to produce best performance. We also compare our method, i.e. R-BiHDM, with pairwise biharmonic distance matrix (BiHDM) and report its performance. Detailed comparison results on the SHREC'11 nonrigid track dataset are shown in Table 3 and Figure 3(a). Detailed comparison results on the second dataset are shown in Table 4 and Figure 3(b).


According to the statistics, our method achieves better performance than all of these methods on both datasets. In particular, our method has obvious improvements in the 1-Tier precision. And according to the PR curves, our method achieves higher precision when recall is above 50%. Furthermore, the retrieval performance is stable when the number of chosen eigenvalues ranges from 15 to 30. The changes of NN, ST, E and DCG measures are not larger than 0.5% (absolutely)

⁴This dataset can be downloaded at <https://code.google.com/p/tri-mesh-toolkit/>

⁵A hybrid track combining SD-GDM and meshSIFT in SHREC'11 did achieve a better performance, but it falls out of scope in our state-of-art evaluation.



(a) Representative query meshes, including 4 bipeds, 3 dinosaurs and 3 quadrupeds, from the second dataset



(b) Ambiguous meshes, including 13 bipeds, 13 dinosaurs and 14 quadrupeds, from the second dataset

Figure 2: Examples from our second dataset

Method	NN	FT	ST	E	DCG
shape-DNA	0.998	0.890	0.952	0.696	0.975
MDS-CM-BoF	0.995	0.913	0.969	0.717	0.982
MeshSIFT	0.998	0.800	0.851	0.627	0.944
SD-GDM	1.000	0.956	0.982	0.728	0.992
BiHDM-L25	0.990	0.915	0.979	0.719	0.984
R-BiHDM-L30	0.998	0.969	0.981	0.729	0.993
R-BiHDM-L25	0.998	0.971	0.982	0.729	0.994
R-BiHDM-L23	1.000	0.971	0.983	0.730	0.994
R-BiHDM-L20	1.000	0.964	0.982	0.728	0.993
R-BiHDM-L15	0.998	0.955	0.981	0.726	0.991

Table 3: Retrieval performance evaluated using five standard measures on the SHREC'11 dataset. The performance of MDS-CM-BoF is cited from the SHREC'11 report.

on both datasets, while the changes of 1-Tier precisions on the original SHREC'11 dataset and the second dataset are smaller than 2% and 1% respectively. The best performance is achieved when $L = 23$ on both datasets simultaneously, further indicating that different datasets could share the same parameter setting. The primary goal of Shape Google [8] is to achieve a high level of robustness in the presence of partial shapes and topological noises. However, it does not address the identical problem as ours. In our experiments (based on implementation provided by authors), its retrieval performance on our two datasets is not comparable to the performance of our method.

Note that the performance gap between our method and the other methods becomes larger on the second dataset. As discussed earlier, the second dataset is more challenging with similar shape instances from the same categories. The results indicate that our method exhibits more discriminative power on datasets with very similar shape instances.

METHOD	NN	FT	ST	E	DCG
shape-DNA	0.985	0.841	0.906	0.666	0.954
MeshSIFT	0.995	0.790	0.890	0.650	0.950
SD-GDM	1.000	0.929	0.986	0.731	0.991
BiHDM-L25	1.000	0.930	0.983	0.722	0.990
R-BiHDM-L30	1.000	0.970	0.996	0.739	0.997
R-BiHDM-L25	1.000	0.975	0.997	0.742	0.998
R-BiHDM-L23	1.000	0.976	0.997	0.742	0.998
R-BiHDM-L20	1.000	0.975	0.997	0.742	0.998
R-BiHDM-L15	1.000	0.970	0.994	0.739	0.997

Table 4: Retrieval performance evaluated using five standard measures on the second dataset.

4.3 Nonrigid Retrieval of Human Models


We have also tested our proposed signature on the SHREC'14 nonrigid dataset for retrieving human models. There are two tracks, one uses synthetic data, and the other uses real data built from point clouds. These two datasets exhibit substantially different properties. The synthetic dataset consists of 15 different human models, each with its own unique body shape, which is determined by a parametric model. The real dataset is composed of 400 meshes, made up of 40 human subjects in 10 different poses. This real dataset is noisy and inaccurate in that certain body features cannot be accurately captured and human poses cannot be precisely controlled. As a participating method, detailed performance statistics of our method have been reported as the R-BiHDM-s method in [31] and more recently in [27], from which we only quote relevant results in this paper.

Method	Synthetic	Real
ISPM [22, 23]	90.2	25.8
HAPT [15]	81.7	63.7
Unsupervised DL [27]	84.2	53.3
Supervised DL [27]	95.4	79.1
ShapeGoogle(VQ) [8]	81.3	51.4
R-BiHDM-s	64.2	64.0


Table 5: Comparison with different retrieval methods in terms of mean average precision (mAP, in%) on the SHREC'14 Human Model datasets (selected results from [27]).

On the synthetic dataset, our method does not perform as well as others which gather statistics from local features. This is perhaps because our proposed signature is a geometric descriptor that captures smooth and global deformations. It has no intention to distinguish between different types of localized deformation, such as pose-driven muscle deformation. Overall, this dataset is relatively easy because a simplistic baseline method, such as the total area of a mesh, can already achieve decent performance [31].

The real dataset is much more challenging than the synthetic



(a) SHREC'11 nonrigid dataset



(b) the second dataset

Figure 3: Precision-recall curves of R-BiHDM, BiHDM, SD-GDM and shapeDNA on two retrieval tasks

one. The performance of our proposed signature is ranked second. The only participating method that outperforms ours is Supervised DL, which is a supervised data-driven method that requires an extra labeled dataset to train the model. Because of this, supervised methods have clear advantages over unsupervised ones, such as ours, in terms of retrieval performance. Nevertheless, our method achieves the best performance among unsupervised methods, including Unsupervised DL (the unsupervised version of Supervised DL), and even performs better than HAPT, another participating method that requires parameter optimization w.r.t. working datasets [31].

4.4 Continuity and Stability


Continuity analysis examines the range or level of shape deformation that a descriptor can capture. Ideally, retrieved objects should be ranked according to their similarity to the query. Note that retrieval performance on a relatively small non-rigid dataset (typically with hundreds of objects) can often be sensitive to particular attributes of the dataset [31], such as the mesh resolution and total surface areas that are largely uninteresting regarding a benchmark study. As a complementary qualitative example, we show our shape signature has a good performance with respect to varied 2D shapes. Compared to 3D shapes, 2D shapes have less geometric features (where we empirically choose less dimensions for shape signatures in experiments) and their boundary can be very noisy, hence global shapes play a key role in similarity based retrieval. We gathered the dataset of 2D contour shapes⁶ used in [44], which contains four subsets (fighter planes, vehicles and two subsets of MPEG-7 CE Shape-1 Part-B, 590 images in total). We used existing software to triangulate these contour shapes, have them cleaned as connected manifold meshes with at least 6k vertices. Then we performed query-based shape retrieval. Note that class labels of these 2D shapes were not used. A comparative example between our method and shapeDNA [35] is given in Figure 4, which shows a ranked list of retrieved shapes when the first shape is used as the query shape. It is observed from such results that our method is able to retrieve shapes according to their global shapes, and the similarity between the query shape and a retrieved shape in the ranked list continuously decreases when we move towards the end of the list.

As for stability analysis, we will show the stability of shape retrieval results when the query shape undergoes deformations. Here we show by examples (Figure 5) that in addition to deformations, our method is also stable when noise and small holes are added to the shape instances. In our experiments, varying degrees of zero mean white noise (.1 ~ .5 average edge length) were added to vertex positions along their normal direction and a certain percentage (1% ~ 3%) of faces were removed from the query meshes. The retrieval results turned out to be even more stable than those under deformations.


4.5 Isospectral vs Isometric.

Shape-DNA is known to be identical on isospectral manifolds, which is a family theoretically larger than isometrics. In fact, it has already been proven that one cannot “hear the shape of drum”. There are many examples of isospectral manifolds which are not isometric. The first example was given in 1964, and later mathematicians proposed several general construction techniques [41] to find non-isometric isospectral manifolds in two and three dimensions. Perhaps the simplest non-isometric isospectral shapes are GWW-prisms [16] (Figure 6).

⁶This dataset can be downloaded at http://visionlab.uta.edu/shape_data.htm



(a) Retrieval results by shapeDNA, with 8 eigenvalues



(b) Retrieval results by R-BiHDM, with $L = 15, m = 60$

Figure 4: The result of a query (the first shape) in a 2D shape dataset, which contains similar shapes under several types of deformations, including bending, torsion, and partial scaling.

Shape-DNA has self-intersections in the spectral space when a shape undergoes non-isometric shape deformations, therefore may not be sufficiently discriminative near those intersections. In contrast, spectra of the reduced biharmonic distance matrices are clearly different for those non-isometric but isospectral pairs [10]. It is of interest to ask if it indicates an isometry that two 2D/3D manifolds have identical R-BiHDM, i.e. not only their λ_i , but also a_j in eq. (5) are the same.

4.6 Limitations

Our method has a few limitations that need further investigation. Although it has strength in identifying global non-rigid shapes, how to enhance our method for retrieving meaningful partial shapes still remains unknown. Furthermore, it assumes the surface, which is a 2D manifold, is connected. It is unclear how to extend our method for matching and retrieving complicated models with many disconnected parts.

5 Conclusions

In this paper, we have first given a brief introduction to existing methods for nonrigid 3D shape retrieval. Our method uses bi-harmonic distance to construct a context-aware integral kernel

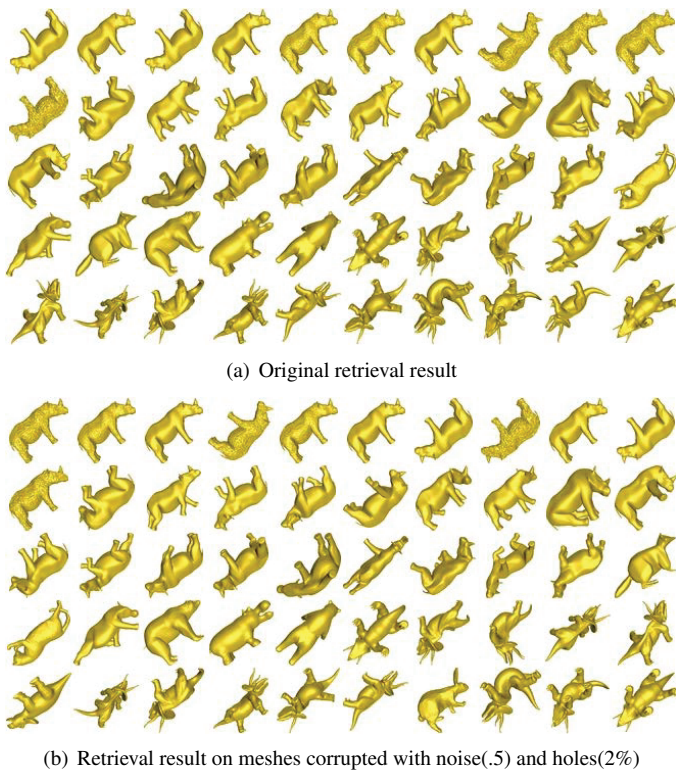


Figure 5: Stability test (the first shape is the query). Our retrieval results remain stable even when the meshes are corrupted with noise and holes. (Set $L = 23$, $m = 60$)


operator on a manifold, then applies modal space restriction to project this operator into a low-frequency representation, and finally computes its spectrum. Our method is simple to implement, isometry-invariant, discriminative, and numerically stable with respect to multiple types of perturbations. Our current implementation is based on FEM. We have evaluated our proposed method on representative datasets, including both the nonrigid track of SHREC'11 and the nonrigid human track of SHREC'14. Evaluation results demonstrate that our shape signature is highly effective in nonrigid shape retrieval.

Acknowledgments

The authors would like to thank Maks Ovsjanikov and Dirk Smeets for sharing their software implementations.

References

- [1] Y. Aflalo and R. Kimmel. Spectral multidimensional scaling. *Proceedings of the National Academy of Sciences*, 110(45):18052–18057, 2013.
- [2] A. Agathos, I. Pratikakis, P. Papadakis, S. Perantonis, P. Azariadis, and N.S. Sapidis. 3d articulated object re-



	shape-DNA		R-BiHDM	
0	3.39832620	3.39691613	2.16530711e-1	2.20749059e-1
1	13.0603602	13.0625695	-1.75187585e-1	-1.75999207e-1
2	16.9851450	16.9743079	-2.23163050e-2	-2.96890457e-2
3	29.8711391	29.8533596	-1.12679120e-2	-7.17493004e-3
4	39.6051131	39.6049419	-2.24496696e-3	-2.25092504e-3
5	43.8536891	43.8563295	-1.27777594e-3	-1.28146228e-3
6	54.7970492	54.6495812	-1.04506137e-3	-1.15918545e-3
7	70.9875083	71.0145921	-6.95901459e-4	-6.89618397e-4
8	79.7177678	79.7141696	-3.98584355e-4	-4.10728542e-4
9	91.4966416	91.7588846	-3.67194501e-4	-3.62850271e-4

Figure 6: Shape-DNA and R-BiHDM spectra computed with linear FEM for GWW prisms ($\sim 1k$ vertices).

trieval using a graph-based representation. *The Visual Computer*, 26(10):1301–1319, 2010.

- [3] A.Y. Alfakih. On the eigenvalues of euclidean distance matrices. *Computational and Applied Mathematics*, 27(3):237–250, 2008.
- [4] W.E. Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Quart. Appl. Math.*, 9(1):17–29, 1951.
- [5] M. Attene and B. Falcidieno. Remesh: An interactive environment to edit and repair triangle meshes. In *Shape Modeling and Applications, 2006. SMI 2006. IEEE International Conference on*, pages 41–41. IEEE, 2006.
- [6] T. Aubin. *Some nonlinear problems in Riemannian geometry*. Springer, 1998.
- [7] E. Bogomolny, O. Bohigas, and C. Schmit. Spectral properties of distance matrices. *Journal of Physics A: Mathematical and General*, 36(12):3595, 2003.
- [8] A.M. Bronstein, M.M. Bronstein, L.J. Guibas, and M. Ovsjanikov. Shape google: geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics (TOG)*, 30(1):1, 2011.
- [9] M.M. Bronstein and A.M. Bronstein. Shape recognition with spectral distances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5):1065–1071, 2011.
- [10] P. Buser, J. Conway, P. Doyle, and K.D. Semmler. Some planar isospectral domains. *International Mathematics Research Notices*, 1994(9):391, 1994.

- [11] M. Castro and V. Menegatto. Eigenvalue decay of positive integral operators on the sphere. *Mathematics of Computation*, 81(280):2303–2317, 2012.
- [12] P. Cignoni, M. Corsini, and G. Ranzuglia. Meshlab: an open-source 3d mesh processing system. *ERCIM News*, 73:45–46, 2008.
- [13] A. Elad and R. Kimmel. Bending invariant representations for surfaces. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages 1–168. IEEE, 2001.
- [14] K.O. Friedrichs. *Perturbation of spectra in Hilbert space*, volume 3. American Mathematical Soc., 1965.
- [15] A. Giachetti and C. Lovato. Radial symmetry detection and shape characterization with the multiscale area projection transform. *Computer Graphics Forum*, 31(5):1669–1678, 2012.
- [16] C. Gordon, D. Webb, and S. Wolpert. Isospectral plane domains and surfaces via riemannian orbifolds. *Inventiones mathematicae*, 110(1):1–22, 1992.
- [17] J. Hu and J. Hua. Salient spectral geometric features for shape matching and retrieval. *The visual computer*, 25(5-7):667–675, 2009.
- [18] T. Kato. *Perturbation theory for linear operators*, volume 132. springer, 1995.
- [19] S. Kurtek, A. Srivastava, E. Klassen, and H. Laga. Landmark-guided elastic shape analysis of spherically-parameterized surfaces. *Computer Graphics Forum*, 32(2pt4):429–438, 2013.
- [20] H. Laga, M. Mortara, and M. Spagnuolo. Geometry and context for semantic correspondences and functionality recognition in man-made 3d shapes. *ACM Transactions on Graphics (TOG)*, 32(5):150, 2013.
- [21] G. Lavoué. Combination of bag-of-words descriptors for robust partial shape retrieval. *The Visual Computer*, 28(9):931–942, 2012.
- [22] C. Li and A.B. Hamza. Intrinsic spatial pyramid matching for deformable 3d shape retrieval. *International Journal of Multimedia Information Retrieval*, 2(4):261–271, 2013.
- [23] C. Li and A.B. Hamza. A multiresolution descriptor for deformable 3d shape retrieval. *The Visual Computer*, 29(6-8):513–524, 2013.
- [24] Z. Lian, A. Godil, B. Bustos, M. Daoudi, J. Hermans, S. Kawamura, Y. Kurita, G. Lavoue, H. van Nguyen, R. Ohbuchi, Y. Ohkita, Y. Ohishi, F. Porikli, M. Reuter, I. Sipiran, D. Smeets, P. Suetens, H. Tabia, and D. Vandermeulen. A comparison of methods for non-rigid 3d shape retrieval. *Pattern Recognition*, 46(1):449–461, 2013.
- [25] V.B. Lidskii. Perturbation theory of non-conjugate operators. *USSR Computational Mathematics and Mathematical Physics*, 6(1):73–85, 1966.
- [26] Y. Lipman, R. M. Rustamov, and T. A. Funkhouser. Biharmonic distance. *ACM Transactions on Graphics, ACM 0730-0301//10-ART*, 2007.
- [27] R. Litman, A. Bronstein, M. Bronstein, and U. Castellani. Supervised learning of bag-of-features shape descriptors using sparse coding. *Computer Graphics Forum*, 33(5):127–136, 2014.
- [28] D.G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- [29] C. Maes, T. Fabry, J. Keustermans, D. Smeets, P. Suetens, and D. Vandermeulen. Feature detection on 3d face surfaces for pose normalisation and recognition. In *Biometrics: Theory Applications and Systems (BTAS), 2010 Fourth IEEE International Conference on*, pages 1–6. IEEE, 2010.
- [30] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30, 2012.
- [31] D. Pickup, X. Sun, P.L. Rosin, R.R. Martin, Z. Cheng, Z. Lian, M. Aono, A.B. Hamza, A. Bronstein, M. Bronstein, et al. Shrec’14 track: Shape retrieval of non-rigid 3d human models. *Proc. 3DOR*, 4(7):8, 2014.
- [32] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical mathematics*, volume 37. Springer Verlag, 2007.
- [33] Dan Raviv and Ron Kimmel. Affine invariant geometry for non-rigid shapes. *International Journal of Computer Vision*, 111(1):1–11, 2015.
- [34] J.B. Reade. Eigenvalues of positive definite kernels. *SIAM Journal on Mathematical Analysis*, 14(1):152–157, 1983.
- [35] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-beltrami spectra as ‘shape-DNA’ of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.
- [36] R.M. Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 225–233. Eurographics Association, 2007.

- [37] P. Shilane, P. Min, M.M. Kazhdan, and T.A. Funkhouser. The princeton shape benchmark. In *Shape Modeling International*, pages 167–178. IEEE Computer Society, 2004.
- [38] D. Smeets, T. Fabry, J. Hermans, D. Vandermeulen, and P. Suetens. Isometric deformation modelling for object recognition. *CAIP 2009, LNCS 5702*, pages 757–765, 2009.
- [39] G. W. Stewart. A Krylov–Schur algorithm for large eigenproblems. *SIAM Journal on Matrix Analysis and Applications*, 23(3):601–614, 2001.
- [40] J. Sun, M. Ovsjanikov, and L.J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Comput. Graph. Forum*, 28(5):1383–1392, 2009.
- [41] T. Sunada. Riemannian coverings and isospectral manifolds. *The Annals of Mathematics*, 121(1):169–186, 1985.
- [42] H. Tabia, M. Daoudi, J.-P. Vandeborre, and O. Colot. A new 3d-matching method of nonrigid and partially similar models using curve analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):852–858, 2011.
- [43] H. Tabia, H. Laga, D. Picard, and P.-H. Gosselin. Covariance descriptors for 3d shape matching and retrieval. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 4185–4192. IEEE, 2014.
- [44] N. Thakoor, J. Gao, and S. Jung. Hidden markov model-based weighted likelihood discriminant for 2-d shape classification. *Image Processing, IEEE Transactions on*, 16(11):2707–2719, 2007.
- [45] H. Weyl. Über die asymptotische verteilung der eigenwerte. *Nachrichten der Königlichen Gesellschaft der Wissenschaften zu Göttingen. Mathematisch-Naturwissenschaftliche Klasse*, pages 110–117, 1911.
- [46] J. Ye, Z. Yan, and Y. Yu. Fast nonrigid 3d retrieval using modal space transform. In *International Conference on Multimedia Retrieval*, pages 121–126, 2013.
- [47] X. Ying, X. Wang, and Y. He. Saddle vertex graph (svg): A novel solution to the discrete geodesic problem. *ACM Trans. Graph.*, 32(6):170:1–170:12, November 2013.
- [48] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 373–380. Ieee, 2009.
- [49] C.T. Zahn and R.Z. Roskies. Fourier descriptors for plane closed curves. *Computers, IEEE Transactions on*, 100(3):269–281, 1972.
- [50] O.C. Zienkiewicz, R.L. Taylor, and J.Z. Zhu. *The Finite Element Method—Its Basis and Fundamentals, volume 1*. Elsevier Butterworth-Heinemann, Amsterdam, London,, 2005.

A Error Bounds under Deformation

In this section, we investigate error bounds of the spectral signature of integral kernel operators in the presence of “smooth” manifold deformations using perturbation theory [14, 18]. The general result, intuitively speaking, is that the total deviation of a spectral signature with an arbitrary number of dimensions is finitely bounded under smooth deformations. Given a general matrix A and a perturbation parameter ϵ within a neighborhood of zero, the perturbed matrix is written as $A + \epsilon B$, where B is an arbitrary matrix. In this situation, each eigenvalue or eigenvectors of $A + \epsilon B$ admits an expansion in fractional powers of ϵ , and the zero-th order term of this expansion is an eigenvalue or eigenvector of the unperturbed matrix A (Lidskii theorem, 1965 [25]). This is a well known result in regular perturbation theory. In particular, one can actually derive a Lipschitz condition number for this continuity:

Theorem A.1 (A special case of Lidskii theorem [18]). *Assuming $\mu(\epsilon)$ is an eigenvalue of the perturbed matrix $A + \epsilon B$ for a sufficiently small ϵ , it admits a first-order expansion*

$$\mu(\epsilon) = \mu + \|\Phi^T B \Phi\|_{\max} \epsilon + o(\epsilon),$$

where Φ is the eigenvector of A corresponding to eigenvalue μ , A and B are conjugated matrices, $\|\cdot\|_{\max}$ represents the largest singular value. The first-order coefficient $\|\Phi^T B \Phi\|_{\max}$ is called the Lipschitz condition number.

For the rest of this section, A is called *characterization*, and B is called *perturbation*. The Lipschitz condition number characterizes the total amount of deviation in a signature with respect to a certain amount of perturbation. In the ideal case of isometric perturbation, where A is isometric and B is zero, the Lipschitz condition number is degenerated. Let M and M' be two compact manifold domains, and H and H' be the spaces of bounded and continuous linear functionals defined on M and M' . In the current context, we consider M' as a “gently” deformed version of M . We assume that implicitly there exists a matching (or registration) between these two domains such that it induces a mapping between H and H' , which is a closed and densely defined operator \mathcal{D} from H to H' . The characterization of this mapping is a self-adjoint operator \mathcal{O} . Suppose $\mathcal{O}'(\epsilon)$ has a first-order approximate perturbed (infinitely-dimensional) matrix representation $A + \epsilon B$ with respect to an underlying domain deformation \mathcal{D}_ϵ .

Definition A.1 (perturbed (α, p) -smoothness). Consider a perturbed matrix $A + \epsilon B$ with a spectral family $\{\mu_i(\epsilon)\}$ admitting first-order expansions: $\mu_i(\epsilon) = \mu_i + |\phi_i^T B \phi_i| \epsilon + o(\epsilon)$. $A + \epsilon B$ is perturbed (α, p) -smooth if

$$\|A^{-\frac{\alpha}{2}} B^p A^{-\frac{\alpha}{2}}\|_1 =: \sum_{i=1}^{\infty} |\mu_i|^{-\alpha} \|\phi_i^T B \phi_i\|^p < C(\alpha, p),$$

for some $\alpha \geq 0$ and $p > 0$, where $\|\cdot\|_1$ is Schatten p -norms.

Perturbed smoothness actually admits a substantially larger class of deformations of practical interest. It allows the characterization to be varied with respect to the domain deformation, but it bounds the amount of deviation.

Theorem A.2 (Error bound of spectral signature). Consider a discrepancy function for a pair of original and perturbed spectra

$$d_n(\mu, \mu'; \alpha, p) := \sum_{i=1}^n |\mu_i|^{-\alpha} |\mu_i - \mu'_i|^p,$$

and two Hilbert spaces H and H' admitting a perturbation path \mathcal{D}_ϵ for $\epsilon \in [0, 1]$. If the perturbation path is uniformly perturbed (α, p) -smooth with a bound $C(\alpha, p)$ with respect to a characterization operator $\mathcal{O}(\epsilon)$, and order preserving, aka $\mu_i \geq \mu_j \iff \mu'_i \geq \mu'_j$ for any i, j , we have

$$d_n(\mu, \mu'; \alpha, p) \leq \kappa^{-\alpha} C(\alpha, p), \forall n$$

where $\kappa = \max_i \left\{ \left| \frac{\mu_i}{\mu_{i+1}} \right|, \left| \frac{\mu_{i+1}}{\mu_i} \right| \right\}$.

Proof. (Sketch) Take the integral w.r.t $\epsilon: \mu'_i - \mu_i = \int_0^1 d\mu_i(\epsilon)$ and use the condition of perturbed (α, p) -smoothness (Definition A.1). \square

Integral Kernel Operator. Experimental evidences suggest that distance maps/matrices, more generally, integral operators, could be a useful source for shape descriptor construction, and are potentially better than differential operators (e.g., LBO). As a rigorous treatment justifying the capability of integral kernel operators, we establish another definition for the smoothness of perturbations:

Definition A.2 (α -Hölder class of perturbed kernel). Consider integral kernel operator $k_\epsilon(\cdot, \cdot)$ with perturbation parameter ϵ , and define the perturbed kernel

$$\gamma_\epsilon(x, y; k) = \frac{1}{\epsilon} (k_\epsilon(\tilde{x}, \tilde{y}) - k(x, y)),$$

where $(\tilde{x}, \tilde{y}) \in M'(\epsilon)$ is the perturbed version of $(x, y) \in M$. Let d_M be the geodesic distance on M . If there exists $C > 0$ such that

$$|\gamma_\epsilon(x, y; k) - \gamma_\epsilon(x, z; k)| \leq C_\alpha \cdot d_M(y, z)^\alpha,$$

and

$$|\gamma_\epsilon(x, y; k)| \leq C_\alpha,$$

$\forall x \in M$ and any $\epsilon > 0$ in a sufficiently small neighborhood of zero.

Remark. If $k(\cdot, \cdot) = d_M(\cdot, \cdot)$, the infimum of bound C_α that makes k_ϵ 's perturbed kernel belong to α -Hölder class, aka

$$C_\alpha^* := \sup_{x, y, z, \epsilon} \left\{ \frac{|\gamma_\epsilon(x, y; k) - \gamma_\epsilon(x, z; k)|}{d_M^\alpha(\cdot, \cdot)}, |\gamma_\epsilon(x, y; k)| \right\},$$

is a ‘‘natural’’ quantity indicator of deformations.

In the sequel, we connect α -Hölder class of perturbed kernel with perturbed smoothness.

Theorem A.3 (main result). Given an integral kernel operator \mathcal{K} as described in Section 2.2, we assume the distance map is in accordance with the geodesics of power q , aka, there exist positive constants C_1, C_2 and some $\delta \geq 0$ such that

$$C_1/n^{1+q+\delta} \leq \lambda_n(k) \leq C_2/n^{1+q}, \forall n, \quad (12)$$

for some integer $q \geq 1$.

If its perturbed kernels with respect to a class of deformations, i.e. $\gamma_\epsilon(\cdot, \cdot; k)$, are of ϵ -Hölder class (Definition A.2) for some $1 \geq \epsilon > 0$, \mathcal{K} or the transformed operator $\tilde{\mathcal{K}} = \Phi^T \mathcal{K} \Phi$ (Section 2.2 Equation 3), as a characterization of the deformed manifold $M'(\epsilon)$, is perturbed (α, p) -smooth (Definition A.1) as long as

$$\left(\frac{1}{2} + \epsilon\right)p > (1 + q + \delta)\alpha + 1.$$

Thus, spectral signatures (with an arbitrary number of dimensions) of both \mathcal{K} and $\tilde{\mathcal{K}}$ have finite error bounds under smooth manifold deformations (Theorem A.2).

Proof. For a general Hilbert-Schmidt operator k satisfying uniformly ϵ -Hölder condition, it has already been shown in [34] that the decay rate of eigenvalues of $k(x, y)$ is at least $O(n^{-1/2-\epsilon})$. And if k is positive definite, the decay rate is improved to at least $O(n^{-1-\epsilon})$. It immediately follows that for $\epsilon > \frac{1}{2}$, ϵ -Hölder continuous kernel γ_ϵ is perturbed $(0, 1)$ -smooth.

For $\alpha > 0$, we require the eigenvalues of the characterization kernel have lower bounds on its decay rate. It is generally not true, for example if $k(x, y) = d_M(x, y)^2$, the resulting Gram matrix has a finite number of nonzero eigenvalues if sample points are uniformly drawn from the Euclidean space. But it is typically tangible, as preliminary work has shown examples where condition (12) holds when $k(x, y) = d_M(x, y)$ (e.g. see [7, 11] examples). Then following the main result on Hölder condition and the decay of eigenvalues [34], one has $\frac{|\lambda_n(\gamma_\epsilon)|^p}{|\lambda_n(k)|^\alpha} \sim O(n^{-(\frac{1}{2}+\epsilon)p+\alpha(1+q+\delta)})$. \square

B Spectral Convergence of R-BiHDM

Since we know eigenvalues of the Laplace-Beltrami operator follow $\lambda_m \sim \frac{4\pi}{A}m$ for surfaces (Weyl-Asymptotic growth of eigenvalue [45]), we can therefore deduce the convergence of a spectral sequence under weak assumptions. Note $\|\cdot\|$ always denotes the L_2 -norm of a vector, a function or a matrix, by default.

Lemma 1. *Let $A \in \mathbb{C}^{n \times n}$ be an Hermitian matrix, and let $(\hat{\lambda}, \hat{x})$ be the computed approximation of an eigenvalue/eigenvector pair (λ, x) of A . By defining the residual*

$$\hat{r} = A\hat{x} - \hat{\lambda}\hat{x}, \hat{x} \neq 0,$$

it follows that

$$\min_{\lambda_i \in \sigma(A)} |\hat{\lambda} - \lambda_i| \leq \frac{\|\hat{r}\|_2}{\|\hat{x}\|_2}.$$

Proof. See [32] page 195. \square

Theorem B.1 (Spectral Convergence). *Let \tilde{K}_m be an $(m+1) \times (m+1)$ square matrix formed by the following steps: i) take the first $m+1$ rows and first $m+1$ columns of \tilde{K} ; ii) when calculating each a_j in this truncated matrix, every infinite summation in (5) is approximated by the first m terms. Set $\mu_i^{(m)}$ be the i -th eigenvalue of \tilde{K}_m , for $i \leq m+1$ and $\mu_i(m) = 0$, for any $i > m+1$. Let $f_m = \sqrt{A} \sum_{i=1}^m \phi_i^2 / \lambda_i^2$ converge (pointwisely) to f . If f is bounded and square integrable, i.e. eigenpairs $\{(\lambda_i, \phi_i)\}_{i=1}^\infty$ satisfy*

$$\sum_{i=1}^{\infty} \frac{\|\phi_i^2\|^2}{\lambda_i^4} < \infty, \quad (13)$$

it then follows that for any i ,

$$\mu_i^{(m)} \rightarrow \mu_i, \quad \text{as } m \rightarrow \infty.$$

Proof. By definition, we have $\tilde{K}_{m+n} - \tilde{K}_{m-1} =$

$$\sum_{k=m}^{m+n} \begin{bmatrix} \frac{2}{\lambda_k^2} & \frac{C_{k,1}}{\lambda_k^2} & \dots & \frac{C_{k,m-1}}{\lambda_k^2} & 0 & \dots & a_k^{(m+n)} & \dots \\ \frac{C_{k,1}}{\lambda_k^2} & 0 & & & & & & \\ \vdots & & \ddots & & & & & \\ \frac{C_{k,m-1}}{\lambda_k^2} & & & 0 & & & & \\ 0 & & & & 0 & & & \\ \vdots & & & & & \ddots & & \\ a_k^{(m+n)} & & & & & & -\frac{2}{\lambda_k^2} & \\ \vdots & & & & & & & \ddots \end{bmatrix}$$

where extra rows and columns of zeros are padded to \tilde{K}_{m-1} , and

$$C_{m,j} = \sqrt{A} \int_M \phi_m^2 \phi_j,$$

which is the Fourier coefficient of $\sqrt{A}\phi_m^2$ w.r.t series $\{\phi_j\}_{j=1}^{m-1}$. By Bessel's inequality we have $1 + \sum_{j=1}^{m-1} C_{m,j}^2 \leq A \|\phi_m^2\|^2$, and we also have

$$\begin{aligned} \sum_{k=m}^{m+n} |a_k^{(m+n)}|^2 &\leq \sum_{k=m}^{\infty} |\int_M f_{m+n} \phi_k|^2 \\ &\leq \|f_{m+n} - f\|^2 + \sum_{k=m}^{\infty} |\int_M f \phi_k|^2. \end{aligned}$$

By assumptions about f and the norm convergence of Fourier series (as shown for example in [1]), we know that $\|f_{m+n} - f\|^2 \rightarrow 0$ and $\sum_{k=m}^{\infty} |\int_M f \phi_k|^2 \rightarrow 0$ as $m \rightarrow \infty$. Since the L_2 -norm of a matrix is less than its Frobenius norm, we have

$$\begin{aligned} &\left\| \tilde{K}_{m+n} - \tilde{K}_{m-1} \right\|_2^2 \\ &\leq \sum_{k=m}^{m+n} \left(\frac{2}{\lambda_k^4} (4 + \sum_{j=1}^{m-1} C_{k,j}^2) + 2 |a_k^{(m+n)}|^2 \right), \\ &\leq \sum_{k=m}^{m+n} \left(\frac{2}{\lambda_k^4} (3 + A \|\phi_k^2\|^2) + 2 |a_k^{(m+n)}|^2 \right) \rightarrow 0 \end{aligned}$$

as $m \rightarrow \infty$. By Lemma 1, we have

$$\begin{aligned} |\mu_i^{(m+n)} - \mu_i^{(m-1)}| &\leq \frac{\left\| (\tilde{K}_{m+n} - \tilde{K}_{m-1}) u_i^{(m+n)} \right\|}{\left\| u_i^{(m+n)} \right\|} \\ &\leq \left\| \tilde{K}_{m+n} - \tilde{K}_{m-1} \right\|_2 \end{aligned}$$

for any $i < m$, where $u_i^{(m)}$ is the eigenvector associated with $\mu_i^{(m)}$ of \tilde{K}_m . $\{\mu_i^{(m)}\}_{m=1}^\infty$ is a Cauchy sequence, that converges. \square

Note that the condition, i.e. Eq. (13), used in the spectral convergence theorem is rather weak. Observe that for two-dimensional Riemannian manifold M , $f(x)$ is given by the green function $G(x, y)$ with $f(x) := G(x, x)$. If M is a compact manifold (without boundary), it satisfies that $G(x, y)$ is bounded [6], hence $f(x)$ is also bounded for any $x \in M$. So is $\|f\|^2$ (square integrable). An experimental validation is also provided in Fig. 7.

C Matrices for Linear and Cubic FEM

Here we provide elementary matrices used in linear and cubic FEM. Every FEM has its own set of nodes. There is a finite element associated with every node. The finite element at a specific node is a piecewise polynomial basis function whose value is equal to 1 at the node and 0 at all other nodes. The support of a basis function includes all triangle faces surrounding the node corresponding to the basis function. Consider a triangle face in a mesh with two node configurations shown in Figure 8. These node configurations are defined for linear and cubic FEM, respectively. Each node in either configuration is associated with a piecewise linear or cubic polynomial




Figure 7: A experimental validation: convergence of the first 30 eigenvalues of an R-BiHDM.

basis function. If we focus on a triangle face, an arbitrary bivariate polynomial over the triangle face can be defined as a linear combination of the basis functions associated with the nodes either inside or on the edges of the triangle. This polynomial can interpolate arbitrary prescribed values at the nodes. It reduces to a univariate polynomial along each triangle edge. The two polynomials defined over two adjacent triangles agree with each other along their shared triangle edge. The stiff and mass matrices used in FEM are given as follows.




Figure 8: Node configuration of linear and cubic FEM

$$\begin{aligned}
J_1 &= \frac{1}{8} \begin{pmatrix} 2 & -1 & -1 \\ -1 & 0 & 1 \\ -1 & 1 & 0 \end{pmatrix} & J_2 &= \frac{1}{8} \begin{pmatrix} 0 & -1 & 1 \\ -1 & 2 & -1 \\ 1 & -1 & 0 \end{pmatrix} \\
J_3 &= \frac{1}{8} \begin{pmatrix} 0 & 1 & -1 \\ 1 & 0 & -1 \\ -1 & -1 & 2 \end{pmatrix} & J_4 &= \frac{1}{12} \begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix}
\end{aligned}$$

Table 6: The four matrices of integrals on the unit triangle for linear FEM.

$$\begin{aligned}
L_{ij}(T) &= \frac{1}{a(T)} \left(\sum_{k=1}^3 l_k^2(T) J_k(t_i, t_j) \right), & i, j &\in N(T) \\
D_{ij}(T) &= a(T) J_4(t_i, t_j), & i, j &\in N(T) \\
L_{ij}(T) &= D_{ij}(T) = 0, & i, j &\notin N(T)
\end{aligned}$$

$$\begin{aligned}
J_1 &= \frac{1}{320} \begin{pmatrix} 0 & 7 & -7 & 57 & -24 & 0 & 0 & 24 & -57 & 0 \\ 7 & 0 & -7 & -24 & 57 & -57 & 24 & 0 & 0 & 0 \\ -7 & -7 & 68 & -6 & -6 & 30 & -51 & -51 & 30 & 0 \\ 57 & -24 & -6 & 135 & 54 & 27 & 27 & 27 & -135 & -162 \\ -24 & 57 & -6 & 54 & 135 & -135 & 27 & 27 & 27 & -162 \\ 0 & -57 & 30 & 27 & -135 & 135 & -108 & -27 & -27 & 162 \\ 0 & 24 & -51 & 27 & 27 & -108 & 135 & 135 & -27 & -162 \\ 24 & 0 & -51 & 27 & 27 & -27 & 135 & 135 & -108 & -162 \\ -57 & 0 & 30 & -135 & 27 & -27 & -27 & -108 & 135 & 162 \\ 0 & 0 & 0 & -162 & -162 & 162 & -162 & -162 & 162 & 324 \end{pmatrix} \\
J_2 &= \frac{1}{320} \begin{pmatrix} 0 & -7 & 7 & -57 & 24 & 0 & 0 & -24 & 57 & 0 \\ -7 & 68 & -7 & 30 & -51 & -51 & 30 & -6 & -6 & 0 \\ 7 & -7 & 0 & 0 & 0 & 24 & -57 & 57 & -24 & 0 \\ -57 & 30 & 0 & 135 & -108 & -27 & -27 & 27 & -135 & 162 \\ 24 & -51 & 0 & -108 & 135 & 135 & -27 & 27 & 27 & -162 \\ 0 & -51 & 24 & -27 & 135 & 135 & -108 & 27 & 27 & -162 \\ 0 & 30 & -57 & -27 & -27 & -108 & 135 & -135 & 27 & 162 \\ -24 & -6 & 57 & 27 & 27 & -135 & 135 & 54 & -162 \\ 57 & -6 & -24 & -135 & 27 & 27 & 54 & 135 & -162 \\ 0 & 0 & 0 & 162 & -162 & -162 & 162 & -162 & -162 & 324 \end{pmatrix} \\
J_3 &= \frac{1}{320} \begin{pmatrix} 68 & -7 & -7 & -51 & 30 & -6 & -6 & 30 & -51 & 0 \\ -7 & 0 & 7 & 24 & -57 & 57 & -24 & 0 & 0 & 0 \\ -7 & 7 & 0 & 0 & 0 & -24 & 57 & -57 & 24 & 0 \\ -51 & 24 & 0 & 135 & -108 & 27 & 27 & -27 & 135 & -162 \\ 30 & -57 & 0 & -108 & 135 & -135 & & 27 & -27 & -27 & 162 \\ -6 & 57 & -24 & 27 & -135 & 135 & 54 & 27 & 27 & -162 \\ -6 & -24 & 57 & 27 & 27 & 54 & 135 & -135 & 27 & -162 \\ 30 & 0 & -57 & -27 & -27 & 27 & -135 & 135 & -108 & 162 \\ -51 & 0 & 24 & 135 & -27 & 27 & -108 & 135 & -162 \\ 0 & 0 & 0 & -162 & 162 & -162 & -162 & 162 & -162 & 324 \end{pmatrix} \\
J_4 &= \frac{1}{6720} \begin{pmatrix} 76 & 11 & 11 & 18 & 0 & 27 & 27 & 0 & 18 & 36 \\ 11 & 76 & 11 & 0 & 18 & 18 & 0 & 27 & 27 & 36 \\ 11 & 11 & 76 & 27 & 27 & 0 & 18 & 18 & 0 & 36 \\ 18 & 0 & 27 & 540 & -189 & -135 & -54 & -135 & 270 & 162 \\ 0 & 18 & 27 & -189 & 540 & 270 & -135 & -54 & -135 & 162 \\ 27 & 18 & 0 & -135 & 270 & 540 & -189 & -135 & -54 & 162 \\ 27 & 0 & 18 & -54 & -135 & -189 & 540 & 270 & -135 & 162 \\ 0 & 27 & 18 & -135 & -54 & -135 & 270 & 540 & -189 & 162 \\ 18 & 27 & 0 & 270 & -135 & -54 & -135 & -189 & 540 & 162 \\ 36 & 36 & 36 & 162 & 162 & 162 & 162 & 162 & 162 & 1944 \end{pmatrix}
\end{aligned}$$

Table 7: The four matrices of integrals on the unit triangle for cubic FEM.

and

$$L = \sum_T L(T), \quad D = \sum_T D(T),$$

where T denotes a triangle face, $N(T)$ denotes the set of incidence nodes on T , t_i denotes the local index (used for T) of the i -th node used in the FEM, $l_k(T)$ ($k = 1, 2, 3$) denotes the length of the k -th edge of T , $a(T)$ denotes the area of T , and $J_k(t_i, t_j)$ denotes the entry with indices (t_i, t_j) of the matrix J_k . J_k ($k = 1, 2, 3, 4$) for linear and cubic FEM are given in Table 6 and Table 7 respectively. Details on the construction of J_k can be found in [50]. Note that these matrices are assembled triangle-wise rather than element-wise in a typical implementation.