

Conformal spherical representation of 3D genus-zero meshes[☆]

Hongdong Li^{*}, Richard Hartley

Research School of Information Sciences and Engineering, The Australian National University, Canberra Research Labs, National ICT Australia Ltd., Australia

Received 1 February 2006; received in revised form 17 December 2006; accepted 15 January 2007

Abstract

This paper describes an approach of representing 3D shape by using a set of invariant spherical harmonic (SH) coefficients after conformal mapping. Specifically, a genus-zero 3D mesh object is first conformally mapped onto the unit sphere by using a modified discrete conformal mapping, where the modification is based on Möbius factorization and aims at obtaining a canonical conformal mapping. Then a SH analysis is applied to the resulting conformal spherical mesh. The obtained SH coefficients are further made invariant to translation and rotation, while at the same time retain the completeness, thanks to which the original shape information has been faithfully preserved.

© 2007 Published by Elsevier Ltd on behalf of Pattern Recognition Society.

Keywords: Conformal mapping; Shape invariant; Spherical harmonics; Shape representation

1. Introduction

Three-dimensional (3D) object recognition is one of the central topics in computer vision and pattern recognition research [1]. A good *shape representation* scheme is at the heart of any practical shape recognition systems. This paper aims at developing a new 3D shape representation and recognition method for general mesh objects. It is well known that in 2D shape recognition the *2D Fourier descriptor (2D-FD)* is an elegant and powerful technique possessing many desirable properties, such as being easy to compute, robust to noise, invariant to rotations, mathematically complete and having good discriminative ability. Our strategy here is to extend the technique to 3D.

Although it is conceptually straightforward, in practice such an extension is somewhat non-trivial. The main challenges arise from two tasks: (1) spherical parametrization; and (2) invariant computation.

We propose the use of a new discrete conformal mapping (DCM) in conjunction with the invariant spherical harmonics

(SH) for the solution. To our knowledge, this paper represents the first demonstration of such techniques to provide a general shape description for mesh surfaces. As the results of the paper will show, the method performs excellently in providing shape descriptors that can be used to represent 3D shapes. In principle, the shape descriptors are complete, in that they can be used to reconstruct the original surface.

In 2D-FD processing, the 2D shape (contour) is mapped onto a *unit circle*, by using an arc-length parametrization. This is followed by Fourier analysis on this circle. Analogously in 3D, the 3D object (surface) should first be mapped onto a *unit sphere*, followed by Fourier analysis on the sphere.

Fourier analysis on a sphere is not a difficult task; SH analysis is such a technique [2] and had been introduced to computer vision for 3D representation decades ago. But the real difficulty comes from *surface parametrization*. Unlike in 2D where arc-length is a natural parametrization, there is no natural way of doing surface parametrization for a general 3D surface, even though it does have a spherical-topology. Some conventional spherical parameterizations techniques exist, but seldom do they provide satisfactory results (as will be explained later). This paper provides a new method to tackle this difficulty. Our method is based on the DCM, which is a recently technique mainly developed in computer graphics area. So far, no systematic result has been reported on using DCM to study the separability of different 3D shapes.

[☆] A preliminary version of the paper was presented on ACCV'06 [1]. NICTA is funded through the Australian Government's Backing Australia's Ability Initiative, in part through the ARC.

^{*} Corresponding author. Tel.: +61 2 6125 8682; fax: +61 2 6125 8660.

E-mail address: Hongdong.li@anu.edu.au (H. Li).

We propose several key improvements to the DCM computation, to make it more suitable and more applicable for shape description and recognition. Among others, a new initialization scheme and a Möbius normalization method are the most notable ones, which substantially improve the efficiency of the computation. The Möbius normalization is very important to guarantee the uniqueness of the parametrization. This idea was actually inspired by the well-known method of projective stratification in vision reconstruction [3].

After DCM mapping, we then apply the SH expansion to derive a *complete* and *invariant* shape representation. We recognize that both the *invariance* (w.r.t. irrelevant transformations) and the *completeness* are equally important issues for a good shape representation. Most existing SH-invariants methods have, however, overlooked the second issue. As a result, using conventional SH invariants often lead to significant information loss, which consequently degrade the performance. We therefore propose a new method here to construct truly *complete* SH invariants, basing on a recent paper of SH computation [4]. By this method the SH coefficients are made invariant to irrelevant transformations, while at the same time retain the completeness.

We confine our analysis to genus-0 (i.e., spherical topology) mesh objects. This does not compromise the goal of the paper at this stage. The reasons are: apart from apparent theoretic value of genus-0, (1) a large portion of commonly encountered objects have spherical topologies. Some visually very complex objects, though look weird, have spherical topologies; (2) other object that does not have a spherical topology can be approximated by a genus-0 object using some recent graphics tools (e.g., alpha-shape or topological surgery); and (3) furthermore, higher genus objects can often be broken into genus-0 objects in some canonical ways.

We conducted experiments on a small set of meshes of complex geometries and different classes. Results show these 3D-FD descriptors are invariant to rotation/translation/scale, and robust to different tessellation, triangulation and resolution, and noise. In addition, they completely preserve the geometric information of the original shape. Shape comparison and recognition experiments have given encouraging results.

2. Shape representation

3D shape representation (description) has a rich history in computer vision and pattern recognition research. There has been a large number of methods existing for 3D shape representations. They can be roughly classified into two categories: *discriminative* and *complete*. The former one is designed to obtain a set of features in order to compare, retrieve, classify, and distinguish between different shapes. For example, a *sphericity* feature is sufficient to distinguish between an apple and a banana. While, the second one is a faithful and invertible representation of the original shape. It serves as a congruent coordinates system. This paper is devoted to this category, a complete representation. Our new 3D-FDs are truly complete in that all necessary information to describe and reconstruct the original shape is fully retained.

3. Previous work

3.1. Spherical parametrization

Because SH is defined on a sphere domain, it is vitally important to find an appropriate *spherical parametrization* for general 3D objects. A *spherical parametrization* is a bijective mapping from the object surface to the unit sphere. Most conventional SH-based surface representation often use naive parametrization method, such as use center-emitted rays to intersect the object surface, they therefore can only process convex objects or star-like objects [5].

Horn's EGI (and its many extensions) is a well-known and nice method for shape description [6,7]. It is based on the theory of Gauss map, hence has a solid theoretical ground. But, in general the Gauss map is not one-to-one for a concave object, therefore its many useful properties only enjoyed by convex objects.

Recent attempts at spherical parametrization of complex (e.g., concave, non-star-like, convolved, folded) 3D objects provided many other interesting approaches. A popular scheme is to gradually deform a surface until it maps onto the sphere (or conversely deform a sphere to the surface). Herbert, Ikeuchi and Delingette's SAI [6], and Sijbers and Dyck's 3D-Fourier [8] are examples of the kind. A shortcoming with them lies in the hardness to analyze the results, due to their heuristic nature. Brechbühler, et al. proposed an interesting method based on solving a heat conduction equation inside the 3D object volume [9]. However, the computation burden is very heavy and the final result depends on some user specified landmark positions.

In 3D model retrieval area some intuitive methods have been proposed. Kazhdan and Funkhouser et al. partitioned a 3D object volume into concentric spherical shells [2]. Similar ideas have been used by Saupe and Vranic [5]. Liu et al. designed a radon-like directional histogram model for 3D shape comparison [10]. The authors also extended the Haar-invariant method to 3D discrete structures [11]. The above methods, though may be capable of describing arbitrary-genus objects, are mainly discriminative, therefore are not suited to the object of the paper.

3.2. Discrete conformal mapping

DCM is a newly developed technique mainly in *computer graphics* and *CAD* area. Although the underlying mathematical principle was known centuries ago, how to apply them to modern digital meshes is still unfamiliar to many practitioners. Recently, several researchers have made remarkable progress along this direction [12–15].

DCM has many nice properties that make it especially suited to the application of surface parametrization. The most notable one is that it preserves angles, and hence local geometries. In addition, it depends only on the surface geometry (the Riemann metric), and therefore is robust to changes of data triangulation, resolution (level-of-detail) and noise. A recent paper provides a method for general 3D shape classification using conformal invariants [16].

However, when try to apply the DCM to 3D shape recognition, existing methods have several difficulties: firstly, many methods were designed for the applications where human interventions are acceptable or even preferable. For example, some methods require doing ‘cuttings’ on the surface, some need to specify some landmarks. Such interaction is obvious not suitable for automatic shape recognition; secondly, the uniqueness of the mapping is not always guaranteed, which subsequently causes ambiguity in shape representation; and thirdly, many methods are computationally expensive. Some take hours to get a single mapping.

In order to better enjoy the nice properties of DCM in shape description, while avoid the shortcomings, we propose a new and efficient DCM algorithm to overcome most of these problems. The efficiency of our new DCM algorithms comes from three innovative procedures: (1) initialization by planar-graph-drawing; (2) diffusion with exponential maps; and (3) normalization based on Möbius factorization.

4. New spherical parametrization

Our new method of spherical parametrization is based on the DCM computation. We first give a brief introduction to some mathematical backgrounds, then explain the new method in detail.

4.1. Mathematical backgrounds

Suppose M_1 and M_2 are two regular surfaces. A bijective differentiable mapping $f : M_1 \rightarrow M_2$ between the two surfaces is said to be *conformal* if it leaves the angle between curves on the surface invariant. A mapping between two surfaces is a conformal mapping if and only if it re-scales the *first fundamental forms* everywhere.

According to the celebrated Riemann Mapping Theorem, there always exists a conformal mapping between any two genus-zero surfaces. In particular, there exists a conformal mapping from any genus-0 surface to the unit sphere S^2 —a spherical parametrization of the surface.

However, such a conformal mapping is not unique since two such maps may differ by a further conformal mapping of S^2 to itself. The set of such mappings from S^2 to S^2 forms a 6D Lie group, the Möbius group, as will be explained now.

We identify the 2D sphere S^2 with the one-point-compactified complex plane $\mathbb{C} \cup \{\infty\}$ via the stereographic mapping

$$\varphi(x, y, z) = (x/(1-z) + iy/(1-z)), \quad (1)$$

where (x, y, z) are the Euclidean coordinates of a point on the unit sphere. The conformal mappings S^2 to itself are then simply the group of Möbius transforms of the complex plane given by

$$m(z) = \frac{az + b}{cz + d} \quad \text{with } ad - bc \neq 0, \quad (2)$$

where z, a, b, c and $d \in \mathbb{C}$. This transform has three complex (six real) parameters, since multiplication of a, b, c and d by a complex number does not change the transform.

To be precise, the set of conformal mappings of S^2 are those mappings of the form $\varphi^{-1} \circ m \circ \varphi$. Another way of thinking of this is to identify $\mathbb{C} \cup \{\infty\}$ with the one-dimensional complex projective plane PC^1 . Points in PC^1 are represented by complex 2-vectors. The space PC^1 has the topology of a real 2-sphere S^2 , and the stereographic mapping $\varphi : (x, y, z) \mapsto (x + iy, 1 - z)$ provides the homeomorphism between these spaces. The Möbius transforms acting on PC^1 are simply the group of projective transforms, represented by non-singular 2×2 complex matrices.

4.2. Harmonic mapping

In practice, the conformal mapping is often approximated by a harmonic mapping, denoted by f . Namely, it satisfies the following harmonic (Laplace) equation:

$$\Delta f = \text{div grad } f = 0. \quad (3)$$

For 3D genus-0 surfaces, these two mappings are essentially the same. Therefore, the problem of finding a spherical conformal parametrization is reduced to a Laplace-on-Manifold problem, where the target manifold is the unit sphere S^2 . Usually this is implemented by minimizing the following harmonic energy function [12,14]:

$$E_H(f) = \frac{1}{2} \int_{M_1} \|\text{grad } f\|^2. \quad (4)$$

The overall procedure is: first find a spherical homeomorphic initialization for the given shape, then iteratively modify this initial mapping by minimizing the above energy function, till it converges to a conformal mapping. The later stage is known as *diffusion*.

Our computation procedures basically follow [13], but made several unique contributions:

1. We introduce a new initialization method based on planar graph-drawing which effectively saves much computation, and alleviates fold-over problems.
2. We use the exponential map in solving the Laplace-on-manifold diffusion problem, thus enlarging the valid convergence basin, and improving convergence.
3. We introduce a Möbius factorization algorithm for Möbius normalization. This algorithm is simple, effective and much faster than other algorithms.

4.3. Initialization by planar graph drawing

We start from a triangulated closed mesh object homeomorphic to a sphere. The minimization of the harmonic energy is based on an iterative updating procedure. It thus requires a good initialization which serves as the starting point for the iteration. This initialization should be an approximation of the final conformal mapping.

Paper [13] provides an initialization method using the Spherical Tutte Embedding, where the Tutte Embedding itself is




Fig. 1. Initialization of a mesh from planar-graph-drawing.

computed by iterating from a Gauss map. Two shortcomings can be identified with it:

1. The spherical Tutte embedding, though having a good convergence property, has a computational requirement as great as that required by the subsequent harmonic diffusion. In other words, computation cost is doubled.
2. It does not converge correctly for some complex meshes, or gives rise to some fold-over triangles. Although theoretically the Tutte embedding should not yield fold-overs, here the reason is due to the use of the Gauss map, which is in general neither one-to-one nor orientation-preserving.

Based on the fact that the connectivity (adjacency) graph of any genus-0 3D object is always a 2D *planar graph* (which by definition is graph that can be drawn on a plane in such a way that there are no edge crossings), we propose a simple method for spherical initialization. Since our diffusion algorithm (described in the next subsection) has a relatively large convergence neighborhood, we do not require the initial mapping to be very close to the final result, as long as it is a homeomorphism.

There exist a number of constructive algorithms that are able to actually draw a graph on a plane without edge crossing. This is the process of *planar graph embedding*. Such a planar graph embedding is obvious a homeomorphism of the original mesh. In fact, every planar graph can be drawn so that all edges are straight, a so-called straight-line planar embedding. In addition, very efficient linear time algorithms for straight line embedding are available now [17].

Our initialization procedure is: first choose an arbitrary surface triangle as the boundary, then apply a straight-line planar graph embedding to the graph, followed by an inverse stereographic mapping to get the initial spherical mapping. Fig. 1 illustrates an example of a wolf mesh. Although the result depends on the specific choice of boundary triangle, this specialization is soon relaxed by the subsequent diffusion process.

4.4. Harmonic diffusion with exp-map

Having a homeomorphic spherical embedding as the initialization, the next step is to diffuse it to a conformal mapping. We accomplish this by solving a diffusion equation on the unit sphere, namely the Laplace–Beltrami equation:

$$\Delta_{S^2} f = \text{div}_{S^2} \text{grad}_{S^2} f = 0. \quad (5)$$




Fig. 2. The orthogonal map p and exponential (logarithm) map p' of vertex v .

Note that the Laplacian is defined here in terms of the local geometry of the target manifold. Instead of directly solving this equation in Cartesian coordinates, which could be very involved, we advocate the use of the tangent-plane-projection method, whereby the Laplace equation remains in its simple form, but acting on the tangent planes. To do this, it is necessary to cover the manifold (S^2) with charts and work locally in these charts. This idea has been used by Gotsman et al. [14] by an orthogonal map. The disadvantage of using the orthogonal map is that the resultant charts are quite small (in fact have radius $\pi/4$ measured in geodesic distance). Thus several charts are required to cover the whole surface.

Instead, we found that creating a chart using the exponential map (or inversely, the logarithm map) gives significantly better results in terms of convergence, and independence of the initial function. The exp-map on a manifold intuitively corresponds to expanding geodesic curves to the tangent plane (see Fig. 2). Using the exponential map, all mesh vertices be mapped one-to-one onto a single chart. The actual computation of such exp-map on the sphere is also very simple thanks to the Rodrigues' formula of matrix exponential (cf. [3,18]).

4.5. Normalization by Möbius factorization

After the previous step, a conformal mapping f from the surface M_1 to S^2 is obtained. However, this mapping is not unique, since it may be followed by another arbitrary conformal mapping of the sphere to itself. Such a conformal mapping can be represented by a Möbius transform. Thus, there exists a 6-parameter family of such mappings. In the current section, we focus on *normalizing* the mapping from M_1 to S^2 so that the remaining ambiguity consists only of 3D rotations of the sphere, a 3-parameter family.

A nested-iteration algorithm is suggested for simultaneous diffusion and normalization [13]. However, this is computationally extremely expensive, especially for large scale meshes. A simplified version by centering the mesh barycenter is thus further proposed, but still very inefficient and sometimes produces degenerate solution as pointed out by Gotsman [14]. Gotsman suggests using anchor points to solve it, but his results depend on particular choice of the anchors.

We propose a new method here, which accomplishes the Möbius normalization task very efficiently, only at the expense of negligible computation. This is done by a process that balances the surface area (or “weight”) distribution on the sphere by a Möbius factorization. Unlike [13,15], we carry out this normalization step after the diffusion process, rather than simultaneously. This leads to significant save of computation.

Consider a surface element dA located at a point $\mathbf{x}$ on M_1 . For the purpose of gaining an intuitive understanding of our method, we assume that this surface element has a “weight” proportional to area, and so the surface element may be thought of as having weight dA . Now, the mapping f maps this to an element of weight dA at point $f(\mathbf{x})$ on the sphere S^2 . The center of gravity of the surface mapped on to S^2 is given by $\int_{M_1} f(\mathbf{x}) dA$. What we want to do is adjust the mapping f so that this center of gravity is at the origin (center of the sphere).

If f_0 is an initial conformal mapping to S^2 , then the most general conformal mapping $f : M_1 \rightarrow S^2$ is of the form $\varphi^{-1} \circ m \circ \varphi \circ f_0$, where m is a Möbius transform. We want to find a Möbius transform m such that

$$\int_{M_1} \varphi^{-1} \circ m \circ \varphi \circ f_0(\mathbf{x}) dA = \mathbf{0}. \quad (6)$$

This could be done by searching over the 6-parameter family of all Möbius transforms. Note, however that applying a rotation to S^2 results in a rotation of the centre of gravity $\int_{M_1} f(\mathbf{x}) dA$, and hence does not change the truth or falsehood of condition (6). Rotations form a subgroup of the Möbius transforms of the sphere, and in seeking to enforce (6) we may factor out the rotations, thus reducing the search to a 3-parameter search.

Formally, it is verified that rotations of S^2 correspond precisely to those Möbius transforms represented by matrices of the form

$$Q = \begin{bmatrix} q_1 & q_2 \\ -\bar{q}_2 & \bar{q}_1 \end{bmatrix}. \quad (7)$$

An arbitrary Möbius transform can be factored as

$$M = Q \cdot R = \begin{bmatrix} q_1 & q_2 \\ -\bar{q}_2 & \bar{q}_1 \end{bmatrix} \cdot \begin{bmatrix} k & z_t \\ 0 & 1 \end{bmatrix}, \quad (8)$$

where $k \in \mathbb{R}$ and $z_t \in \mathbb{C}$. Thus, in enforcing (6), we may ignore the left-hand rotation matrix Q , and constrain our search to Möbius transforms of the form given by the right-hand matrix above. Such a transformation is of the type $z \mapsto kz + z_t$, which represents a scaling, followed by a complex translation (by z_t) in the complex plane. Transformations of this type form a 3-parameter family.

The above discussion was in terms of continuous surfaces. In the case of a triangulated mesh surface M_1 , we may consider just the vertices $\mathbf{v}_i$ of the mesh, and to each one assign a weight equal to the area of the corresponding region in a dual

tessellation. We then seek the solution to

$$\sum_i w_i \varphi^{-1} \circ m \circ \varphi \circ f_0(\mathbf{v}_i) = 0, \quad (9)$$

over all Möbius transforms of the form

$$m(z) = kz + z_t. \quad (10)$$

At first sight, this equation is nonlinear. By some elementary algebras, however, one soon reduces it to an equivalent linear system, for which a least square technique suffices.

Once this is solved, applying the corresponding spherical affine transformation R to the diffused result will give us a unique solution up to rotation. For example, for the Stanford “bunny” mesh, one of our experiments obtained the following affine factor:

$$R = \begin{bmatrix} 0.21491 & -0.00932 + 0.00583i \\ 0.00000 & 1.00000 \end{bmatrix},$$

whose effect (as can be directly ascertained) is approximately re-scaling in the radial direction.

5. Construct complete SH invariants

Any bounded $\mathbf{L}^2$ continuous function (real or complex) $f(\theta, \phi)$ defined on a sphere can always be decomposed into a finite set of SH coefficients C_ℓ^m , where $|m| \leq \ell$, $\ell = 0, 1, 2, 3, \dots, \ell_{\max}$, ℓ is called the degree (or frequency) of the SH expansion.

The key equation for SH expansion is the following, where $f(\theta, \phi)$ is any smooth function defined on the unit sphere ($0 \leq \theta \leq \pi$, $0 \leq \phi \leq 2\pi$):

$$f(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} C_\ell^m Y_\ell^m(\theta, \phi). \quad (11)$$

For these assumptions, the spherical expansion exists and converges. The function $Y_\ell^m(\theta, \phi)$ is called a SH function, and is given by

$$Y_\ell^m(\theta, \phi) = \sqrt{\left(\frac{(2\ell+1)(\ell-m)!}{4\pi(\ell+m)!} \right)} P_\ell^m(\cos(\theta)) e^{im\phi}.$$

The functions $P_\ell^m(x)$ are called *associated Legendre functions*, and are a set of orthogonal polynomials (i.e., *associated Legendre polynomials*).

The computed surface points are then used to calculate the coefficients C_ℓ^m , which depend on both n and m according to the following definition:

$$C_\ell^m = \int_0^{2\pi} \int_0^\pi d\phi d\theta \sin(\theta) f(\theta, \phi) Y_\ell^{m*},$$

where the asterisk denotes the complex conjugate.

The SH expansion has been employed in many different areas. Its definition and fast computation can be found elsewhere. In this paper we mainly address the issue of how to construct *complete SH invariants* in the context of 3D shape representation.

As its 2D counterpart of Fourier transform, 3D SH also has the nice property that the coefficients can be made invariant with respect to translation, rotation, and scale change. We are most interested in the rotational invariance, because others can be easily eliminated by a trivial pre-alignment operation, whereas eliminating the rotation is not so easy. The PCA-pre-alignment technique [5], though has been popularly adopted, proves to be unstable for noisy cases or shapes with high-order symmetries.

Previously, many authors suggest the following energy SH-invariants (EIs) [2,19]:

$$EI(\ell) = \sum_{|m| \leq \ell} \|C_\ell^m\|^2, \quad (12)$$

which is based on the fact that the squared magnitude of the SH coefficients at every frequency ℓ is independent of rotation. This method has drawbacks:

1. These invariants are not complete. This results in difficulty in discriminating shapes. For example, distinct shapes may have the same descriptors, and similar shapes may not be distinguishable. Moreover, it may not be possible to reconstruct the original shape from the invariants.
2. There is not only information loss but serious computation waste. For SH coefficients up to degree ℓ_{max} , there should be $((\ell_{max} + 1)^2)$ independent complex invariants. However, only $(\ell_{max} + 1)$ real energy invariants are obtained by the conventional method.

5.1. Our complete SH invariants

This paper provides a method of constructing a *complete* set of SH invariants. Completeness implies that the shape descriptors suffer no ambiguity in shape classification and recognition. We make use of a recent algorithm of estimating orientation from SHs [4]. The principle is the fact that SH coefficients at every frequency ℓ , $\ell \geq 1$, form an irreducible representation of the $SO(3)$ group. In other words, when a rotation is applied to the original function, the resulting SH coefficients will transform among themselves in exactly the same way.

Specifically, if we apply a rotation denoted by the Euler angles (α, β, γ) , we get new C_ℓ^m from the original $C_\ell^{m'}$ defined by

$$C_\ell^m = e^{-im(\alpha+\pi/2)} \cdot \sum_{|m'| < \ell} e^{-im'(\gamma+\pi/2)} C_\ell^{m'} \cdot \sum_{|k| < \ell} \mathcal{P}_\ell^{mk}(0) \mathcal{P}_\ell^{m'k}(0) e^{-ik(\beta+\pi)}, \quad (13)$$

where the two $\mathcal{P}_\ell(\cdot)$ are the *associated Legendre polynomials*.

Specifying some of the SH coefficients with some *canonical* values, we can estimate a *canonical* rotation $\mathbf{R}(\alpha^*, \beta^*, \gamma^*)$ by which the SH coefficients are transformed into complete rotation invariants. For instance, when $\ell = 2$, we can use the

following *canonical* values [20]:

$$C_2^1(\alpha^*, \beta^*, \gamma^*) = 0,$$

$$C_2^2(\alpha^*, \beta^*, \gamma^*) \text{ real, positive and maximal,}$$

$$\text{Re}(C_1^1(\alpha^*, \beta^*, \gamma^*)) \geq 0,$$

$$\text{Im}(C_1^1(\alpha^*, \beta^*, \gamma^*)) \geq 0.$$

We use $\ell = 2, 3, 4, 5$ in a least square fashion in our experiments for robust rotation-estimation. The subspace $\ell = 1$ is discarded as it is equivalent to the PCA-pre-alignment.

The use of such invariant definition is theoretically not entirely satisfactory, because it relies on identifying a canonical rotation and we suspect if this leads to a 2-fold ambiguity. Nevertheless, it has given good results in our experiments. We continue to look for better ways of defining rotationally invariants SH coefficients.

5.2. Shape functions in use

Now that we know how to compute a set of rotation-invariant SH coefficients of *shape functions*, we need to specify which function to use. One choice is the density, or *area ratio* function on S^2 induced by the conformal mapping $f : M_1 \rightarrow S^2$. Let T' be a facet in the dual mesh of M_1 , corresponding to a vertex $\mathbf{v}$ of the triangulation, and let $f(T')$ be the corresponding facet on S^2 . We define a function g on S^2 facet by facet on the mesh. The value assigned to each point of a facet $f(T')$ is equal to the area ratio $\text{area}(T')/\text{area}(f(T'))$. Note that this is essentially independent of the triangulation of the surface M_1 . For computational purposes, a delta-function of weight $\text{area}(T')$ placed at $f(\mathbf{v})$ may be used instead.

Other functions are possible, such as the radius function, in which each point $\mathbf{x}$ on S^2 is assigned the value of the radial distance of $f^{-1}(\mathbf{x})$ from the centroid of the surface M_1 . Yet another possibility is the mean-curvature function, whereby each point $\mathbf{x}$ on S^2 is assigned the value of the mean-curvature of the surface M_1 at the point $f^{-1}(\mathbf{x})$. Since M_1 is a mesh surface, it is necessary to consider the mean curvature of a smooth approximation to the mesh. In studying the general 3D shape similarity, the EGI/SAI family gave some valuable suggestions [6,7].

We are interested in a more rigorous and deep result based on the following theorem, also noted by [21,22]:

Theorem. A closed surface $f(u, v)$ in $\mathbb{R}^3$ with parameter (u, v) is determined by its conformal factors and its mean curvatures uniquely up to rigid motions.

Therefore, by using the density ratios and mean-curvature as the shape functions, the resultant SH invariants are indeed *complete*.

Nevertheless, in the future we will pursue a more ambitious target as follows. It seems that specifying both the density-ratio and mean-curvature functions on the sphere provides redundant information for a global genus-0 closed surface. We are not aware, however, how to specify a *minimal* amount of

information to determine the surface. We think it an interesting *inverse problem*, where we argue that a *regularization* approach might help.

6. The proposed 3D-FD algorithm

The overall 3D-FD algorithm proceeds as:

Input: a genus-zero 3D mesh object.

Output: a complete set of 3D Fourier descriptors.

1. Pre-normalize the object with respect to translation and scale.
2. Get initial spherical parametrization by planar-graph-drawing and inverse stereographic mapping.
3. Harmonic diffuse on the unit sphere with exp-map by minimizing formula (4).
4. Solve the normalization (9) by Möbius factorization, and apply the resulting spherical affine transformation.
5. Compute the complete SH invariants for the chosen shape functions on the sphere.

7. Reconstruction algorithm

Theoretically, the descriptors we have developed here are *complete* descriptors of the surface, from which it should be possible to reconstruct the surface, within accuracy limited by the degree of the SH expansion. Fig. 3 illustrates the processing stages of our method.

Here we give a preliminary reconstruction algorithm. In essence, the reconstruction is no other than *graph-embedding*. A graph embedding is a particular drawing of a graph, so that the embedding has no crossing edges. In other words, the graph is planar. Here we adopt the spring-based embedding algorithm [23]. We use it to draw (embed) the mesh in 3D Euclidean space. For our application, the graph to be embedded is guaranteed to be planar because it has a spherical topology. What we have now are all the neighboring distances (i.e., the edge lengths) and second order distances, which are in fact computed from the conformal factors and mean curvatures of the mesh (which can be recovered from the FD coefficients for proper shape functions).

The total springs energy is written as

$$W = \sum_{i=1}^{n-1} \sum_{j \in N_i} \frac{1}{2} k_{ij} (v_i - v_j - l_{ij})^2, \quad (14)$$

where the v_i, v_j are vertex coordinate vectors, and l_{ij} represent both the edge lengths and second order distances, n is the total number of vertices, N_i represents the up to second order neighborhood. k_{ij} is the spring constant—here we set it to the inverse of the corresponding distance.

We use a Gauss–Newton method to solve this minimization problem of Eq. (14). Our experiments show that the convergence is very fast. The iteration finishes usually within 10 steps for a meshes with about 2000 edges.

Recently, the authors are considering modifying the semi-definite embedding (SDE) algorithm [24], which could be adopted here and may result in a more efficient algorithm.

8. Experimental results

8.1. Spherical DCM embedding

We test our method on 26 genus-0 meshes of different classes. Some of them have complex geometries. These meshes are verified to have valid genus-zero topology. A topological checking and editing procedure can be employed if it is not the case.

We implement the algorithm in C++, and test it on an Intel-P4 3 GHz PC using win-XP OS. The experiments show that all the 26 meshes converge quickly with little human interaction.




Fig. 4. Some results of spherical conformal mapping.




Fig. 3. The processing flow-chart of our 3D-FD algorithm.




Fig. 5. This figure show some mesh objects, their corresponding energy-invariants (row-2) and our 3D-FD shape descriptors (row-3, magnitude part only). For better illustration, we only depict the first 36 coefficients. The actual number is about 200.

We have verified that the results of our spherical conformal mappings are indeed unique up to rotation, by applying the algorithm to other randomly rotated versions of the same object. Fig. 4 gives the conformal spherical parametrization results for some mesh objects used in our experiments.

8.2. Compute the new invariants

The *shape function* we use in experiments is a complex-valued function with the real and imaginary parts given by the radius and area-ratios, resp. Fig. 5 gives some resulting SH

invariants. The first row shows the original meshes. The second row shows the obtained EI invariants up to degree 15. The bottom row shows our complete invariant shape descriptors (magnitude part only). Note the similarity between alike shapes and the difference between unlike ones. While the EI descriptors also capture the shape geometry, ours retains more information due to its completeness nature.

We test the *robustness* of our algorithm to different triangulations, resolution and noise. This was done by first perform both subdivision-based *refinement* and edge-collapse-based *simplification* on the original meshes, to alter the resolution




Fig. 6. Obtained 3D-FD shape descriptors of the same object but with different triangulations and resolutions. Here the two horses have 1000 and 800 facets, respectively. The two rabbits have 2000 and 5000 facets and with different levels of noise added.

and tessellation, then introduce isotropic Gaussian noise to the vertex coordinates. Apply our algorithm again to the distorted meshes, the newly obtained 3D-FD descriptors are still very stable, which indicates that both the parametrization process and the invariant computation are stable. These results are more robust than that by Gu et al. [22] in term of degenerate handling.

Fig. 6 gives results on the horse and the “bunny” meshes. Note that even the phase parts of the new descriptors are rather consistence. Table 1 gives the Euclidean distance matrix for classifying several objects.

8.3. Reconstruction experiment: completeness

In order to test the completeness of the obtained shape invariants, we conduct shape reconstruction experiment. The algorithm we use is the spring-embedding. In Fig. 7 we show some results of shape reconstruction (Ref. [25] for details), where the inputs are the mean curvatures and conformal factors, and the outputs are 3D vertex coordinates. It clearly shows the reconstructions are almost identical to the original versions. The residual errors, in term of the spring’s energy at convergency,

Table 1
Euclidean distance matrix

						
	0.0000	0.0636	0.1469	0.5116	0.4682	0.5977
	0.0636	0.0000	0.1670	0.4227	0.3817	0.4078
	0.1469	0.1670	0.0000	0.3645	0.3391	0.4381
	0.5116	0.4227	0.3645	0.0000	0.2172	0.1839
	0.4682	0.3817	0.3391	0.2172	0.0000	0.1452
	0.5977	0.4078	0.4381	0.1839	0.1452	0.0000



Original Shape			
Reconstructed Shape			
W	4.23*e-08	0.00873	0.02412

Fig. 7. Shape reconstruction from spring-embedding [25]. (row-1: ground-truth meshes; row-2: reconstructed meshes; W is the resulting minimal springs energy)

are very small (e.g., $4.23e - 8$, 0.00873, and 0.02412 for the pawn, bunny and horse meshes, resp.).

We also test the relationship between the reconstruction error with the degree of the SH expansions. In other words, we want to find how many shape invariants are needed in order to achieve a prescribed representation error. Fig. 8 show the result for the bunny mesh with 2000 facets. The y-axis gives the RMS error in the recovered 3D coordinates at each vertex, and x-axis is the degree of the SH expansion. It is clear that when the SH degree is greater than 10 the RMS reconstruction error is very small. Fig. 9 gives a comparison of the original and the reconstructed bunny mesh at degree 10. There is visually little difference (Fig. 7).

8.4. Computation time

Our software code have not been carefully optimized. In this test the intention is only to compare the time used for the diffusion and the normalization. For the wolf meshes with 308 facets, the diffusion cost 12.1 s, while the normalization costs only 0.002 s. For computing the horse-2k and bunny-2k meshes, which has 2k facets each, cost about 86 and 62 s, respectively, in diffusion, but the time spent on normalization are 0.006 and 0.005 s, respectively.




Fig. 8. Shape reconstruction RMS error vs. SH degree. The size of the original shape is normalized to within a bounding box of side-length 1.0.




Fig. 9. Shape reconstruction RMS error vs. SH degree.

The following Table 2 shows the overall computation time for some of the meshes used in experiments. All the trials are repeated 10 times to get average estimations.

Because both the planar graph drawing and the SH expansion have fast algorithmic implementations (e.g., [4,17]), it is expected our algorithm will run more efficiently after appropriate optimization.

9. Conclusion and future work

The method of using SH applied to a new conformal spherical parametrization, proposed in this paper, works well in discriminating mesh objects of different shape, while being invari-

Table 2
Time used for computing the invariant conformal mapping

Mesh	V	E	F	Time (s)
bunny-1k	502	1500	1000	35.2
horse-500	252	750	500	23.1
fish-800	402	1200	800	28.5
pawn-304	154	456	304	16.0
wolf-534	269	801	534	26.8

Execution time = diffusion time plus normalization time, V , E , F are the number of vertices, edges and facets, respectively.

ant to rotation, robust to triangulation and noise in the model description.

The proposed new algorithm provides practical improvements over existing methods, and already shows good results. All these have added to the confidence of applying the DCM to various vision problems. Of course, more work still need to be done for many real-world applications, for example, medical anatomic shape comparison or 3D model retrieval. For the purpose of general 3D model retrieval, we do *not* recommend our method, as it does not work for non-genus-zero object, while in practice 3D meshes almost never satisfy the strong topological condition. To solve it, one would used some re-meshing tools, but such procedures are generally computational expensive and involve laborious operations. However, for the application of medical shape comparison, our method is promising. Currently, the proposed DCM model is under investigation for designing a new shape representation of 3D brain anatomical structure (e.g., hippocampus) for aiding medical diagnosis.

One of the benefits of the discrete conformal mapping is that it gives a unique spherical parametrization. An alternative scheme could be the spherical equiareal mapping (aka., area-preserving mapping) [12], which is more appropriate for the computation of the discrete SH expansion because the discrete SH expansion often implicitly assumes a uniform distribution of the vertices. However, because of the multiplicity of the equiareal mapping, without non-trivial modification this mapping can not be used for shape parametrization.

Another fascinating and challenging problem is to represent shapes using *minimal* information. This can find valuable applications in geometric compression, and will be a priority in our future work in this area.

Acknowledgments

The authors wish to thank F. Kahl for many insightful discussions regarding the conformal mapping. Ms. L. Zhou has helped the reconstruction experiments. Thank the anonymous reviewer who give invaluable suggestions. The mesh objects used in this paper are download from C. Grimm's homepage ([//www.cs.wustl.edu/cmng](http://www.cs.wustl.edu/cmng)), and is gratefully acknowledged.

References

- [1] H. Li, R. Hartley, New 3D Fourier descriptors for genus-0 mesh objects, in: Proceedings of the ACCV-2006, 2006.
- [2] M. Kazhdan, T. Fukhouser, S. Rusinkiewicz, Rotation invariant spherical harmonic representation of 3D descriptors, Eurographics'03 Symposium on Geometry Process, 2003.
- [3] R. Hartley, A. Zisserman, Multiple View Geometry in Computer Vision, second ed., Cambridge University Press, Cambridge, 2004.
- [4] A. Makadia, K. Daniilidis, Rotation estimation from spherical images, in: Proceedings of the IEEE-ICPR-2004, 2004.
- [5] D. Vranic, An improvement of rotation invariant 3D shape descriptor, in: Proceedings of the IEEE-ICIP-2003, 2003.
- [6] H. Shum, M. Hebert, K. Ikeuchi, On 3D shape similarity, in: Proceedings of the IEEE-CVPR-1996, 1996, pp. 526–531.
- [7] S.B. Kang, K. Ikeuchi, The complex EGI: a new representation for 3-D pose determination, IEEE-T-PAMI, 1993, pp. 707–721.
- [8] J. Sijbers, T. Ceulemans, D.V. Dyck, Algorithm for the computation of 3D Fourier descriptors, in: Proceedings of the ICPR-2002, vol. 2, 2002, pp. 11–15.
- [9] M. Quicken, C. Brechbller, J. Hug, H. Blattmann, G. Szekely, Parameterization of closed surfaces for parametric surface description, in: Proceedings of the IEEE-CVPR-2000, 2000.
- [10] X. Liu, R. Sun, S. Kang, H. Shum, Directional histogram model for 3D shape similarity, in: Proceedings of the IEEE-CVPR, 2003, pp. 813–820.
- [11] H. Burkhardt, M. Reisert, H. Li, Invariant for discrete structures—an extension of Haar invariant over Dirac delta functions, in: Proceedings of the 26th Pattern Recognition Symposium, DAGM-2004, September 2004, Germany.
- [12] M. Floater, K. Hormann, Surface parameterization: a tutorial and survey, Advances in Multiresolution for Geometric Modelling, Springer, Berlin, 2005 pp. 157–186.
- [13] X. Gu, et al., Genus zero surface conformal mapping and its application to brain surface mapping, IEEE Trans. Med. Imaging 23 (8) (2004).
- [14] C. Gotsman, X. Gu, A. Sheffer, Fundamentals of spherical parameterization for 3D meshes, ACM Trans Graphics 22 (3) (2003) 358–363.
- [15] M. Jin, et al., Optimal global conformal surface parameterization, in: Proceedings of the IEEE-Vis'04, 2004, pp. 267–274.
- [16] X. Gu, S. Yau, Surface classification using conformal structures, in: Proceedings of the ICCV-2003, Nice, France, 2003, pp. 701–708.
- [17] J. Boyer, W. Myrvold, Stop minding your p 's and q 's: a simplified $O(n)$ planar embedding algorithm, in: Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, 1999, pp. 140–146.
- [18] J.J. Koenderink, Solid Shape, MIT Press, Cambridge, MA, 1990.
- [19] M. Novotni, R. Klein, 3D Zenike descriptors for content based shape retrieval, in: Proceedings of the ACM Solid Model-03, 2003.
- [20] G. Burel, H. Henocq, Determination of the orientation of 3D objects using spherical harmonics, CVGIP:GMIP, vol. 57(5), 1996, pp. 400–408.
- [21] G. Kamberov, G. Kamberova, Conformal method for quantitative shape extraction: performance evaluation, in: Proceedings of the ICPR-2004, 2004.
- [22] X. Gu, Y. Wang, et al., Geometric compression using Riemann surface structure, Commun. Info. Sys 3 (3) (2004) 171–182.
- [23] T. Fruchterman, E. Reingold, Graph drawing by force-directed placement, Software Pract. Exper. 21 (11) (1991) 1129–1164.
- [24] K. Weinberger, L. Saul, Unsupervised learning of image manifolds by semidefinite programming, in: Proceedings of the CVPR-2004, 2004.
- [25] H. Li, R. Hartley, H. Burkhardt, Discrete Conformal Shape Representation and Reconstruction of 3D Mesh Objects, ICIAP-2005.

About the Author—HONGDONG LI, obtained his PhD degree from Zhejiang University, China, majored in Information and Electronic Engineering. He is currently a Research Fellow with the Research School of Information Sciences and Engineering (RSISE) at Australian National University (ANU). He is also a seconded Researcher to National ICT Australia Ltd.

About the Author—RICHARD HARTLEY, Professor, Fellow of Australian Academy of Sciences, Research School of Information Sciences and Engineering, Institute of Advanced Studies, Australian National University. VISTA Program, National ICT Australia, Canberra Labs.