

A 3D shape descriptor based on spectral analysis of medial axis [☆]



Shuiqing He ^{a,*}, Yi-King Choi ^a, Yanwen Guo ^b, Xiaohu Guo ^c, Wenping Wang ^a

^a Department of Computer Science, The University of Hong Kong, Hong Kong

^b National Key Lab for Novel Software Technology, Nanjing University, Nanjing, China

^c Department of Computer Science, The University of Texas at Dallas, Richardson, TX, USA

ARTICLE INFO

Article history:

Received 7 July 2015

Accepted 30 August 2015

Available online 9 October 2015

Keywords:

Spectral analysis

Laplace–Beltrami operator

Medial axis

Shape descriptor

Shape matching

Shape retrieval

ABSTRACT

The medial axis of a 3D shape is widely known for its ability as a compact and complete shape representation. However, there is still lack of a generative description defined over the medial axis directly which limits its actual application to 3D shape analysis such as shape matching and retrieval. In this paper, we propose a new spectral shape descriptor that directly applies spectral analysis to the medial axis of a 3D shape, which we call the *medial axis spectrum* for a 3D shape. We develop a newly defined Minkowski–Euclidean ratio inspired by the Minkowski inner product to characterize the geometry of the medial axis of a 3D mesh. We then generalize the Laplace–Beltrami operator to the medial axis, and take the solution to a Laplacian eigenvalue problem defined on it as the medial axis spectrum. The medial axis spectrum is invariant under rigid transformation and isometry of the medial axis, and is robust to shape boundary noise as shown by our experiments. The medial axis spectrum is finally used for 3D shape retrieval, and its superiority over previous work is shown by extensive comparisons.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

The medial axis of a shape comprises the set of interior points, each of which has two or more closest points on the shape boundary (Siddiqi and Pizer, 2008). Each point on the medial axis can also be associated with the local thickness, symmetry information and part-structure of the object, an advantage not possessed by boundary shape representation. This makes the medial axis based shape representation compact and complete in the sense that the original shape can be reconstructed from its medial axis. For this reason, the medial axis shape representation is widely used for shape editing and processing.

However, there is still lack of a generative description defined over the medial axis directly. This limits the actual application of the medial axis to 3D shape analysis. As a matter of fact, in shape indexing, matching and retrieval, the medial axis has not been fully exploited directly but is only used for derivation of a high-level graph abstraction to which graph matching or graph spectral techniques are applied (Shokoufandeh et al., 2005; Siddiqi et al., 1999, 2008).

[☆] This paper has been recommended for acceptance by Rida Farouki.

* Corresponding author. Tel.: +852 95388869.

E-mail addresses: qingshui5889@gmail.com (S. He), ykchoi@cs.hku.hk (Y.-K. Choi), ywguo.nju@gmail.com (Y. Guo), xiaohuguo@gmail.com (X. Guo), wenping@cs.hku.hk (W. Wang).




Fig. 1. Medial axes of a 2D bird (left) and a 3D eagle (right).

In this paper, we propose a new medial axis based spectral shape descriptor which we call the medial axis spectrum for a 3D shape. The eigenfunctions of the Laplace–Beltrami operator described by Levy (2006) have the capacity of understanding the global surface property of a shape. This can be used for obtaining an isometry-invariant descriptor known as the Shape-DNA (Reuter et al., 2006), based on the boundary shape representation. We show that the spectral analysis approach can be directly applied to the medial axes of 3D shapes. This yields a generative medial axis shape descriptor which could be used for shape analysis tasks including shape matching and retrieval.

Applying spectral analysis to the medial axis of a 2D shape has been explored by the previous work (He et al., 2014), which makes use of the derivative of medial axis radius function that describes the local geometry of the medial axis. It is, however, not intuitive to generalize it to a 3D shape. Fig. 1 shows a comparison between the medial axes of a 2D bird and a 3D eagle. The medial axis of a 2D shape is a 1D curve even though medial curves may connect at branch points. In contrast, the medial axis of a 3D shape is a surface in 3D. Although the dimension is just one degree higher, parametrization of the medial surface is not readily available. This makes it challenging to find a derivative similar metric on the radius function of a 3D shape for describing the local geometry of the medial surface, which needs to be fed to our derived Laplacian eigenfunctions. To this end, we develop a newly defined Minkowski–Euclidean ratio inspired by the Minkowski inner product to characterize the geometry of the medial axis surface of a 3D mesh. We then generalize the Laplace–Beltrami operator to the medial axis, and take the solution to a Laplacian eigenvalue problem defined on the surface as the medial axis spectrum.

The medial axis spectrum we propose has several favorable properties. On the one hand, it is invariant to rigid transformation of the input shape as well as isometry of the medial axis. On the other hand, the medial spectrum is robust to shape boundary noise and remains stable upon small perturbation of the medial axis. The medial axis itself is known for its instability. That is, small variations of object boundary may lead to a big fluctuation to the medial axis, a phenomenon that hinders severely the effective utilization of medial axes. We measure quantitatively the bound of the medial axis spectrum for the perturbed medial axis. Our experiments also show that the spectrum is resistant to the perturbation.

We will first review the related work in Section 2. We then present the formulation of the medial axis spectrum in Section 3 and the computation details with robustness analysis in Section 4. Experimental results and an application of the medial axis spectrum to shape retrieval will be given in Sections 5 and 6. We finally conclude the paper in Section 7.

2. Related work

2.1. Spectral shape analysis

Spectral shape analysis aims to capture the intrinsic properties of a shape by studying the spectrum of the Laplace–Beltrami operator of a shape. Invariant descriptors are thus derived. The pioneer work by Reuter and his colleagues (Reuter et al., 2006) proposed to use a subset of eigenvalues of the Laplace–Beltrami operator to serve as a shape descriptor, namely the Shape-DNA. The sequence of the first k eigenvalues is also called the k -spectrum. The spectrum is invariant to translation of the shape and is isometric invariant. Additionally the spectrum can be normalized (Reuter et al., 2006) so that scaling factors for the geometric shape can be obtained easily (Reuter et al., 2005). It has been known that the k -spectrum is effective in distinguishing different shapes in practice, though in theory there exist different shapes (known as the isospectral shapes) yielding the same spectrum (Reuter et al., 2006). We follow the same spirit here: while the Shape-DNA method considers the Laplace–Beltrami operator on a shape boundary, we extend the operator to the medial axis of a 3D shape.

Rustamov (2007) presented the global point signature (GPS) embedding by combining the Laplace–Beltrami eigenvalues and eigenfunctions. This signature is robust to local topology changes. Both Shape-DNA and GPS are global shape descriptors. Sun et al. (2009) introduced the heat kernel signature, a multi-scale local shape descriptor which is insensitive to topological noise at small scales and can be used for partial shape matching. Due to the Laplace–Beltrami formulation, these shape descriptors are invariant under isometric deformation and therefore allow shape bending without distorting the local metric on a shape, which is commonly observed in natural articulated objects.




Fig. 2. Instability of the 3D medial axis. Any small perturbation of the shape boundary causes drastic changes to the medial axis. Left: a shape (upper-left) with smooth boundary and its medial axis (bottom-left). Right: the same shape (upper-right) with slightly perturbed boundary and its medial axis (bottom-right).

Jain and Zhang presented an approach for 3D shape retrieval (Jain and Zhang, 2007). Each shape is represented by the eigenvalues of an appropriately defined matrix, which encodes pairwise relations between the data points. Retrieval is then performed in the spectral domain using shape descriptors based on the dissimilarity distance.

2.2. Medial axis for shape analysis

Blum (1967) proposed the medial axis as an intrinsic shape representation. It consists of the set of points, also known as the medial loci, in the interior of the shape each of which has two or more closest points on shape boundary. The medial axis together with the radius function of the medial balls is called the medial axis transform. The original shape can be recovered from its medial axis based shape representation. However, the medial axis is prone to boundary noise and is known for its instability (Siddiqi and Pizer, 2008). That is, any small perturbation of the shape boundary may cause drastic changes to the medial axis (see Fig. 2), leading to a lot of unstable spikes and making it highly non-manifold. The excessive geometric complexity and pathological topology of such an unstable medial axis severely hinder its effective application to shape analysis, even though researchers (Miklos et al., 2010; Sud et al., 2005) have made great efforts to generate simplified and clean medial axes.

A series of methods (Siddiqi et al., 1999, 2008) are developed to perform shape matching based on the medial axis of a 3D shape. Graph Laplacian is applied to a directed acyclic graph, derived from the medial axis and encoding its topological structure, for obtaining an eigenvalue characterization of the graph. This is finally used for shape matching. Note that in these methods, the discrete Laplacian is involved. The discrete Laplacian strives to retain the properties of the continuous Laplace–Beltrami operator. Our method works quite differently from these medial graph spectra approaches. Though working also on medial axis, our method is based on the continuous Laplace–Beltrami operator. Our medial axis spectrum thus generated, as a result, better captures the intrinsic geometric and topological properties of the domain.

3. Medial axis spectrum of a 3D shape

We shall use a triangular mesh to represent the medial axis surface of a 3D shape in this paper. We first review some related concepts about the medial axis spectrum. Then we show how to define the medial axis spectrum as the 3D shape descriptor.

Let f be a real-valued function defined on a Riemannian manifold. The Laplace–Beltrami operator Δ is defined as

$$\Delta f = \text{div}(\text{grad } f),$$

where div is the divergence and $\text{grad } f$ is the gradient of f on the manifold. The Laplace–Beltrami operator is a linear operator that maps functions to functions on the manifold. We can therefore study its eigenvalues and eigenfunctions as denoted by the Helmholtz equation (Auteri and Quartapelle, 2000; Bao et al., 2004)

$$\Delta f = -\lambda f,$$

which is also known as the Laplacian eigenvalue problem. Here, a non-trivial solution of $\lambda \in \mathbb{R}$ is an eigenvalue of Δ and f is its corresponding eigenfunction. The eigenvalues of Δ are non-negative and can be arranged in non-decreasing order (repeated by multiplicities)

$$0 \leq \lambda_0 \leq \lambda_1 \leq \dots,$$

which constitute the *spectrum* of the manifold.

For a 2D shape, the medial axis spectrum is defined in He et al. (2014) as the solution to the following equation,

$$\frac{r^2}{1-r'^2} \Delta f = -\lambda f, \quad (1)$$

where f is a real-valued function, r is the radius function with r' its first derivative. We also call r' the slope of radius function r . The medial axis spectrum for a 2D shape defined above has several good properties such as invariance under isometry and uniform scaling, and robustness to boundary noise of a 2D shape.

It is however a challenging and non-trivial task to extend this definition from the medial axis of a 2D shape to that of a 3D shape. As shown by Fig. 1, the medial axis of a 2D shape is a 1D curve over which we can easily compute r' , the slope of r . In comparison, the medial axis of a 3D shape is a surface in 3D. It is thus not intuitive to find a slope similar metric on the medial axis radii of a 3D shape for characterizing the local geometry of the medial surface, which needs to be fed to our derived Laplacian eigenfunctions.

For the medial axis of a 2D shape, it is not hard to prove the relationship given by Eq. (2) (He et al., 2014) between the approximated hyperbolic distance d_h and the Euclidean distance d .

$$\frac{d_h}{d} = \sqrt{1-r'^2}. \quad (2)$$

It can be seen that the relationship is only related to r' . According to such a relationship, Eq. (1) could be represented as

$$\frac{r^2}{\left(\frac{d_h}{d}\right)^2} \Delta f = -\lambda f. \quad (3)$$

In the above equation, the denominator on the left hand side encodes the relationship between the approximated hyperbolic distance and the Euclidean distance, which essentially converts the distance from the Euclidean metric to the hyperbolic or Lorentz metric (Choi and Seidel, 2001). Although it is not easy to find a slope similar metric to capture this relationship for 3D shapes, we may be able to define a ratio that has a similar effect as $\sqrt{1-r'^2}$ for 2D shapes. We will define such an expected ratio for 3D shapes with the help of the Minkowski inner product (Catoni et al., 2011; Naber, 2003).

Definition 1. The Minkowski inner product of two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{n+1}$ is defined in the following form:

$$\langle \mathbf{a}, \mathbf{b} \rangle_M = a_1 b_1 + a_2 b_2 + \dots + a_n b_n - a_{n+1} b_{n+1}.$$

Note: We will use the Minkowski distance between two medial points \mathbf{a} and \mathbf{b} below, whose square is defined by $|\mathbf{a} - \mathbf{b}|_M^2 = \langle \mathbf{a} - \mathbf{b}, \mathbf{a} - \mathbf{b} \rangle_M$. The Minkowski inner product defined above might be negative, and so is the square of Minkowski distance. In this paper, if the Minkowski inner product $\langle \mathbf{a} - \mathbf{b}, \mathbf{a} - \mathbf{b} \rangle_M$ is negative, we always take $|\mathbf{a} - \mathbf{b}|_M = -\sqrt{-\langle \mathbf{a} - \mathbf{b}, \mathbf{a} - \mathbf{b} \rangle_M}$, i.e., we consider the Minkowski distance to be negative in this case.

Notice that the approximated hyperbolic distance (He et al., 2014) between \mathbf{p}_i and \mathbf{p}_j has the form $d_h = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 - (r_i - r_j)^2}$, which could also be represented as

$$d_h^2 = \langle \mathbf{p}_i - \mathbf{p}_j, \mathbf{p}_i - \mathbf{p}_j \rangle_M \quad (4)$$

in the Minkowski inner product form. As it is the Minkowski inner product of the difference between two medial axis points, we also call this the *Minkowski distance*. Similar to the factor $(1-r'^2)$ that makes the medial axis spectrum of a 2D shape robust to noise (He et al., 2014), the relationship between the Minkowski distance and Euclidean distance also makes the medial axis spectrum robust to noise.

For a 2D shape, we may measure the Minkowski distance on the medial axis composed of a series of line segments. For a 3D shape, since the basic elements of the medial axis are triangles, we use some kind of area metric instead.

Definition 2. Given a triangle $\triangle ABC$ on the medial axis surface of a 3D shape, its Minkowski area can be defined as

$$S_M = \frac{1}{2} |AB|_M h_{MC},$$

where $|AB|_M$ is the Minkowski distance between A and B , and h_{MC} is the Minkowski distance from C to the line segment AB , which is given by the minimum Minkowski distance between C and the points on AB .




Fig. 3. The negative Minkowski area problem. The gray spheres depict the medial points.

For a triangle $\triangle ABC$ on the medial axis of a 3D shape, we also have the equalities given in Eq. (5), whose proof can be found in Appendix A:

$$S_M = \frac{1}{2}|AB|_M h_{MC} = \frac{1}{2}|AC|_M h_{MB} = \frac{1}{2}|BC|_M h_{MA}. \quad (5)$$

Hence, the Minkowski area of triangles on the medial axis is independent of which side is chosen as the base. From Eq. (5), we can also draw the conclusion that as long as one of the Minkowski distance of the three sides is very small, the Minkowski area also remains small. Choi and Seidel (2001) considered the hyperbolic Hausdorff distance as an appropriate metric to measure the differences between medial axes. The hyperbolic distance takes into account the radius function such that an instable branch corresponding to boundary noise contributes little to the distance measure. By the relationship between the approximated hyperbolic distance and the Minkowski distance (Eq. (4)), it follows that the Minkowski area is very sensitive to the variation of noises on the medial axis of a 3D shape.

One point worth noting is that in the Euclidean space, for a triangle, the sum of the lengths of any two sides must be bigger than the length of the third side. In the Minkowski distance space, however, this property does not hold, even when all three edges have positive Minkowski distances and would result in negative Minkowski areas. On the highly instable medial axis of a 3D shape, there exist a few such triangles sometimes. Fig. 3 illustrates the negative Minkowski area problem. The figure on the left is the normal envelope structure of a medial surface with positive Minkowski area. For the figure on the right, although it has positive Minkowski distances between any two medial points, it has a negative Minkowski area as the Minkowski distance from one medial point to the opposite side is negative. A brief proof is given in Appendix A.

We further define the expected ratio on the medial axis of a 3D shape, which has a similar effect as $1 - r'^2$ on the medial axis of a 2D shape.

Definition 3. For a triangle $\triangle ABC$ on the medial axis of a 3D shape, the Minkowski–Euclidean ratio is defined as

$$\text{ratio}_{ME} = \frac{S_M}{S_E},$$

where S_M is the Minkowski area and S_E is the area in the Euclidean space.

We use this newly defined Minkowski–Euclidean ratio to substitute the term $1 - r'^2$ in Eq. (1) to define the medial axis spectrum of a 3D shape as follows.

Definition 4 (Medial axis spectrum of a 3D shape). Given a 3D shape S in \mathbb{R}^3 and let \mathcal{M} be the medial axis of S , the medial axis spectrum of S is defined as the solution to the following extended Laplacian eigenvalue problem on \mathcal{M} :

$$\frac{r^2}{\text{ratio}_{ME}} \Delta f = -\lambda f, \quad (6)$$

where f is a real-valued function, r is the radius function on \mathcal{M} , and ratio_{ME} is the Minkowski–Euclidean ratio.

The medial axis spectrum of a 3D shape defined above inherits the properties of that of a 2D shape, such as invariance to isometry of the medial axis and uniform scaling, and robustness to boundary noise. We will demonstrate these properties with some examples in Section 5.

4. Computation details and robustness analysis

We compute the medial axis of a 3D shape using the Voronoi-based method (Amenta and Bern, 1998). However, the medial axis thus obtained is highly instable due to boundary noise, which might sometimes cause the negative Minkowski

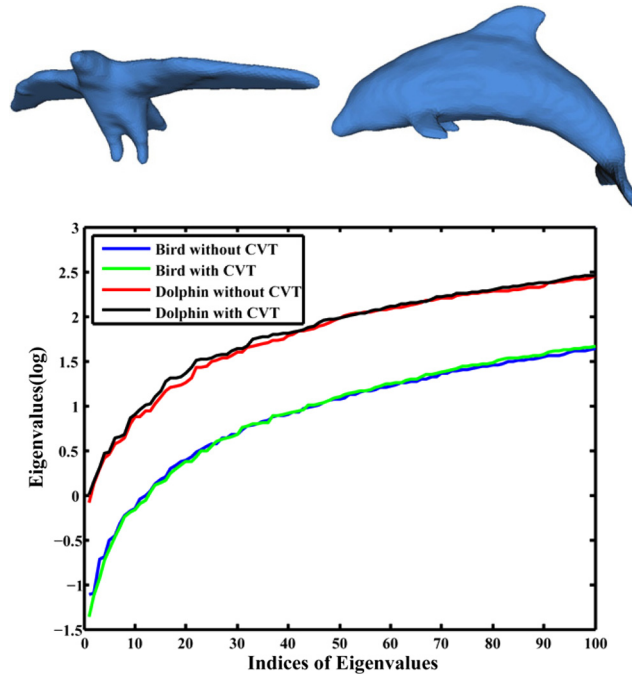


Fig. 4. Influence of CVT remeshing: The medial axes of the two shapes are simplified by Sun et al. (2013) until 4000 points remained. We then use the CVT remeshing method to sample them also to 4000 points. The bottom shows the difference between their medial axis spectrum without and with CVT remeshing.

area problem as shown in Fig. 3. To resolve this problem, we further use the medial axis simplification method (Sun et al., 2013), which contracts the edge with minimum post-contraction error (which measures the shape approximation error) in each step, to remove the negative Minkowski area triangles on the 3D medial axis. For such a triangle with negative Minkowski area on the medial axis as in Fig. 3(right), the edge connecting a big medial ball and a small one has a very small post-contraction error. Therefore, edges on triangles with negative Minkowski areas are contracted sooner and as a result, these triangles are prone to be removed firstly. The simplified medial axes are then taken as the input to spectrum computation in order to avoid the negative Minkowski area problem.

We use the method in Meyer et al. (2003) to solve the discretization of the eigenvalue problem defined by Eq. (6). The accuracy of the spectrum depends on both the number and the quality of the sample medial points. Different shapes usually produce different numbers of medial points by the Voronoi-based method (Amenta and Bern, 1998). The medial points obtained by the medial axis simplification method (Sun et al., 2013) are usually distributed on the medial axis unevenly. The unsatisfactory triangle quality thus induced may pose challenges for further processing. Moreover, medial axes of different shapes normally contain different number of points. When this difference is large, the accuracies of the spectrum of differ a lot as well. In order to have a comparable accuracy for each shape, we need to resample the medial axis into similar numbers of sample medial points with high quality. In this paper, we choose the CVT remeshing method to work on the 3D medial axis surface as in Zhu (2014). As the CVT remeshing method is designed for manifold surface, non-manifold edges are detected and labeled as features. Details of remeshed 3D medial axis surface extraction and radius estimation can be found in Zhu (2014), Yan et al. (2009). Fig. 4 shows that the CVT remeshing method has only slight influences to the medial axis spectrum.

In Eq. (6), r is defined on each medial point, while ratio_{ME} is defined on each triangle. We assign the value ratio_{ME} for each medial point the average ratio_{ME} of all the triangles associated to this point. In this paper, we use the Dirichlet boundary condition for the discretization of Eq. (6) in all the experimental results presented below. The boundary values are set to zero, and the eigenvalue problem is then solved to obtain the medial axis spectrum.

4.1. Non-manifold handling

The medial axis of a 3D shape comprises branch surfaces connecting at intersection curve segments, at which the medial balls are tangent to the shape boundary at three or more points. While the Laplace–Beltrami operator is defined on manifold domain (Meyer et al., 2003), we extend its computation to the non-manifold medial axis surface. Given an edge E on the non-manifold curve segment at which n triangles meet, the adapted Laplace–Beltrami operator is defined as the average of the ordinary Laplace operator acting on all possible manifolds that can be derived at E . Specifically, let X_1, \dots, X_n be the n neighbor triangles meeting at E in some order. The possible manifolds derived at E are $X_1EX_2, X_2EX_3, \dots, X_nEX_1$. The adapted Laplace–Beltrami operator at E is then the average of Laplace operators on these n manifolds. Let O be an

endpoint of the above non-manifold edge E , and hence O is a non-manifold point. The adapted Laplace–Beltrami operator at O corresponding to edge E is the average of the ordinary Laplace operators on these n derived manifolds. The Laplace operator at O then involves all of its incident edges.

4.2. Eigenvalue problem

The matrix discretized by Meyer et al. (2003) corresponding to the Laplace–Beltrami operator is sparse, symmetric, and semi-definite. Its eigenvalues are therefore non-negative. However, the irrational expression involving the radius function r and the Minkowski–Euclidean ratio $ratio_{ME}$ in Eq. (6) renders the resulting matrix non-symmetric. In order to solve this, the Laplace–Beltrami eigenvalue problem is then transformed to a generalized eigenvalue problem,

$$Af = \lambda Bf, \quad (7)$$

where both A and B are $m \times m$ matrices with m being the number of medial axis points. Here, A is a sparse, symmetric, and semi-definite matrix obtained from discretizing the Laplace–Beltrami operator. Essentially, it is a Laplacian matrix which encodes the medial axis structure. The matrix B is a diagonal positive definite matrix with the diagonal entries being $b_{ii} = \frac{ratio_{ME_i}}{r_i^2}$ for the i -th sample point on the medial axis. The value $ratio_{ME_i}$ is the averaged Minkowski–Euclidean ratio of all the triangles meeting at the i -th sample medial point.

Since A is a symmetric semi-definite matrix and B is a diagonal positive definite matrix, the generalized eigenvalue problem (7) is equivalent to an eigenvalue problem. Let matrices C and D be constructed by $C = DAD$ and $D^2 = B^{-1}$, respectively, the generalized eigenvalue problem (7) is transformed to

$$Cf = \lambda f. \quad (8)$$

This eigenvalue problem is finally solved by the numeric solver ARPACK (Lehoucq et al., 1997).

4.3. Analysis of the bounds of eigenvalues for perturbed medial axis

In this subsection, we discuss the robustness of the medial axis spectrum to noises via eigenvalue analysis. We first briefly analyze the influence on the spectrum when a noisy ridged surface is added to the original medial axis. Conclusions of quantitative analysis on the bounds of eigenvalue variations due to perturbation of the medial axis then follows directly.

4.3.1. Influence of adding a new ridged surface

For the medial axis of a shape, the eigenvalues can be obtained through Eq. (7). We then add a new noisy ridged surface to the medial axis, so that the generalized eigenvalue problem with the noisy surface can be represented by

$$\begin{pmatrix} A & Z \\ Z^T & A_n \end{pmatrix} = \lambda \begin{pmatrix} B & 0 \\ 0 & B_n \end{pmatrix},$$

where A and B are the matrices in Eq. (7) obtained from the original medial axis, A_n and B_n are the matrices corresponding to the newly added noisy surface, and Z is a sparse matrix connecting the original medial axis and the added noisy surface. Considering only the generalized eigenvalue problem on the noisy surface, we have

$$A_n f = \lambda B_n f = \left(\frac{\lambda}{k} \right) (kB_n) f.$$

In this sense, scaling the matrix B_n by a factor of k results in the scaling of the eigenvalues by a factor of $\frac{1}{k}$. Therefore, very small values on the diagonal of B_n will generate very large eigenvalues. A noisy surface on the medial axis has very small Minkowski–Euclidean ratio, which corresponds to very small values on the diagonal of B_n . If the connection between the original medial axis and the noisy surface is ignored, which also means $Z = 0$, some very big eigenvalues will be added to the medial axis spectrum.

When we consider the problem with connection between the medial axis and the noisy surface, the sparse matrix Z will have the effect of propagating the influence of the noisy surface to the original medial axis. If the Minkowski–Euclidean ratio approaches 0, this influence is negligible. The experiments in Section 5 demonstrate the influences on the medial axis spectrum when noisy surfaces are introduced.

4.3.2. Quantitative bounds of eigenvalues

Let \tilde{A} , \tilde{B} and $\{\tilde{\lambda}_i\}$ denote the matrices and eigenvalues of the perturbed system. We have the following relationship

$$|\tilde{\lambda}_i - \lambda_i| < \|\tilde{B}^{-1}\| (\|\lambda_i\| \cdot \|\tilde{B} - B\| + \|\tilde{A} - A\|), \quad (9)$$

where $\|\cdot\|$ represents the spectral matrix norm (Crawford, 1976). This implies that bigger eigenvalues exhibit larger variations than smaller ones.

Table 1
Computational time (time unit: second).

Shapes	MAT	CVT	10 eigens	20 eigens	50 eigens	100 eigens
Homer	160.479	19.045	0.396	1.068	4.431	8.642
Armadillo	99.219	17.295	0.42	1.318	4.466	11.552
Bunny	164.854	16.455	0.965	2.124	9.921	19.434



Fig. 5. The three shapes have about 5000 points.

Let \tilde{C} denote the matrix of the perturbed system corresponding to C , the total perturbation bound of all eigenvalues holds as

$$\sqrt{\sum_{i=0}^n (\tilde{\lambda}_i - \lambda_i)^2} \leq \|\tilde{C} - C\|_F, \quad (10)$$

where $\|\cdot\|_F$ denotes the Frobenius norm (Li, 1994). Please refer to Crawford (1976), Li (1994) for more details of the analysis deduction.

It follows that the medial axis spectrum is robust to noise.

5. Evaluation and experiments

In this section, we first report the computational time of the method, and then evaluate the effectiveness of our proposed medial axis spectrum with intensive experiments.

5.1. Computational time

We first need to compute the medial axis transform (MAT) and after that, the algorithm executes a CVT remeshing process. The final step is to compute the medial axis spectrum. In the three steps, computing the MAT dominates the computation. As solving the eigenvalue problem is also not trivial, the time consumed by this step depends on the number of eigenvalues needed. Table 1 shows the time spent on each step, where the three shapes with about 5000 points for collecting the time statistics are shown in Fig. 5. All these are run on a Windows 8 workstation with an Intel(R) Xeon(R) CPU (2 processors) and 12.0 GB main memory. From this table, it can be seen that computing the first 10 eigenvalues only takes less than 1 second, while computing the first 100 eigenvalues normally needs 8 to 20 seconds.

5.2. Invariance to uniform scaling

The medial axis spectrum is invariant to uniform scaling. Fig. 6 shows four 3D boxes. Shape b is generated by a uniform scaling of shape a , shape c and shape d are generated from shape a by two different non-uniform scalings. From the plotted graph of the eigenvalues, the medial axis spectrum successfully differentiates shape a from shape c and d .

5.3. Robustness to boundary noise

In Figs. 7 and 8, for both the plane and the ant, the boundary surfaces have about 5000 points, while the coarse medial axes obtained from the Voronoi-based method (Amenta and Bern, 1998) have about 15,000 points. MAs 1–4 are the simplified medial axes by Sun et al. (2013) and have about 100, 500, 2000 and 8000 points, respectively. These medial axis surfaces are then remeshed by about 4000 points using the method in Zhu (2014). The subfigures on the right show the eigenvalues of these medial axes with different levels of noises. We can see that the medial axis spectrum is robust to noise.




Fig. 6. The medial axis spectrum under uniform scaling. Shapes *b*, *c* and *d* are obtained from shape *a* by a uniform scaling, a non-uniform scaling and another non-uniform scaling, respectively.

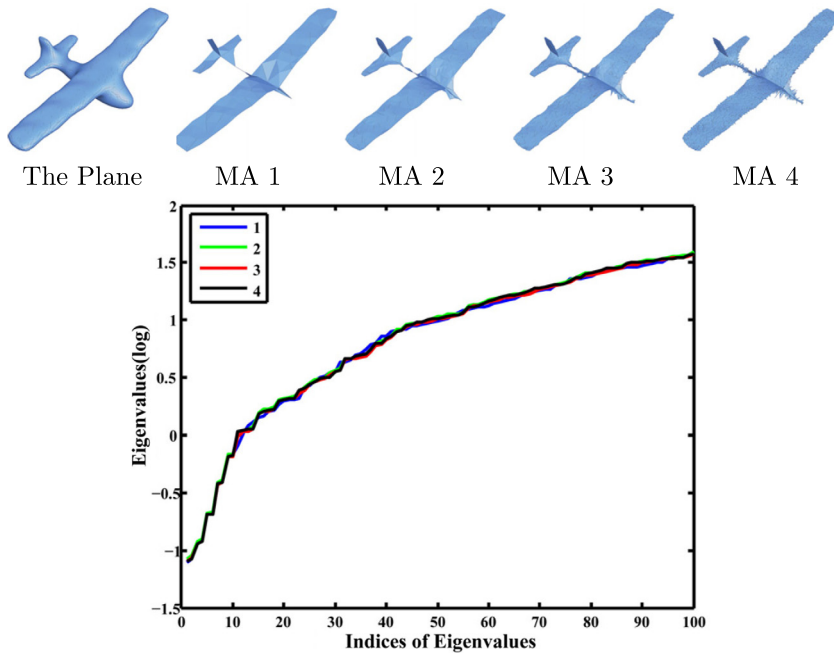


Fig. 7. Medial axis spectrum with different noise (The Plane).

As discussed in Subsection 4.3.1, adding one noisy surface to the medial axis introduces some bigger eigenvalues to the original eigenvalue set with negligible influence to the smaller eigenvalues. Fig. 9 shows such an example. The box on the left is smooth, while the right one is with a ridge added. The medial axis of this ridged box, as a result, shows a ridged surface in the middle part. For the medial axis of the smooth box on the left, there are about 3500 eigenvalues in total. Adding one noisy ridged surface introduces an extra 100 bigger eigenvalues. We can see that these 100 eigenvalues are much bigger than the original eigenvalues, and the addition of the noisy bump surface does not influence the original eigenvalues much.

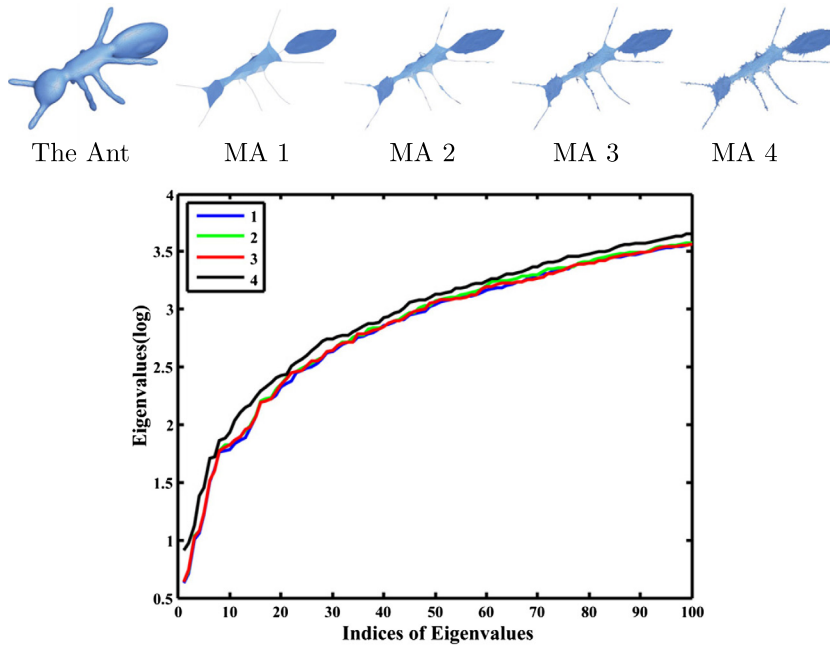


Fig. 8. Medial axis spectrum with different noise (The Ant).




Fig. 9. The addition of a noisy surface to the medial axis of a box introduces only some bigger eigenvalues, without influencing the original eigenvalues much.

5.4. Shape classification

In this subsection, we apply the medial axis spectrum to shape classification. Fig. 10 contains three groups of 3D shapes. It also shows eigenvalues (in log scale) of these shapes, where lines with the same color represent the first 100 eigenvalues of the shapes in the same group. We can see that using our method the shapes in the same group have similar eigenvalues. By comparison, the medial axis spectrum method is superior over the Shape-DNA method in classifying these 3D shapes.

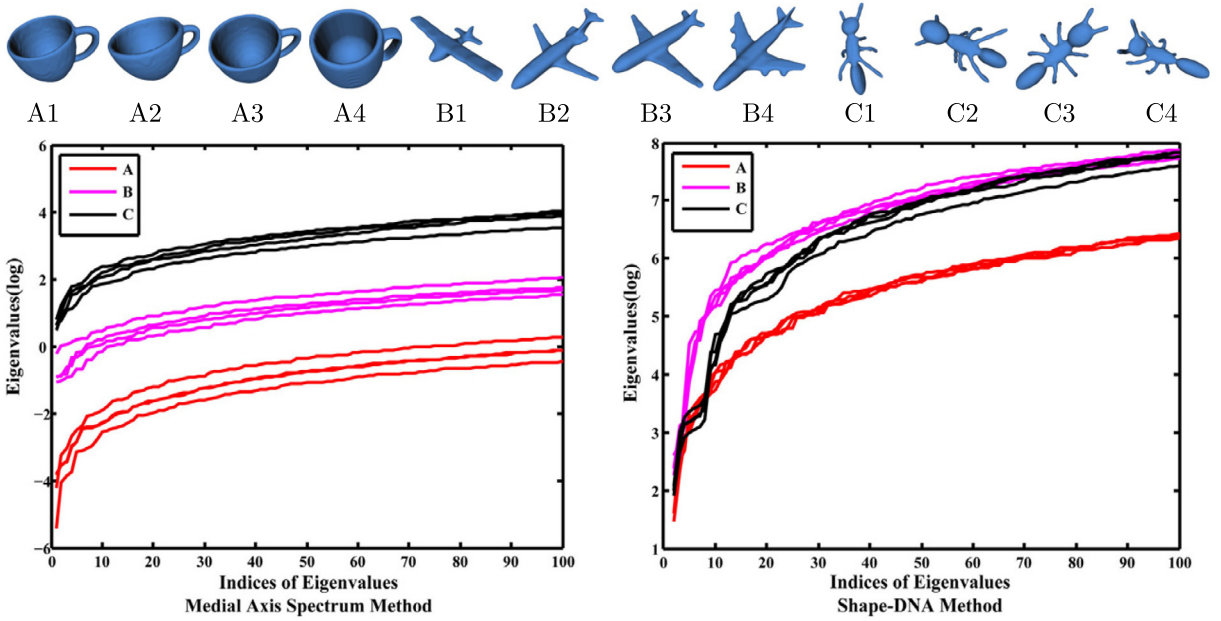


Fig. 10. Shape classification of medial axis spectrum and Shape-DNA.

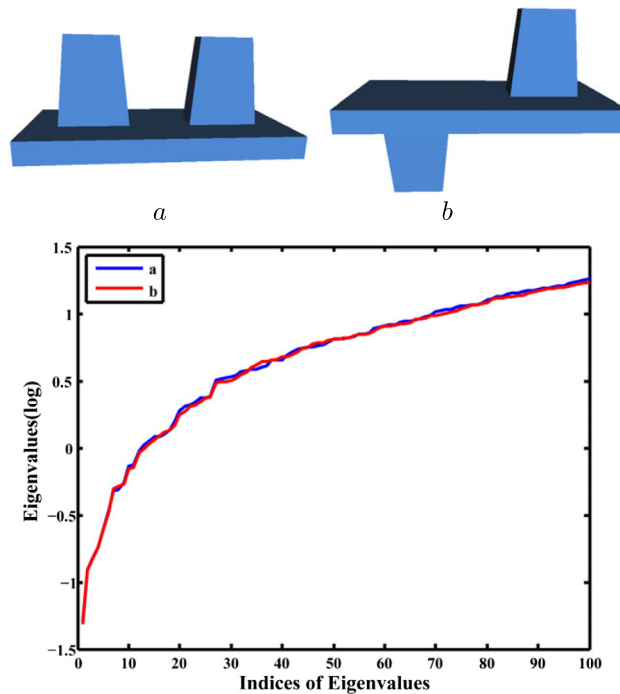


Fig. 11. The medial axis spectrum fails to distinguish between the above two shapes.

5.5. Limitations

The medial axis spectrum encodes local thickness and symmetry information as well as the part-structure of a shape. Hence, for those shapes having the same aforementioned properties, the medial axis spectrum will fail to distinguish them. Fig. 11 shows such an example. The medial axis of shape *a* presents the same topological structure and local thickness to the medial axis of shape *b*, except for the local extrusions in opposite directions. Their medial axis spectra are identical, as evidenced by the visualization of their eigenvalues. Another limitation of our method lies in the difficulty in handling partial shapes, due to the global nature of the spectral shape descriptor.

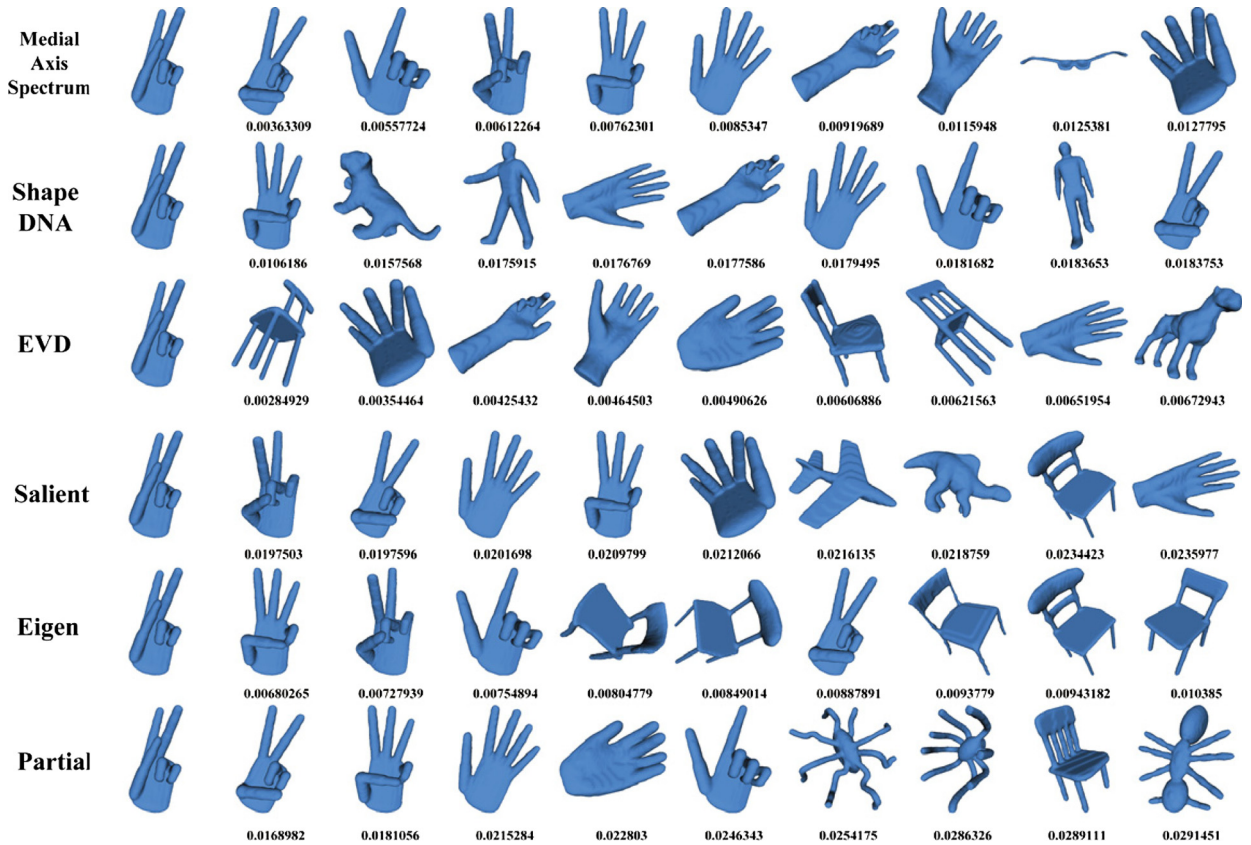


Fig. 12. A shape retrieval result using 400 shapes from McGill database. The first column shows the query shape, while the retrieval shapes are shown on the right in ascending order of dissimilarity distance.

6. Medial axis spectrum for shape retrieval

We further apply the medial axis spectrum to shape retrieval. Shape retrieval has received considerable attention during recent years. Many techniques such as Shape-DNA method (Reuter et al., 2006), spectral approach method (Jain and Zhang, 2007; Ruggeri et al., 2010; Isaacs, 2011), medial surface method (Siddiqi et al., 2008), the methods that can handle partial shapes (Hu and Hua, 2009; Jiang et al., 2013) and so on, have been developed to improve the performance of shape retrieval. It is necessary to compute distances between pairs of descriptors using a dissimilarity measure to capture how dissimilar two objects are. Different dissimilarity measures are used in shape matching to capture the dissimilarity between different 3D shapes (Reuter et al., 2006; Jain and Zhang, 2007; Tangelder and Velkamp, 2008). In order to compute the dissimilarity between different shapes, we use the L_2 norm to measure the difference of two k -dimensional vectors comprising the first k eigenvalues of the medial axis spectrum as the spectrum distance. According to the inequality in Eq. (9), large eigenvalues might have greater variations due to shape perturbation. In this sense, the difference between two k -dimensional vectors measured by the L_2 norm would be dominated by those larger eigenvalues. Hence, k should not be too large in order to yield useful spectrum comparison. In our experiments, by taking dense samples of about 2000 for the discretization of the medial axis of each shape, we find that the value of k around 50 works well for spectral comparison and therefore $k = 50$ is used for our results.

6.1. Shape retrieval comparison

We test our method on the McGill database (McGill, 2005) for verifying the shape discriminating ability of the medial axis spectrum in shape matching. A total of 400 shapes from the data set are used. All these shapes are classified into 18 different shape sets.

We compare our method with the Shape-DNA method (Reuter et al., 2006), the spectral approach with eigenvalue descriptor (EVD) (Jain and Zhang, 2007), the salient spectral geometric feature based method (Salient) (Hu and Hua, 2009), the Laplace–Beltrami eigenfunctions method (Eigen) (Isaacs, 2011) and the partial shape matching method (Partial) (Jiang et al., 2013). For Shape-DNA, we also use $k = 50$ for computing the spectrum distance as recommended by Reuter et al. (2006). For the spectral approach with EVD, we use the χ^2 -distance (Jain and Zhang, 2007) defined below to measure dissimilarity,

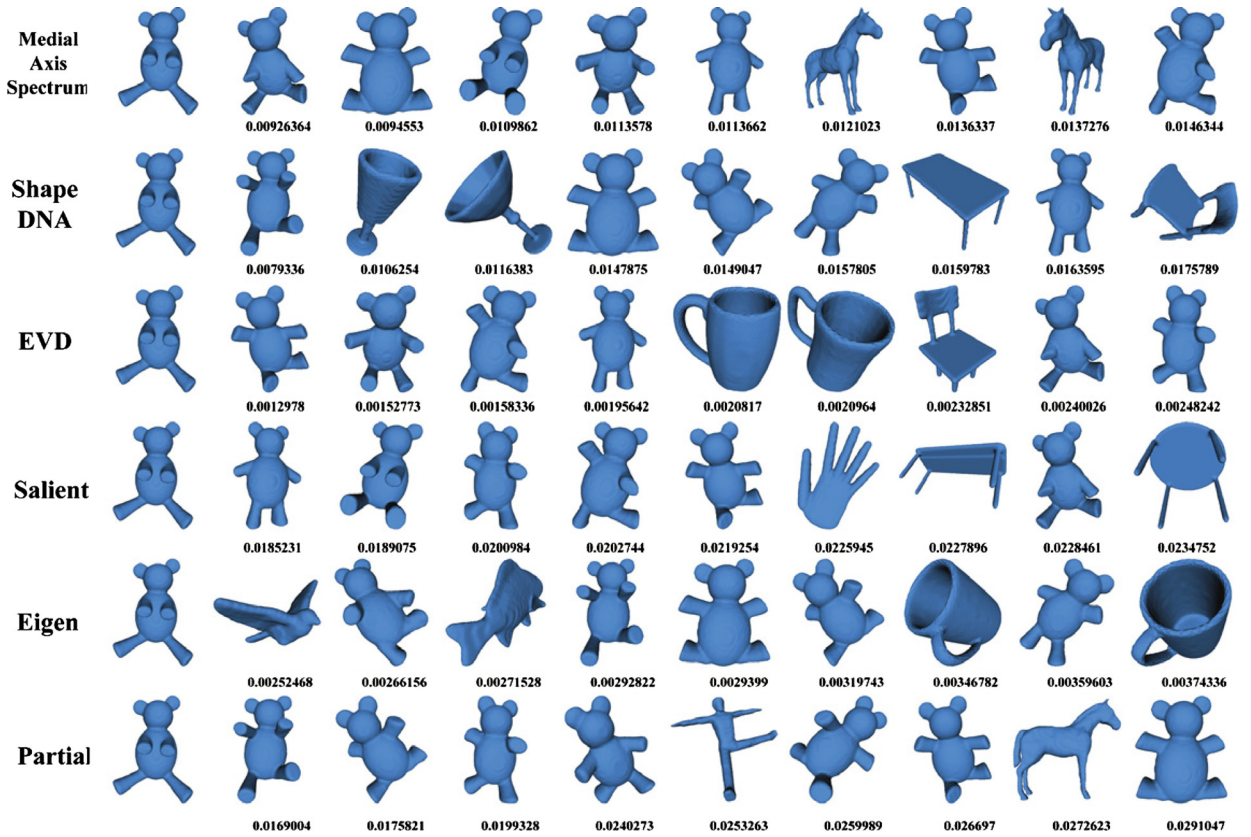


Fig. 13. Another shape retrieval result using the same database as Fig. 12.

$$Dist_{EVD}(P, Q) = \frac{1}{2} \sum_{i=1}^k \frac{[|\lambda_i^P|^{\frac{1}{2}} - |\lambda_i^Q|^{\frac{1}{2}}]^2}{|\lambda_i^P|^{\frac{1}{2}} + |\lambda_i^Q|^{\frac{1}{2}}}, \quad (11)$$

where $k = 20$ as proposed in Jain and Zhang (2007). The salient spectral geometric feature based method first extracts feature points, and then establishes correspondences between features of different shapes (Hu and Hua, 2009). A matching score, which can be figured out through solving an Integer Quadratic Programming problem, is used for retrieval. Then a normalized dissimilarity distance can also be defined based on the matching score. The Laplace–Beltrami Eigenfunction method also defines a feature metric for shape comparison (Isaacs, 2011). The partial shape matching method firstly segments the shape into several regions through the K-means clustering, and then finds the optimal solution of an assignment problem based on the Hungarian method (Jiang et al., 2013), which is a kind of distance between two shapes. Figs. 12, 13 and 14 show some of the results by the six methods. In each figure, the first column shows the query, and the nine closest matching results are shown on the right. These retrieved shapes are listed in ascending order by the dissimilarity distances, which are shown under the corresponding shapes. The results show that the medial axis spectrum method performs better than other methods.

Quantitative comparisons among these methods are also done by the receiver operating characteristic (ROC) curves (Fawcett, 2006; Manning et al., 2008). For a database with N shapes in total, there is a retrieval list $\{S_1, \dots, S_N\}$, which is ordered by ascending dissimilarity distances, for each query S . All the N shapes in the database can be classified into four categories for each number i , where $1 \leq i \leq N$. Two categories are from $\{S_1, \dots, S_i\}$, while the other two categories are from $\{S_{i+1}, \dots, S_N\}$. For a shape in the i closest matching $\{S_1, \dots, S_i\}$, if it is from the same shape set as S , it belongs to the category true positives (TP_i). Otherwise it falls into false positives (FP_i). Similarly, for a shape in $\{S_{i+1}, \dots, S_N\}$, it belongs to false negatives (FN_i) if it is from the same set as S . And it falls into true negatives (TN_i) otherwise. Two metrics can therefore be defined,

$$\text{False Positive Rate (FPR)}_i = \frac{FP_i}{FP_i + TN_i}, \quad (12)$$

$$\text{True Positive Rate (TPR)}_i = \frac{TP_i}{TP_i + FN_i}. \quad (13)$$

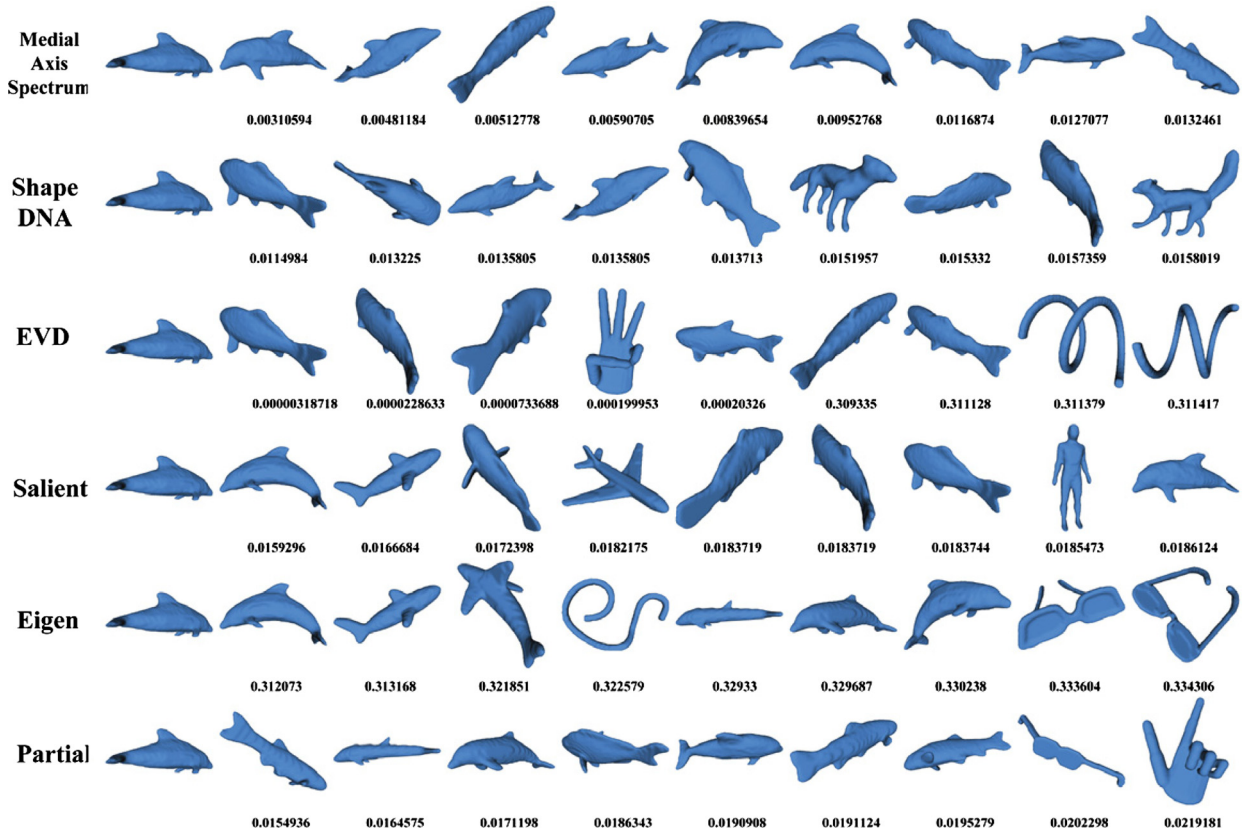


Fig. 14. Another shape retrieval result using the same database as Fig. 12.




Fig. 15. ROC curves of the medial axis spectrum, Shape-DNA, EVD, Salient, Eigen and Partial methods for retrieval of 400 shapes from the McGill database.

We can see that the fraction of wrong examples that are misclassified as a match can be measured by False Positive Rate (FPR_i), while the fraction of positive examples that are correctly labeled can be measured by True Positive Rate (TPR_i). The ROC curve of S is then got by all the (FPR_i, TPR_i) pairs, where i is traversed from 1 to N . We construct a collection of shapes by choosing two shapes from each shape set. Averaging the ROC curves of the collection of shapes forms the overall ROC curve. Comparisons among these methods are shown in Fig. 15. It is obvious that our medial axis spectrum method has a higher TPR value than all the other methods. Hence our method performs consistently better than other methods and has a superior shape retrieval performance.

7. Conclusion and further work

We have proposed the medial axis spectrum of a 3D shape, which is a novel medial axis based shape descriptor for 3D shape analysis. We demonstrate the several useful properties of the medial axis spectrum including invariance to uniform scaling and isometry of the medial axis, and consistency in the presence of shape boundary noises. As shown by our experiments, the medial axis spectrum works effectively in shape classification and retrieval. Based on the Minkowski–Euclidean ratio developed in this paper, we may be able to develop a framework to simplify the medial axis of a 3D shape. This will be our future work.

Acknowledgements

We would like to thank all the reviewers for their valuable comments and suggestions. This work is supported by the GRF grant of Hong Kong Research Grant Council (HKU 717813E), the National Natural Science Foundation of China under Grants 61272019 and 61332015, the 973 Program (2011CB302400), and the Shenzhen Science and Technology Program (JCYJ20140903112959962). Yanwen Guo is supported by the National Natural Science Foundation of China under Grants 61373059 and 61321491. Xiaohu Guo is partially supported by the National Science Foundation (NSF) under Grant Nos. IIS-1149737 and CNS-1012975.

Appendix A

We use \mathbf{a} , \mathbf{b} and \mathbf{c} to represent the edge vectors $B - C$, $C - A$ and $A - B$ of $\triangle ABC$ respectively. $|\mathbf{e}|_M$ denotes the Minkowski distance between the two endpoints of the edge \mathbf{e} . $|\mathbf{e}|_M^2 = \langle \mathbf{e}, \mathbf{e} \rangle_M$ might be negative according to the definition of Minkowski inner product.

Property 1. *The Minkowski area of a triangle $\triangle ABC$ on the medial axis is unique, no matter which side is chosen as the base.*

Proof. We first choose the edge \mathbf{c} as the base, so the height is the Minkowski distance from C to \mathbf{c} . Let $\mathbf{n} = \frac{\mathbf{c}}{|\mathbf{c}|_M}$. For any point P on the edge AB , it can be represented as $P = A + t\mathbf{n}$. So the squared distance between C and P is

$$dis_{MPC}^2 = |C - P|_M^2 = \langle t\mathbf{n} - \mathbf{b}, t\mathbf{n} - \mathbf{b} \rangle_M.$$

As h_{MC} is the minimum Minkowski distance between C and the point on \mathbf{c} , we can determine h_{MC} by the following equation

$$\frac{\partial(dis_{MPC}^2)}{\partial t} = 0.$$

We thus have

$$h_{MC}^2 = \langle \mathbf{b}, \mathbf{b} \rangle_M - \langle \mathbf{b}, \mathbf{n} \rangle_M^2.$$

According to the definition of Minkowski area, we have

$$4S_M^2 = \langle \mathbf{b}, \mathbf{b} \rangle_M \langle \mathbf{c}, \mathbf{c} \rangle_M - \langle \mathbf{b}, \mathbf{c} \rangle_M^2. \quad (14)$$

Then the squared Minkowski area can be represented as a form of the Minkowski inner product of A , B , and C . No matter which side is chosen as the base, the squared Minkowski area is always the same. \square

Property 2. *Assume that all the 3 edges of a triangle $\triangle ABC$ have positive Minkowski distances. If the sum of two lengths of any two sides is not greater than the third side, then the Minkowski area of $\triangle ABC$ is not greater than 0.*

Proof. Without loss of generality, we can assume

$$|\mathbf{a}|_M + |\mathbf{b}|_M \leq |\mathbf{c}|_M.$$

As all distances are positive, squaring both sides of the above inequality, we have

$$2|\mathbf{a}|_M |\mathbf{b}|_M \leq |\mathbf{c}|_M^2 - |\mathbf{a}|_M^2 - |\mathbf{b}|_M^2$$

Squaring the above inequality again, we get

$$2|\mathbf{a}|_M^2 |\mathbf{b}|_M^2 + 2|\mathbf{a}|_M^2 |\mathbf{c}|_M^2 + 2|\mathbf{b}|_M^2 |\mathbf{c}|_M^2 - |\mathbf{a}|_M^4 - |\mathbf{b}|_M^4 - |\mathbf{c}|_M^4 \leq 0. \quad (15)$$

Considering that

$$\mathbf{a} + \mathbf{b} + \mathbf{c} = \mathbf{0},$$

we have

$$|\mathbf{a}|_M^2 = |\mathbf{b} + \mathbf{c}|_M^2,$$

$$2 < \mathbf{b}, \mathbf{c} >_M = |\mathbf{a}|_M^2 - |\mathbf{b}|_M^2 - |\mathbf{c}|_M^2.$$

Putting the above expression into Eq. (14), we get

$$16S_M^2 = 2|\mathbf{a}|_M^2|\mathbf{b}|_M^2 + 2|\mathbf{a}|_M^2|\mathbf{c}|_M^2 + 2|\mathbf{b}|_M^2|\mathbf{c}|_M^2 - |\mathbf{a}|_M^4 - |\mathbf{b}|_M^4 - |\mathbf{c}|_M^4.$$

According to Eq. (15), we can reach the conclusion that the Minkowski area of $\triangle ABC$ is not greater than 0. \square

References

- Amenta, N., Bern, M., 1998. Surface reconstruction by Voronoi filtering. In: Proceedings of the Fourteenth Annual Symposium on Computational Geometry. SCG '98. ACM, New York, NY, USA, pp. 39–48.
- Auteri, F., Quartapelle, L., 2000. Galerkin–Legendre spectral method for the 3D Helmholtz equation. *J. Comput. Phys.* 161 (2), 454–483. <http://dx.doi.org/10.1006/jcph.2000.6504>.
- Bao, G., Wei, G.W., Zhao, S., 2004. Numerical solution of the Helmholtz equation with high wave numbers. *Int. J. Numer. Methods Eng.* 2004, 359–389.
- Blum, H., 1967. A transformation for extracting new descriptors of shape. In: Models for the Perception of Speech and Visual Form, pp. 362–380. <http://pageperso.lif.univ-mrs.fr/~edouard.thiel/rech/1967-blum.pdf>.
- Catoni, F., Boccaletti, D., Cannata, R., Catoni, V., Zampetti, P., 2011. Geometry of Minkowski Space–Time. Springer Briefs in Physics. Springer. <http://books.google.com.hk/books?id=D0ZNegIhosYC>.
- Choi, S.W., Seidel, H.-P., 2001. Hyperbolic Hausdorff distance for medial axis transform. *Graph. Models* 63 (5), 369–384. <http://dx.doi.org/10.1006/gmod.2001.0556>.
- Crawford, C.R., 1976. A stable generalized eigenvalue problem. *SIAM J. Numer. Anal.* 13 (6), 854–860. <http://www.jstor.org/stable/2156448>.
- Fawcett, T., 2006. An introduction to ROC analysis. *Pattern Recognit. Lett.* 27 (8), 861–874. <http://dx.doi.org/10.1016/j.patrec.2005.10.010>.
- He, S., Choi, Y.-K., Guo, Y., Wang, W., 2014. Spectral analysis on medial axis of 2D shapes. *Comput. Graph. Forum* 33 (8), 109–120. <http://dx.doi.org/10.1111/cgf.12404>.
- Hu, J., Hua, J., 2009. Salient spectral geometric features for shape matching and retrieval. *Vis. Comput.* 25 (5–7), 667–675.
- Isaacs, J.C., 2011. Laplace–Beltrami eigenfunctions for 3D shape matching. In: SPIE Defense, Security, and Sensing. International Society for Optics and Photonics, 80170Q.
- Jain, V., Zhang, H., 2007. A spectral approach to shape-based retrieval of articulated 3D models. *Comput. Aided Des.* 39 (5), 398–407. <http://dx.doi.org/10.1016/j.cad.2007.02.009>.
- Jiang, L., Zhang, X., Zhang, G., 2013. Partial shape matching of 3D models based on the Laplace–Beltrami operator eigenfunction. *J. Multimed.* 8 (6), 655–661.
- Lehoucq, R.B., Sorensen, D.C., Yang, C., 1997. ARPACK users' guide: solution of large scale eigenvalue problems by implicitly restarted Arnoldi methods.
- Levy, B., 2006. Laplace–Beltrami eigenfunctions towards an algorithm that “understands” geometry. In: Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006. SMI '06. IEEE Computer Society, Washington, DC, USA, pp. 1–13.
- Li, R.-C., 1994. Relative perturbation theory: (i) eigenvalue variations. Tech. Rep. UCB/CSD-94-855. EECS Department, University of California, Berkeley. <http://www.eecs.berkeley.edu/Pubs/TechRpts/1994/5481.html>.
- Manning, C.D., Raghavan, P., Schütze, H., 2008. Introduction to Information Retrieval. Cambridge University Press, New York, NY, USA.
- McGill, 2005. McGill 3D Shape Benchmark. Available at <http://www.cim.mcgill.ca/~shape/benchMark/>.
- Meyer, M., Desbrun, M., Schröder, P., Barr, A., 2003. Discrete differential-geometry operators for triangulated 2-manifolds. In: Hege, H.-C., Polthier, K. (Eds.), Visualization and Mathematics III. In: Mathematics and Visualization. Springer, Berlin, Heidelberg, pp. 35–57.
- Miklos, B., Giesen, J., Pauly, M., 2010. Discrete scale axis representations for 3D geometry. *ACM Trans. Graph.* 29 (4), 101:1–101:10. <http://dx.doi.org/10.1145/1778765.1778838>.
- Naber, G., 2003. The Geometry of Minkowski Spacetime: An Introduction to the Mathematics of the Special Theory of Relativity. Applied Mathematics Sciences. Dover Publications. <http://books.google.com.hk/books?id=pNfRHZwdVZ0C>.
- Reuter, M., Wolter, F.-E., Peinecke, N., 2005. Laplace-spectra as fingerprints for shape matching. In: Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling. SPM '05. ACM, pp. 101–106.
- Reuter, M., Wolter, F.-E., Peinecke, N., 2006. Laplace–Beltrami spectra as ‘Shape-DNA’ of surfaces and solids. *Comput. Aided Des.* 38 (4), 342–366. <http://dx.doi.org/10.1016/j.cad.2005.10.011>.
- Ruggeri, M.R., Patanè, G., Spagnuolo, M., Saupe, D., 2010. Spectral-driven isometry-invariant matching of 3D shapes. *Int. J. Comput. Vis.* 89 (2–3), 248–265. <http://dx.doi.org/10.1007/s11263-009-0250-0>.
- Rustamov, R.M., 2007. Laplace–Beltrami eigenfunctions for deformation invariant shape representation. In: Proceedings of the Fifth Eurographics Symposium on Geometry Processing. SGP '07. Eurographics Association, pp. 225–233. <http://dl.acm.org/citation.cfm?id=1281991.1282022>.
- Shokoufandeh, A., Macrini, D., Dickinson, S., Siddiqi, K., Zucker, S., 2005. Indexing hierarchical structures using graph spectra. *IEEE Trans. Pattern Anal. Mach. Intell.* 27 (7), 1125–1140. <http://dx.doi.org/10.1109/TPAMI.2005.142>.
- Siddiqi, K., Pizer, S., 2008. Medial Representations: Mathematics, Algorithms and Applications. Springer.
- Siddiqi, K., Shokoufandeh, A., Dickinson, S.J., Zucker, S.W., 1999. Shock graphs and shape matching. *Int. J. Comput. Vis.* 35 (1), 13–32. <http://dx.doi.org/10.1023/A:1008102926703>.
- Siddiqi, K., Zhang, J., Macrini, D., Shokoufandeh, A., Bouix, S., Dickinson, S., 2008. Retrieving articulated 3-D models using medial surfaces. *Mach. Vis. Appl.* 19 (4), 261–275. <http://dx.doi.org/10.1007/s00138-007-0097-8>.
- Sud, A., Foskey, M., Manocha, D., 2005. Homotopy-preserving medial axis simplification. In: Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling. SPM '05. ACM, New York, NY, USA, pp. 39–50.
- Sun, J., Ovsjanikov, M., Guibas, L., 2009. A concise and provably informative multi-scale signature based on heat diffusion. In: Proceedings of the Symposium on Geometry Processing. SGP '09. Eurographics Association, pp. 1383–1392. <http://dl.acm.org/citation.cfm?id=1735603.1735621>.
- Sun, F., Choi, Y.-K., Yu, Y., Wang, W., 2013. Medial meshes for volume approximation. arXiv:1308.3917. <http://dblp.uni-trier.de/db/journals/corr/corr1308.html#SunCYW13>.

- Tangelder, J.W., Veltkamp, R.C., 2008. A survey of content based 3D shape retrieval methods. *Multimed. Tools Appl.* 39 (3), 441–471. <http://dx.doi.org/10.1007/s11042-007-0181-0>.
- Yan, D.-M., Lévy, B., Liu, Y., Sun, F., Wang, W., 2009. Isotropic remeshing with fast and exact computation of restricted Voronoi diagram. In: *Proceedings of the Symposium on Geometry Processing. SGP '09, Aire-la-Ville, Switzerland*. Eurographics Association, pp. 1445–1454. <http://dl.acm.org/citation.cfm?id=1735603.1735629>.
- Zhu, Y., 2014. Compact representation of medial axis transform. Ph.D. thesis. The University of Hong Kong.