

Temperature Distribution Descriptor for Robust 3D Shape Retrieval

Yi Fang
Purdue University
West Lafayette, IN, USA
fang4@purdue.edu

Mengtian Sun
Purdue University
West Lafayette, IN, USA
sun84@purdue.edu

Karthik Ramani
Purdue University
West Lafayette, IN, USA
ramani@purdue.edu

Abstract

Recent developments in acquisition techniques are resulting in a very rapid growth of the number of available three dimensional (3D) models across areas as diverse as engineering, medicine and biology. It is therefore of great interest to develop the efficient shape retrieval engines that, given a query object, return similar 3D objects. The performance of a shape retrieval engine is ultimately determined by the quality and characteristics of the shape descriptor used for shape representation. In this paper, we develop a novel shape descriptor, called temperature distribution (TD) descriptor, which is capable of exploring the intrinsic geometric features on the shape. It intuitively interprets the shape in an isometrically-invariant, shape-aware, noise and small topological changes insensitive way. TD descriptor is driven by heat kernel. The TD descriptor understands the shape by evaluating the surface temperature distribution evolution with time after applying unit heat at each vertex. The TD descriptor is represented in a concise form of a one dimensional (1D) histogram, and captures enough information to robustly handle the shape matching and retrieval process. Experimental results demonstrate the effectiveness of TD descriptor within applications of 3D shape matching and searching for the models at different poses and various noise levels.

1. Introduction

1.1. Background

With the advancement in data acquisition techniques, we have observed an exponential increase in the 3D meshed surface models across a variety of fields, such as engineering, entertainment and medical imaging [28, 24, 20, 10, 9, 3, 13, 22, 14, 27, 6, 1]. Although shape retrieval has been studied extensively, new capability of descriptors we have identified, such as awareness, intrinsic is less developed. To better understand the shape, a shape descriptor should have as many of the following properties as possible: 1)

shape-awareness: capture all possible shape information of the meshed surface as opposed to those in the local regions, 2) isometrically-invariant: consistent representation along isometric transformations of the meshed surface, 3) insensitivity to both topological and numerical noise such as consistent behavior even with topological short-circuits and numerical noises, 4) a compact representation for fast and effective shape matching.

1.2. Review of Related works

The shape descriptor for similarity measurements and matching between shapes has been extensively studied across areas as diverse as computer vision, robotics, structural biology and others [34, 17, 12, 36, 30]. We mainly review the related works that partly address the new emerging requirements for an robust shape descriptor. For a review of earlier shape descriptors, we refer the reader to [24] for D2 shape distribution, [12] for statistical moments of the model, [36, 29] for Fourier descriptor, [7] for Light Field Descriptor, [18] for Eigenvalue Descriptor (EVD) and [21] for harmonic-based methods for rotation invariant descriptors, [31, 33, 8] for methods of structure extraction based on graphs. Although these shape descriptors can represent the shape effectively, they are either sensitive to non-rigid transformation or topological changes, or computationally inefficient.

To be invariant to the isometric transformation, a widely used strategy is to represent the shape through the local geometric features. The shapes are matched based on the local descriptors, such as spin images, or the distribution of local geometric features, such as histograms [19, 16, 15, 23]. These shape descriptors are able to consistently describe deformable shapes, however they are sensitive to local geometric noise and are not able to provide a good picture of the overall shape. In addition, some of these descriptors require large amounts of space to store all of the local descriptors in order to facilitate a global shape matching. One of the representative works is the pose oblivious signature based on the local diameter of a 3D shape in the neighborhood of each vertex. However, because of the dependence




Figure 1. Flowchart of Computing TD descriptor.

of the local geometric features, the signature is sensitive to perturbation of local geometric features, for example small bumps on the surface and numerical noise.

The latest efforts to circumvent the problem of a robust shape signature are diffusion based approaches [32, 5, 28, 25]. The concept is first presented by recently proposed shape signature, global point signature (GPS). It shows that the eigenfunctions of the Laplace-Beltrami operator are able to robustly characterize the point on meshed surface. Through GPS, each vertex is represented by a high-dimensional vector with its components of scaled eigenfunctions of the Laplace-Beltrami operator evaluated at the vertex. GPS gains its advantage by discarding the use of geodesic distances or any of its variation but the eigenvalues and eigenfunctions of the Laplace-Beltrami. Another concise and provable signature is heat kernel signature (HKS). The HKS has gained attention because of the concise form and other properties. HKS is invariant to isometric deformations, insensitive to the small perturbations of the surface, and especially a multi-scale interpretation of the shape. Both GPS and HKS are point based signatures, which characterizes each vertex on the meshed surface using a vector. However, the authors in [32, 28] did not demonstrate how to retrieve shapes using HKS and GPS, although they allude to the future potentials in shape retrieval applications. Shape google proposed in [25] develops a shape retrieval strategy based on HKS.

1.3. Temperature Distribution Descriptor

We introduce TD descriptor based on the evaluation of temperature distribution after applying a unit heat at each vertex. Every vertex on the meshed surface is assigned a value that measures the average temperature of all of the vertices on the meshed surface after applying a unit heat on that vertex. TD descriptor is built in the form of a one dimensional histogram from the distribution of the average temperature values. TD naturally inherits the properties of the heat kernel. Figure 1 illustrates the pipeline of computing the TD descriptor for a 3D horse model. There are three main steps included in the flowchart. 1) Obtaining point temperature distribution for every vertex on the meshed surface: a unit heat is applied on one vertex of the 3D meshed surface and the temperature distribution is recorded. This

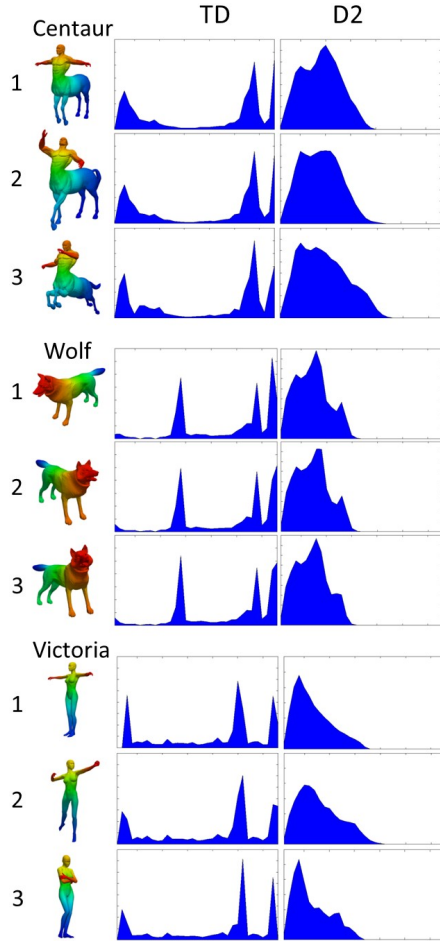


Figure 2. Comparison of TD and D2 descriptors for shapes with pose changes.

is repeated for every vertex of the mesh; 2) Computing the average temperature distribution: the value of the average temperature is first computed for each vertex and all of these values together produce the average temperature distribution for the whole model; 3) Building temperature distribution (TD) descriptor: build a one dimensional histogram based on the average temperature distribution.

In addition, Figure 2 displays comparisons of TD descriptor for deformable 3D models across different categories with different poses. This figure is an overall presentation of our shape descriptor and its performance on deformed models. TD descriptor describes each 3D model in a concise form of 1D histogram as shown in the figure. Within the same category, we can find that the TD descriptors of the deformable models are consistent with pose changes. Across different categories, TD descriptor differentiates the models from each category. For comparison, the shape distribution (D2) is displayed on the right of the TD descriptor. We can find that D2 has difficulty in maintaining

consistency with pose changes. The detailed analysis of the Figure 2 will be presented in the experiment section. The contribution of this work is centered on the development of a novel intrinsic shape descriptor and the corresponding shape retrieval scheme for efficient and effective 3D shape retrieval.

2. Methods

2.1. Heat Kernel

The heat kernel $H_t(x, y)$ measures the amount of the heat that passes from the vertex x to any other vertex y within a certain amount of time. Since heat kernel aggregates heat flow through all possible flowing paths between two vertices on the meshed surface, it can capture much of the structure encapsulated by the meshed surface. If a vertex has more paths to other vertices, the heat flows faster across the manifold. Consequently, the average temperature is higher when the heat is applied at vertices with a large number of paths to other vertices. The 3D model is represented as a graph $G = (V, E, W)$, where V is the set of vertices, E is the set of edges and W is the weigh value for each edge. The widely used cotangent weight for the edges is adopted in our work [28]. The heat flow across the surface is governed by the heat equation $u(x, t)$, where x denotes one vertex on the meshed surface and t denotes the time after the application of unit heat. The heat kernel provides the fundamental solution of the heat equation, which is closely associated with Laplace-Beltrami operator by:

$$\frac{\partial H_t}{\partial t} = -LH_t \quad (1)$$

where H_t denotes the heat kernel, L denotes the Laplace-Beltrami and t denotes time. The heat kernel can be further expressed as the eigenfunction expansion by the Laplace-Beltrami described below.

$$H_t = \exp(-tL) \quad (2)$$

By the Spectral Theorem, the heat kernel can be further expressed as follows:

$$H_t(x, y) = \sum_i e^{-\lambda_i t} \phi_i(x) \phi_i(y) \quad (3)$$

where λ_i is the i_{th} eigenvalue of the Laplacian, ϕ_i is the i_{th} eigenfunction, and x and y denotes vertex x and y respectively.

2.2. Proposed shape descriptor

2.2.1 Average temperature function

The average temperature at a vertex on the surface is defined according to the point temperature distribution after the application of a unit heat at this vertex. The average

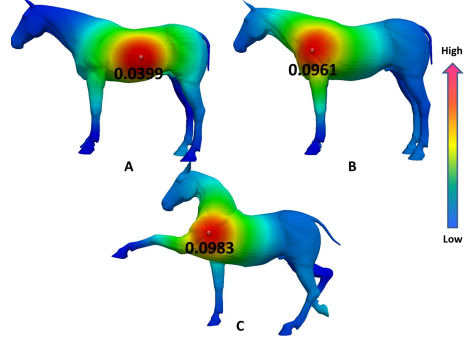


Figure 3. Illustration of point temperature distribution and average temperature.

temperature function is derived from heat kernel and defined in Eq. 3 for each vertex. Figure 3 illustrates the point temperature distribution after applying a unit heat on a vertex (on the body of the horse and marked by a red ball). Figure 3(A) and (B) display the same horse model, and (C) display a horse model with pose change. The color denotes its the temperature value which is mapped with a colorbar on the right. For a vertex on the shape, the closer the color is to red, the higher temperature value is, while the closer the color is to the blue, the lower the temperature value is. The heat is applied on the different vertices on Figure 3(A) and (B), and corresponding vertices of horse models in Figure 3(B) and (C). In (A) and (B), we can see that the point temperature distribution varies because of applications of heat at different places. The point temperature distribution is consistent in (B) and (C) even with the pose change. The normalized average temperatures, which are listed under the red balls, are numerically distinct for (A) and (B) and consistent for (B) and (C). Note that, the average temperature is normalized to [0 1]. It is easy to find that point temperature distribution is invariant to pose changes, and the average temperature for the corresponding vertices are therefore highly consistent with pose changes. We define the Average temperature function as follows:

$$\begin{aligned} AVG_t(x) &= \frac{1}{N-1} \sum_{y, y \neq x} H_t(x, y) \\ &= \frac{1}{N-1} \sum_{y, y \neq x} \sum_i e^{-\lambda_i t} \phi_i(x) \phi_i(y) \end{aligned} \quad (4)$$

where $AVG(x)$ denotes the average temperature for vertex x , and (x, y) denote the vertex x and y on the meshed surface respectively, t denotes the heat dissipating time and N is the total number of vertices. The parameter t is often chosen empirically [11]. In our work, we choose a relatively smaller t to preserve a higher resolution version of the original surface [35].

The definition of average temperature (AVG) function

is invariant to isometric transformation and robust against topological changes and local numerical noise. In addition, the average temperature function naturally evaluates the influence of a vertex, which means, how important the vertex is among all of the vertices. This is because the higher the average temperature, the faster the heat flows across the meshed surface, therefore, the more influential the vertex is. The influence of a vertex is critical for the definition of a shape descriptor, which is normally used to weight the histogram. For example, the authors [15] define the influence as the area of the triangles surrounding a vertex divided by the whole surface area, and then use the influence to weight the shape descriptor.

2.2.2 Temperature Distribution Descriptor

The average temperature function is able to robustly characterize the vertices on the meshed surface. However, each vertex from the model is described by an average temperature value, which requires large amount of storage space for a large model (over 20K vertices) and large database. In addition, the retrieval in 3D database would be computationally expensive and inefficient. To effectively retrieve the 3D model from a large database, we need to develop shape descriptors that can convert the set of average temperature values to a compact and concise representation of the entire 3D model. That is, as we will demonstrate further, the whole 3D model can indeed be described by a succinct representation. A solution for this conversion is obtained by approximating the probability density function of the values. For example, shape distribution is built based on the distribution of the distance values between two random points on the surface (D2) or a point to the centroid of the model (D1) [24]. The pose-oblivious shape signature is built based on the distribution of the local diameters on the surface [15]. We adopt the similar strategy to develop our shape descriptor. TD descriptor is built based on the distribution of the values of average temperature for all of the vertices on the meshed surface. As average temperature naturally integrates the influence of each vertex, there is no need to weight the value of the average temperature based on the influence like weight scheme presented in [15].

2.3. Relationship among existing approaches

We briefly review the relationship among the existing approaches which are developed based on the eigensystem of Laplace-Beltrami operator. The given heat kernel has the following eigen-decomposition:

$$H_t(x, y) = \sum_{i=0} e^{-\lambda_i t} \phi_i(x) \phi_i(y) \quad (5)$$

where ϕ_i and λ_i are the i_{th} eigenfunction and eigenvalue of Laplace-Beltrami operator, respectively.

Global point signature The global point signature (GPS) is defined as:

$$GPS(x) = \left(\frac{\phi_1(x)}{\sqrt{\lambda_1}}, \frac{\phi_2(x)}{\sqrt{\lambda_2}}, \dots, \frac{\phi_i(x)}{\sqrt{\lambda_i}}, \dots \right) \quad (6)$$

GPS is a global point signature and invariant to isometric deformation of the shape. The global representation of the shape is useful for shape matching and global intrinsic symmetry detection [26, 28].

Heat kernel signature Heat kernel signature is defined as:

$$HKS_t(x) = H_t(x, x) \quad (7)$$

$$= \sum_{i=0} e^{-\lambda_i t} \phi_i(x) \phi_i(x)$$

HKS is a point signature and also captures the intrinsic geometry of the shape. In addition, it is able to represent the shape in a multi-scale way, which is able to detect partial symmetries or to perform partial matching [11, 32].

Diffusion distance The diffusion distance is defined below:

$$D_t^2(x, y) = H_t(x, x) + H_t(y, y) - 2H_t(x, y) \quad (8)$$

$$= \sum_i e^{-\lambda_i t} (\phi_i(x) - \phi_i(y))^2$$

Diffusion distance is a distance metric between a pair of vertices on a meshed surface. It is also a multi-scale representation metric and has been used in [10] for the hierarchical segmentation of the articulated shapes.

Note that all of GPS, HKS, Diffusion distance and TD are able to intrinsically describe the 3D shape in the manner of shape-awareness, isometrically-invariant and insensitivity to both topological and numerical noise. The distinct feature of the TD descriptor over GPS and HKS is that TD is a concise and succinct form of the representation for a whole shape while GPS and HKS are point based signature for describing the individual vertex on the shape.

3. Experimental Results

We carried out two experiments for shape matching and retrieval and assessed TD's performance from the experimental results. The models used in the experiments were chosen from the following database: TOSCA 3D database [9, 3, 2, 4]. We define a universal time unit $\tau = 0.01$ for the computation of TD descriptor. In all of the following tests, 80τ is used to compute the TD descriptor of models (t in Eq. 4).

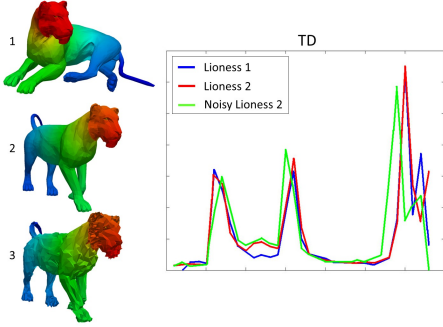


Figure 4. TD descriptor for shape matching for deformed and noisy models.

3.1. Shape Matching Performance

The shape matching is the first step of database retrieval of 3D model. A good shape descriptor should be able to consistently represent the 3D model with pose changes, topological changes, and noise corruption. We assess the performance of TD descriptor from these three aspects.

Consistency over articulated bodies In this test, we test the performance of TD descriptor on the deformed shape models. We choose three ‘Centaur’ models, three ‘Victoria’ models and three ‘Wolf’ models with different poses. We compare the TD descriptor with a classical descriptor, the shape distribution (D2). The comparison results are shown in Figure 2. We can see from the figure that TD descriptor is able to consistently represent the shapes with pose changes. In addition, as indicated in the figure, ‘Victoria’ model 3 has both hands crossed and touching the chest. This test on three ‘Victoria’ models demonstrate another attractive property for TD descriptor of being insensitive to topological noise (short-circuit touching) besides of pose oblivious property. Furthermore, an interesting observation from the comparison is that TD descriptor includes much richer information than D2. Compared to D2, TD has more features as shown in Figure 2 that there are about three major peaks in the TD descriptor while D2 is mostly centralized in the one major peak. This makes TD descriptor not only consistent in representing the shapes in the same category but also distinctly describes the shapes across different categories. It is important for a shape descriptor to have both capabilities for effective database retrieval. The test demonstrates that TD descriptor performs well on consistent representation on the articulated shapes.

Resistance to noise In this test, we will demonstrate that TD descriptor is robust to numerical noise. We experiment by perturbing the vertices of the original model by various levels of numerical noise. The numerical noise, a 3-dimensional vector, is randomly generated from a mul-

tivariate normal distribution, which can be written in the following notation.

$$Noise \sim N_3(\mu, NR * \Sigma) \quad (9)$$

where $\mu = [E[X_1], E[X_2], \dots, E[X_k]]$ is the 3-dimensional mean vector (the coordinates of all vertices) and $\Sigma = [Cov[X_i, X_j]]_{i=1,2,\dots,k;j=1,2,\dots,k}$ (the 3×3 covariance matrix for all vertices), and NR denotes the ratio between the variance of noise and variance of original signal (coordinates of the vertices).

In Figure 4, Lioness 2 is the original model and Lioness 3 is a noise corrupted version of Lioness 2. We are interested in the consistency of TD descriptor between the clean and noisy models. In the noisy model Lioness 3 (NR 0.001), geometric features of the original model have moderately deteriorated. As indicated by the results in Figure 4, at moderate level of noise, TD descriptors for the clean and noisy models (plotted with red and green curve respectively) align reasonably well. If the descriptors heavily rely on the local geometric features, they would be overwhelmed by the geometric corruption due to the noise. The test demonstrates that TD descriptor is resistant to noise.

3.2. 3D Shape Retrieval Performance

To further investigate the effectiveness, we test the TD descriptor on a benchmark dataset of 3D models extracted from TOSCA 3D database. We design two tests for the verification. In the first test, we compute the TD descriptor on the clean 3D models. In the second test, we compute the TD descriptor on the 3D models whose vertices are perturbed with different levels of numerical noise. Before searching, we pre-calculate shape descriptors of all queries in the database and stored them as index files. Therefore, every shape is transformed to a compact 1D vector (e.g. 128 dimension). With this new form of representation, shape retrieval becomes extremely fast since only the distances among 1D vectors need to be compared. We use L_2 norm to compute the distance between two TD descriptors.

We compare TD with two competitive shape descriptors in the retrieval test: Eigenvalue descriptor (EVD) which is a state of the art spectral technique for articulated shape retrieval, and a classical rigid shape descriptor, D2. The TOSCA database is a collection of 3D nonrigid shapes in a variety of poses. The 3D shapes are divided into different groups. The models from the same group are basically the same type of objects but with different poses (see Figure 5 for examples of models in different groups). To prepare the noisy database, we corrupt the original model at two different noise levels (NR=0.001 and NR=0.008). Figure 7 illustrates an example of the original model (sitting dog) and its noisy versions (NR 0.001 and NR 0.008).




Figure 5. Retrieval result by TD descriptor.

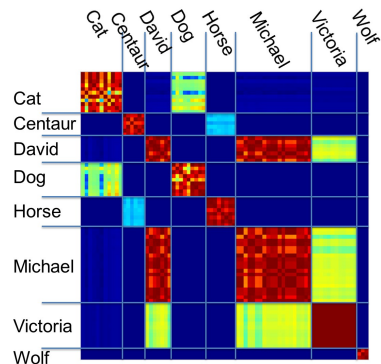


Figure 6. Illustration of distance matrix.

3.2.1 3D retrieval on clean model

In the experiment, we first compute the distance for every pair of the models in the dataset, and then rank the distance value for each query model. The precision and recall for every query model is recorded then averaged to generate the final precision and recall curve. We compare the retrieval performance between TD, EVD and D2 descriptors. The retrieval results are plotted in Figure 7. The red plots in Figure (A), (B) and (C) are the precision-recall curves for TD, EVD and D2 for clean models, respectively. A perfect search retrieves all relevant objects consistently at each recall level, generating a horizontal line at precision equal to 1.0. However, practically, the precision decreases with the increasing recall. The closer the curve tends to the horizontal line at precision equal to 1.0, the better the re-

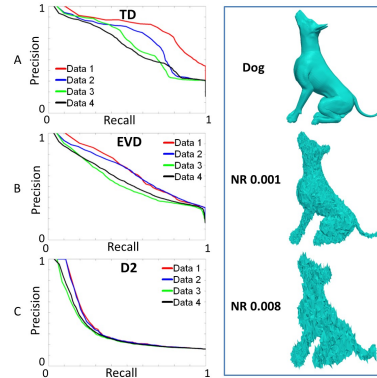


Figure 7. Shape retrieval result for both clean and noisy models.

trieval method. Therefore, based on the comparison of the three curves, TD descriptor performs better than both EVD and D2 shape descriptor since the precision-recall curves for EVD and D2 methods drops faster.

To further investigate the retrieval results, we select one query model from each group and show the corresponding retrieval result (see Figure 5). We only list the top five retrieval results with incorrect retrieval boxed in red. All models are colored according to the values of the average temperature. It can be observed that the color distribution remains consistent within each group. TD descriptor fails to differentiate part of the models from groups ‘David’ and ‘Michael’ (represented by query models in the third and sixth row, respectively), which results in the incorrect retrievals in the third row and the sixth row. However, we find that shapes from these two groups are really similar and thus are difficult to be distinguished from each other even by human observation. In addition, because there are only three models in group ‘Wolf’, two models from group ‘Cat’ are returned in the top five retrieval list by our method. This is interesting as we can visually see that the shape of wolf and cat are similar to each other in this database.

Furthermore, results in Figure 5 can be confirmed by the pairwise distance matrix for all models (see Figure 6). The distance is mapped by color. The more similar two shapes are, the smaller the distance is, and the redder the color is. It is clearly observed that diagonal areas are closer to red because distances among the models from the same groups tend to be smaller. The red off-diagonal areas are interesting as they indicate the overlapping of the models across different groups. The red overlapping area for groups ‘David’ and ‘Michael’ explains the retrieval results in Figure 5, as it indicates the shapes are geometrically similar to each other in these two groups.

3.2.2 3D retrieval on noisy model

To further demonstrate the performance of the TD descriptor, we design three retrieval tests on noisy models and com-

pare the results to EVD and D2 descriptor results. In the first test, test data consists of the original models and the noisy models with NR equal to 0.001. In the second test, test data consists of the original models and the noisy models with NR equal to 0.008. In the third test, we mix all of the three sets of models above. A sitting dog in Figure 7 is used to illustrate the noise effect. Precision-recall curves are produced using the same setting in the tests above. All precision-recall curves are plotted in Figure 7 with different colors. In Figure 7(A), we find that the retrieval performance of the TD descriptor is slightly affected by moderate noise corruption. The performance of TD descriptor remains reasonable even when NR increases to 0.008, which is in fact worse than the noise levels in real scenario. The performance of TD is compared to EVD and D2 descriptors. It is not surprising that D2 fails to have a reasonable retrieval performance because it is a rigid shape descriptor and sensitive both isometric transformation and numerical noise. Although EVD gives a reasonable retrieval performance, it performs worse than TD at different levels of the noises.

4. Discussion and Conclusions

The data-richness of 3D models urges us to focus on robust and intelligent model analysis and understanding. We develop a previously undescribed shape descriptor that assists shape matching and retrieval. TD descriptor demonstrates its performance in various tests for matching and retrieving 3D shapes with deformation, topological short-circuits and numerical noises. The tolerance to noise is highlighted by the experiments in the paper. One of the critical properties for TD descriptor that has not been fully discussed in the experiment is the heat dissipating time (t in Eq.4) which is crucial to determine how locally or globally the geometric structure is exploited. In our current implementation of TD descriptor, a relatively smaller scale is used to preserve a higher resolution version of original shape. In the future work, We will carry out the sensitivity analysis to find how the TD descriptor correlates to the time scale and then develop the algorithms to determine which time is the optimal time to describe the shape. On the other hand, we will explore the multi-scale properties of TD descriptor for shape matching by varying the heat dissipating time across the meshed surface. In addition, we will test on more datasets (SHREC) to evaluate the performance of our descriptor and add more comparisons to other methods for example, ShapeGoogle [25].

Acknowledgement

We would like to thank the Institute for Pure & Applied Mathematics (IPAM), a National Science Foundation (NSF) Math Institute at the University of California at Los Ange-

les (UCLA) for the fellowship to support Karthik Ramani. The experimental images are provided by the whole brain database from Harvard university. This work is partly based upon work supported by the National Science Foundation Partnership for Innovation Grant# 0917959 (3D Hub), NSF IIS Grant# 0535256, Donald W. Feddersen Chair Professorship support, Purdue University support for Faculty Study in a second discipline, and the School of Mechanical Engineering. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] A. Bronstein, M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov, G. Patane, M. Spagnuolo, and R. Toldo. Shrec 2010: robust large-scale shape retrieval benchmark. *Eurographics workshop on 3D object retrieval 2010*, 2010. 9
- [2] A. Bronstein, M. Bronstein, and R. Kimmel. Calculus of nonrigid surfaces for geometry and texture manipulation. *IEEE Transactions on Visualization and Computer Graphics*, 13:902–913, 2007. 12
- [3] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Sci. Comput.*, 28:1812–1836, September 2006. 9, 12
- [4] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Numerical geometry of non-rigid shapes. *Springer*, 2008. 12
- [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. *Int. J. Comput. Vision*, 89:266–286, 2010. 10
- [6] M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. *IEEE Computer vision and pattern recognition (CVPR) 2010*, pages 1704–1711, 2010. 9
- [7] D.-Y. Chen, X.-P. Tian, Y. te Shen, and M. Ouhyoung. On visual similarity based 3d model retrieval. *Computer Graphics Forum*, 22:223–232, 2003. 9
- [8] D.-Y. Chen, D. Yun Chen, and M. Ouhyoung. A 3d object retrieval system based on multi-resolution reeb graph. *Proc. Computer Graphics Workshop*, pages 16–20, 2002. 9
- [9] X. Chen, A. Golovinskiy, and T. Funkhouser. A benchmark for 3D mesh segmentation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 2009. 9, 12

- [10] F. De Goes, S. Goldenstein, and L. Velho. A hierarchical segmentation of articulated bodies. *Computer Graphics Forum*, 27:1349–1356, 2008. 9, 12
- [11] T. Dey, K. Li, C. Luo, P. Ranjan, I. Safa, and Y. Wang. Persistent heat signature for pose-oblivious matching of incomplete models. *Computer Graphics Forum*, 29:1545–1554, 2010. 11, 12
- [12] M. Elad, A. Tal, and S. Ar. Content based retrieval of vrmf objects - an iterative and interactive approach. *Proc. Sixth Eurographics Workshop Multimedia*, pages 97–108, 2001. 9
- [13] Y. Fang, Y.-S. Liu, and K. Ramani. Three dimensional shape comparison of flexible proteins using the local-diameter descriptor. *BMC Structural Biology*, 9(1), 2009. 9
- [14] Y. Fang, M. Sun, M. Kim, and K. Ramani. Heat-mapping: A robust approach toward perceptually consistent mesh segmentation. *IEEE computer and vision and pattern recognition (CVPR) 2011*, pages 2145–2152, 2011. 9
- [15] R. Gal, A. Shamir, and D. Cohen-Or. Pose-oblivious shape signature. *IEEE Transactions on Visualization and Computer Graphics*, 13:261–271, 2007. 9, 12
- [16] D. Huber, A. Kapuria, R. Donamukkala, and M. Hebert. Parts-based 3d object classification. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 2:82–89, 2004. 9
- [17] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman, and K. Ramani. Three-dimensional shape searching: state-of-the-art review and future trends. *Computer-Aided Design*, 37(5):509 – 530, 2005. Geometric Modeling and Processing 2004. 9
- [18] V. Jain and H. Zhang. A spectral approach to shape-based retrieval of articulated 3d models. *Computer-Aided Design*, 39(5):398 – 407, 2007. Geometric Modeling and Processing 2006. 9
- [19] A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):433–449, 1999. 9
- [20] S. Katz, G. Leifman, and A. Tal. Mesh segmentation using feature point and core extraction. *The Visual Computer*, 21:649–658, 2005. 9
- [21] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164, 2003. 9
- [22] Y.-S. Liu, Y. Fang, and K. Ramani. Idss: deformation invariant signatures for molecular shape comparison. *BMC Bioinformatics*, 10, 2009. 9
- [23] R. Ohbuchi, K. Osada, T. Furuya, and T. Banno. Salient local visual features for shape-based 3d model retrieval. *Shape Modeling and Applications, 2008. SMI 2008. IEEE International Conference on*, pages 93–102, 2008. 9
- [24] R. Osada, T. Funkhouser, B. Chazelle, and D. Dokin. Shape distributions. *ACM Transactions on Graphics*, 33:133–154, 2002. 9, 12
- [25] M. Ovsjanikov, A. Bronstein, and M. Bronstein. Shape google: a computer vision approach to invariant shape retrieval. *Proc. NORDIA*, 2009. 10, 15
- [26] M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. *Proceedings of the Symposium on Geometry Processing*, pages 1341–1348, 2008. 12
- [27] G. Patane and B. Falcidieno. Multi-scale feature spaces for shape processing. *Shape Modeling International*, pages 113–123, 2010. 9
- [28] R. M. Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. *Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 225–233, 2007. 9, 10, 11, 12
- [29] D. Saupe and D. V. Vranic. 3d model retrieval with spherical harmonics and moments. *DAGM*, pages 392–397, 2001. 9
- [30] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The princeton shape benchmark. *In Shape Modeling International*, pages 167–178, 2004. 9
- [31] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, pages 13–32, 1998. 9
- [32] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *SGP '09: Proceedings of the Symposium on Geometry Processing*, pages 1383–1392, 2009. 10, 12
- [33] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. *Proceedings of the Shape Modeling International 2003*, page 290, 2003. 9
- [34] J. W. H. Tangelder and R. C. Veltkamp. A survey of content based 3d shape retrieval methods. *In Shape Modeling International*, pages 145–156, 2004. 9
- [35] A. Vaxman, M. Ben-Chen, and C. Gotsman. A multi-resolution approach to heat kernels on discrete surfaces. *ACM Trans. Graph.*, 29:121:1–121:10, July 2010. 11
- [36] D. V. Vranic, D. Saupe, and J. Richter. Tools for 3d-object retrieval: Karhunen-loeve transform and spherical harmonics. *IEEE MMSP 2001*, pages 293–298, 2001. 9