Fast Nonrigid 3D Retrieval Using Modal Space Transform

Jianbo Ye
The University of Hong Kong
bobye@hku.hk

Zhicheng Yan
University of Illinois at Urbana-Champaign
zyan3@illinois.edu

Yizhou Yu
The University of Hong Kong
yizhouy@acm.org

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 first analyze the advantages of functional operators, and further propose a framework to design novel shape signatures for encoding nonrigid object structures. 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. Our method is fast, isometry-invariant, discriminative, and numerically stable with respect to multiple types of perturbations.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*Retrieval models*; H.4 [**Information Systems Applications**]: Miscellaneous; I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling—*Geometric algorithms, languages, and systems*

Keywords

Content-Based Object Retrieval, Shape Retrieval, Biharmonic Distance, Functional Map, Shape Signature

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. Important criteria for measuring the performance of nonrigid object retrieval techniques include *isometry invariance*, *discrimination power*, *efficiency* and *stability*. In addition, there exist other considerations, including *extensibility*, *applicability* and *complexity of implementation*.

In this paper, we first analyse functional operators over modal space and further introduce spectrum-based shape signatures to encode the structure of a nonrigid shape. 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. To further boost retrieval performance, we incorporate a pseudo-relevance feedback mechanism to iteratively improve similarity ranking among retrieved object instances.

Our experiments demonstrate that our new shape signature for nonrigid objects can outperform all non-hybrid methods participating in the nonrigid track of the SHREC'11 contest [9] on a representative mixed dataset combining the dataset from the nonrigid track of SHREC'11 and another custom built dataset. In particular, our method achieves obviously higher precision than other methods when recall is above 50%.

The rest of the paper is organized as follows. Related work will be discussed in Section 1.1. In Section 2, we discuss the fundamental mathematics behind our signature design. In Section 3, we briefly sketch our numerical implementation. In Section 4, we present experimental results to validate our shape signature. Section 5 concludes our paper.

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 [8], where researchers began to consider bending invariant or insensitive 3D shape recognition.

In our knowledge, two major classes of approaches are proposed for nonrigid shape retrieval during the last decade. The first class

ICMR'13, April 16–20, 2013, Dallas, Texas, USA.

includes all local feature based approaches. Inspired by the success of the SIFT feature descriptor in image retrieval, researchers proposed analogous descriptors based on extrinsic curvatures and tangential fields, such as meshSIFT [12] and MeshHOG [25], for representing local features on mesh surfaces. ShapeGoogle [14, 4] emphasizes the robustness of association and classification, especially for objects with missing parts and topological noises. It integrates the local Heat Kernel Signature (HKS) [22, 6] with the bag-of-word framework. By extracting isometrically invariant dense point descriptors and quantizing them into binary codes, shapes are registered for efficient indexing, comparison and association.

The second class emphasizes coarser-scale or global structures of a 3D nonrigid shape. The skeleton-based method in [23] encodes the geometric and topological information in the form of a skeleton graph and uses graph matching to retrieve similar skeletons. Another method based on Gromov-Hausdorff distance has been proposed in [5] for matching nonrigid shapes. Statistical techniques, such as histograms and D2 distributions [18], are also popular in designing descriptors in respect of coarse-scale structures.

Among those methods concerning global shapes, spectrum-based techniques became popular in recent years. Shape-DNA [17] proposes to use the spectrum of the Laplace-Beltrami operator as an isometry-invariant shape descriptor. Another similar method, SD-GDM [20] proposes to compute a singular value decomposition (spectra) for the geodesic distance matrix, which seems to outperform shape-DNA [9]. 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 [24, 15], and a uniform mesh decimation is typically employed in practice to speed up this process.

Local features, such as SIFT [11] and its variants, have been successfully adopted in image retrieval. However, in the context of 3D nonrigid shape retrieval, local feature based methods have been outperformed by shape descriptors emphasizing coarser-scale structures. We provide the following explanations for this phenomenon. First, pixel values captured by a camera are related to photometric properties of real scenes and objects. Photometric properties tend to have higher frequencies than pure geometry. For example, there could be high-frequency texture patterns over a geometrically flat surface. Given such high-frequency photometric properties, 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, thus unsuitable for retrieval tasks.

2. OUR APPROACH

2.1 Laplace-Beltrami Operator

Shape-DNA [17] exploits eigenvalues to achieve an impressive shape retrieval performance, while there are a number of other methods, such as [18, 22], 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,

$$-\Delta_M u = \lambda u. \tag{1}$$

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 [13], 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 which are experimentally shown to be more robust in presence of their spectrums than Laplace-Beltrami operator.

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

$$\mathcal{K}f(y) = \int_M k(x,y)f(x)\mathrm{d}x. \tag{2}$$

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) [9], the retrieval performance of SD-GDM [20], 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.

The major important contribution of our approach is instead naively computing transforms matrix in complexity of geometry (i.e. pairwise 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 onto Φ). We can write the transform matrix by

$$\widetilde{K} = \Phi^T\mathcal{K}\Phi, \tag{3}$$

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

It is worth noting that this restriction is not a simple efficient approximation, the functional transforms before and after restriction are different in both theory and numerics. Precisely speaking, the original functional space before restriction is generally a Banach space (and is numerically approximated in a Sobolev space, i.e. $W_{q,p}(M) = \{f \in L^p(M) : D^\alpha f \in L^p(M), \forall\alpha \le q\}$[1]); while the functional space Φ in restriction is an infinite dimensional Hilbert space or inner product space (and is numerically approximated by truncating lower eigenvectors, converges in a weak sense). Hence the spectrums of $\mathcal{K}$ in these two functional spaces can be different (although they are identical for Laplace-Beltrami operator).

In our experiments (see comparison of track BiHDM and R-BiHDM in section 4.2), it is indicated spectrum computed by modal space restriction can better tolerate manifold deformations, and outperforms the spectrum of a pairwise kernel in retrieval tasks.

2.3 Distance Map Based on Modals

[1]Here, the notation $D^\alpha f = \frac{\partial^{|\alpha|} f}{\partial x_1^{\alpha_1} \ldots \partial x_i^{\alpha_i}}$, for $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_i)$

Computing the geodesic distance matrix is very computationally expensive for large meshes. We choose to compute (squared) biharmonic distance [10] 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,

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

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 $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \ldots$ 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 [10] 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.

Additionally note that we can have different weight choices instead of $1/\lambda_i^2$, e.g. $\lambda_i \exp(-\lambda_i t)$. It is of course interesting to justify specific choices. But in our experiments, we in "hard-code" choose biharmonic weights to achieve favorable results. See section **??** for a more extensive discussion.

2.4 Functional Biharmonic Distance Map

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

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

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

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

We also have

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

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 $\widetilde{K}$, called reduced biharmonic distance matrix(R-BiHDM),

$$\widetilde{K} = \begin{bmatrix} a_0 & a_1 & a_2 & \cdots \\ a_1 & -2/\lambda_1^2 & & \\ a_2 & & -2/\lambda_2^2 & \\ \vdots & & & \ddots \end{bmatrix}. \tag{6}$$

The above matrix is infinite. Let $\widetilde{K}_m$ be an $(m+1) \times (m+1)$ square matrix formed by the following two steps: i) take the first $m+1$ rows and first $m+1$ columns of $\widetilde{K}$; ii) when calculating each a_j in this truncated matrix, every infinite summation in (5) is approximated by the first m terms. As $m \to \infty$, the largest tens of eigenvalues of $\widetilde{K}_m$ enjoy quick convergence rate. Figure 1 shows the maximum error of the first 30 eigenvalues versus m, the number of eigenpairs of the Laplace-Beltrami operator. It is observed that the asymptotic eigenvalues converge linearly.

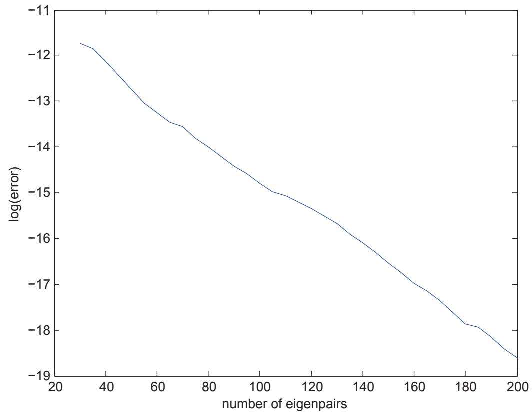

Figure 1: Convergence of the first 30 eigenvalues of an R-BiHDM.

Note that $\mathrm{tr}(\widetilde{K}) = 0$. We denote all eigenvalues of $\widetilde{K}$ 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 [1] and references therein) Hence a scale invariant spectrum can be defined as

$$\bar{\mu}_j = \left| \frac{\mu_j}{\mu_0} \right|. \tag{7}$$

Our shape signature is defined as a vector $S = [\bar{\mu}_1, \bar{\mu}_2, \ldots, \bar{\mu}_L]^T$, 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 [20]. 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

$$D_2^2 = \sum_{j=1}^{L} \frac{(S_j^p - S_j^q)^2}{S_j^p S_j^q}, \tag{8}$$

which performs better for our approach by experiments.

3. IMPLEMENTATION

Computing eigenspace of the Laplace-Beltrami operator on a manifold is well studied in the literature. It involves space discretization and a sparse eigensolver. In experiments, we use the finite element method [26] (also adopted in [17, 16]) to discretize manifolds. It can handle meshes with a wide range of simplicial degrees and topologies, including non-manifolds. We found a linear FEM enough for our approach. 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 [7, 3] 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) degenerate faces.

Once the stiff and mass matrices have been assembled, the rest is to solve a sparse symmetric generalized eigenvalue problem that has efficient solvers, such as IRAM [2] and Krylov-Schur [21].

4. EXPERIMENTAL RESULTS

4.1 Efficiency

Efficiency is usually important for shape retrieval techniques to be practical on large datasets. Very few works in 3D shape retrieval reports their timings. Most existing techniques in shape analysis and higher-level geometry processing are computationally expensive. Computational intensity is a major limitation for methods based on optimization, geodesic computation, 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 $O(n)$.

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 1. Such running times are based on our implementation of R-BiHDM and SD-GDM [20][2], and the original authors' implementation of meshSIFT [12][3].

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 1: 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.

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").

4.2 Signature Based Retrieval

We have tested the nonrigid shape retrieval performance of our method on a representative large dataset, which mixes the dataset from the nonrigid track of SHREC'11 with another dataset custom built by ourselves. The SHREC'11 nonrigid track dataset serves as a background dataset. It has 600 watertight triangle meshes that were derived from 30 original models. The custom built dataset contains 200 deformed meshes derived from 4 biped models, 3 dinosaur models and 3 quadruped models (Figure 3(a)) as well as 40 background meshes (other biped, dinosaur and quadruped models, see Figure 3 (b)). This mixed dataset was designed to be more challenging and practical than the SHREC'11 nonrigid track dataset 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. We have performed retrieval tests on the 200 deformed meshes in the mixed dataset with a total of 840 (200+40+600) meshes.

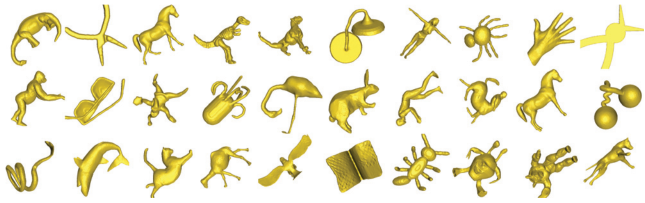

Figure 2: 30 classes of nonrigid models in the nonrigid track of SHREC'11

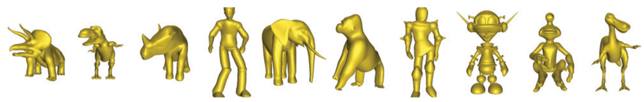

(a) Representative query meshes, including 4 bipeds, 3 dinosaurs and 3 quadrupeds

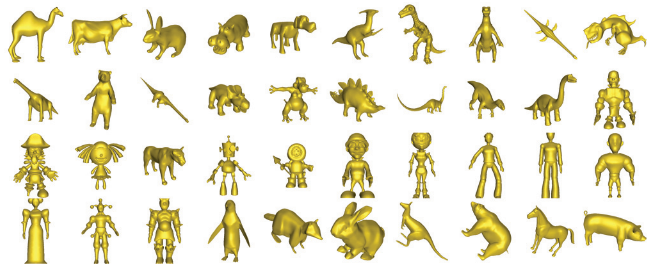

(b) Ambiguous meshes, including 13 bipeds, 13 dinosaurs and 14 quadrupeds

Figure 3: Examples from our custom built dataset

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 [19] for detailed definitions. In our method, we use Normalized Euclidean distance to measure similarity among R-BiHDM signatures (see eq. (7)).

We have compared the retrieval performance of our method with that of Shape-DNA (OrigM-n12-norm1) [17], meshSIFT [12], and SD-GDM [20]. These are the best performing non-hybrid methods in the nonrigid track of the SHREC'11 contest[4]. The parameters in these methods were set empirically to produce best performance. We have also compared our method, i.e. R-BiHDM, with pairwise

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

[3]This code can be downloaded at `https://mirc.uzleuven.be/MedicalImageComputing/downloads/meshSIFT.php`

[4]A hybrid technique 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.

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-n25	1.000	0.930	0.983	0.722	0.990
R-BiHDM-n30	1.000	0.970	0.996	0.739	0.997
R-BiHDM-n25	1.000	0.975	0.997	0.742	0.998
R-BiHDM-n23	**1.000**	**0.976**	**0.997**	**0.742**	**0.998**
R-BiHDM-n20	1.000	0.975	0.997	0.742	0.998
R-BiHDM-n15	1.000	0.970	0.994	0.739	0.997

Table 2: Retrieval performance evaluated using five standard measures on the mixed dataset.

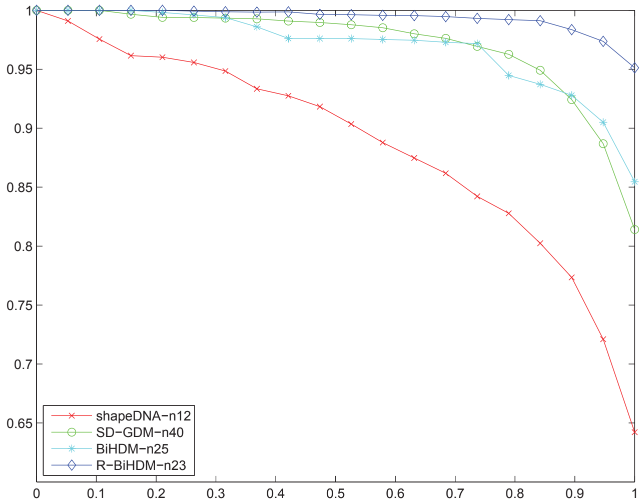

Figure 4: Precision-recall curves of R-BiHDM, SD-GDM and shapeDNA on the mixed dataset.

biharmonic distance matrix (BiHDM). Detailed comparison results on the mixed dataset are shown in Table 2 and Figure 4.

According to the statistics, our method achieves better performance than all of these methods. 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 primary goal of shapeGoogle [14] 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 is not comparable to the other methods we compared with. The results indicate that our method exhibits strong discriminative power on datasets with very similar shape instances.

5. CONCLUSIONS

In this paper, we have first given a brief introduction to existing methods for nonrigid 3D shape retrieval, and paid special attention to global structure coding. We have further proposed a novel method for building shape signatures. Our method uses biharmonic distance to construct a context-aware integral kernel operator on a manifold, then applies modal restriction to project this operator into a low-frequency representation, and finally computes its spectrum. Our method is isometry-invariant, discriminative, and numerically stable with respect to multiple types of perturbations. Our current implementation is based on FEM. We have evaluated our newly proposed method on representative datasets. Evaluation results indicate our method is surprisingly promising as a structural descriptor.

Our method has a few limitations that need further investigation. Although it has strength in identifying global non-rigid structures, 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.

Acknowledgments

The authors would like to thank anonymous reviewers for their comments and suggestions. This work was partially supported by Hong Kong Research Grants Council under General Research Funds (HKU718712).

6. REFERENCES

[1] A.Y. Alfakih. On the eigenvalues of euclidean distance matrices. *Computational and Applied Mathematics*, 27(3):237–250, 2008.

[2] W.E. Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Quart. Appl. Math*, 9(1):17–29, 1951.

[3] 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.

[4] 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.

[5] 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. *International Journal of Computer Vision*, 89(2):266–286, 2010.

[6] Michael Bronstein and Iasonas Kokkinos. Scale-invariance in local heat kernel descriptors without scale selection and normalization, 2009.

[7] P. Cignoni, M. Corsini, and G. Ranzuglia. Meshlab: an open-source 3d mesh processing system. *ERCIM News*, 73:45–46, 2008.

[8] 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 I–168. IEEE, 2001.

[9] Z. Lian, A. Godil, B. Bustos, M. Daoudi, J. Hermans, S. Kawamura, Y. Kurita, G. Lavoué, H. V. Nguyen, R. Ohbuchi, Y. Ohkita, Y. Ohishi, F. Porikli, M. Reuter, I. Sipiran, D. Smeets, P. Suetens, H. Tabia, and D. Vandermeulen. SHREC '11 Track: Shape Retrieval on Non-rigid 3D Watertight Meshes. pages 79–88.

[10] Y. Lipman, R. M. Rustamov, and T. A. Funkhouser. Biharmonic distance. *ACM Transactions on Graphics, ACM 0730-0301//10-ART*, 2007.

[11] D.G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.

[12] 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.

[13] 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.

[14] M. Ovsjanikov, A.M. Bronstein, M.M. Bronstein, and L.J. Guibas. Shape google: a computer vision approach to isometry invariant shape retrieval. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pages 320–327. IEEE, 2009.

[15] Gabriel Peyré and Laurent D. Cohen. Heuristically driven front propagation for fast geodesic extraction. *Intl. Journal for Computational Vision and Biomechanics 1(1)*, pages 55–67, 2009.

[16] Martin Reuter. Hierarchical shape segmentation and registration via topological features of laplace-beltrami eigenfunctions. *International Journal of Computer Vision*, 89(2):287–308, 2010.

[17] Martin Reuter, Franz-Erich Wolter, and Niklas Peinecke. Laplace-beltrami spectra as 'shape-DNA' of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.

[18] 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.

[19] Philip Shilane, Patrick Min, Michael M. Kazhdan, and Thomas A. Funkhouser. The princeton shape benchmark. In *Shape Modeling International*, pages 167–178. IEEE Computer Society, 2004.

[20] Dirk Smeets, Thomas Fabry, Jeroen Hermans, Dirk Vandermeulen, and Paul Suetens. Isometric deformation modelling for object recognition. *CAIP 2009, LNCS 5702*, pages 757–765, 2009.

[21] G. W. Stewart. A Krylov–Schur algorithm for large eigenproblems. *SIAM Journal on Matrix Analysis and Applications*, 23(3):601–614, 2001.

[22] Jian Sun, Maks Ovsjanikov, and Leonidas J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Comput. Graph. Forum*, 28(5):1383–1392, 2009.

[23] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. In *Shape Modeling International, 2003*, pages 130–139. IEEE, 2003.

[24] Vitaly Surazhsky, Tatiana Surazhsky, Danil Kirsanov, Steven J. Gortler, and Hugues Hoppe. Fast exact and approximate geodesics on meshes. *ACM Transactions on Graphics*, 24(3):553–560, July 2005.

[25] 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.

[26] OC Zienkiewicz, RL Taylor, and JZ Zhu. *The Finite Element Method–Its Basis and Fundamentals, volume 1*. Elsevier Butterworth-Heinemann, Amsterdam, London,, 2005.