Frequency-Based Representation of 3D Models using Spherical Harmonics

M. MOUSA R. CHAINE S. AKKOUCHE
L.I.R.I.S : Lyon Research Center for Images and Intelligent Information Systems
Bâtiment Nautibus, 8 boulevard Niels Bohr
69622 Villeurbanne Cedex, FRANCE
{mmousa, rchaine, sakkouch}@liris.cnrs.fr

ABSTRACT

3D meshes are the most common representation of 3D models. However, surfaces represented by 3D meshes may contain noise or some unrequired details. Multiresolution representations and filtering techniques are very useful in this case. In this paper, we propose a new and compact representation for the surface of a general 3D mesh using the spherical harmonics. This representation can be useful in many applications such as filtering, progressive transmission and compression of 3D surfaces. First, we present a basic framework for star-shaped objects. Then, we show how to extend this framework to general form meshes using certain segmentation techniques in combination with implicit surface techniques. An interesting feature of our approach is that the computation of the involved spherical harmonics transform is decomposed into the computation of spherical harmonics transforms based on elementary triangles which compose the mesh. This feature shows that the complexity of the computation of the used spherical harmonics transform linearly dependant on the number of triangles of the mesh. We present some experimental results which demonstrate our technique.

Keywords
Spherical harmonics, triangulated mesh, implicit surfaces, filtering, geometric modeling.

1 INTRODUCTION

Polygonal meshes remain the primary representation of 3D models. The recent development tools for scanning and modelling devices allows these meshes to contain finer details. However, some applications need not these high details to be kept, especially when they cannot be distinguished from noise. These details add geometric and topological complexity to the 3D models, which affects model retrieval applications as well as visualization applications. Due to this demand, the multiresolution representation of 3D meshes has been raised in many research areas, especially in the field of digital geometry processing (DGP) [SS01]. Furthermore, multiresolution representation is interesting for compression and progressive transmission purposes.

WSCG 2006 conference proceedings, ISBN 80-86943-03-8
WSCG'2006, January 30 – February 3, 2006
Plzen, Czech Republic.

Multiresolution representation and surface filtering have received a renewal of interest during the recent years [Bül02, CDR00, DMSB99, GSS99, ZBS04]. In signal processing community, the fundamental step is based on the construction of the spectrum of frequencies of the surface with respect to a set of basis functions. The low frequency components in the signal correspond to smooth features, and the high frequency components correspond to finer details such as creases, folds and corners. Retaining just lower frequency components suffices to represent and to capture the overall perceptual shape of the model.

Analyzing 3D meshes for retrieval purpose using the spherical harmonics has been developed during the last decade. As in [FMK+03, KFR03, KFR04, SV01], the spherical harmonics transform of spherical functions induced by the mesh can be calculated using a voxelization of the mesh as a preprocessing step. However, representing and filtering the 3D models using the spherical harmonics have become very interesting thanks to the efforts of Zhou et al. [ZBS04]. They calculate a frequency-based representation of 0-genus meshes using a spherical harmonics transform of spherical conformal parametrization of the mesh. The extension to higher genus meshes is performed by

operating a prior surface cutting along some user-specified closed paths to form a surface with the same topology as the sphere. However, cracks may occur along the cutting boundaries after filtering. Moreover, they consider a separate spherical function for each coordinate component x, y and z of the points of the mesh. They filter each function of $x(\theta,\varphi)$, $y(\theta,\varphi)$ and $z(\theta,\varphi)$ independently of the others without considering the dependencies between these three functions on the surface. Moreover, using three independent functions without regarding the correlation between these functions on the surface causes information redundancy.

Spherical harmonics are not the only approach used to represent the surface at several levels of details. There is also the spherical wavelet techniques [EDD+95, JDBP04] which rely on the same framework as Zhou et al.[ZBS04] except that they applied the spherical wavelet transform instead of spherical harmonics transform. Another example is the Laplacian operator [Tau95] which can smooth large meshes quickly, however, due to the huge computation of the eigenvectors decomposition, this method is limited to low pass filtering and cannot be used for general filter design.

Our contribution

In this paper, the spherical harmonics transform of a radial function induced by a mesh is computed in a cumulative manner. That is, the transform is applied independently to each triangle representing the mesh and then the results are summed up.

0-genus object can be represented by a set of three spherical functions [ZBS04] thanks to a previous conformal parametrization, however, in the case of a star-shaped object, a single spherical function is enough and does not require any previous parametrization. Therefore, there no information redundancy as the case of using three independent spherical functions in [ZBS04]. Moreover, it allows to exploit the correlation between the x, y and z coordinates on the surface. In case that the object is not star-shaped, we decompose it into star-shaped parts by a robust segmentation method, and show how a frequency-based description of the whole object can be obtained from the frequency-based description of the parts. For that, we revert to an implicit formulation for each part and blending techniques to extend to the entire object.

This paper is organized as follows. Section 2 recalls a brief mathematical background of the spherical harmonics. Section 3 shows how to represent star-shaped objects using the spherical harmonics transform that is calculated directly on the mesh without voxelization. Section 4 gives a generalization of our method to represent general meshes using the spherical harmonics by the means of the segmentation and the implicit framework techniques. Section 5 shows the use of our surface representation to filter general models using the spherical harmonics. We give some experimental results that demonstrate our approach in section 6. Finally, we conclude in section 7.

2 BACKGROUND

Spherical harmonics $\{Y_l^m(\theta,\varphi) : |m| \leq l \in \mathbb{N}\}$ are special functions defined on the unit sphere $\mathbb{S}^2$ [Bye59, Hob55] as :

$$Y_l^m(\theta,\varphi) = (-1)^m k_{l,m} P_l^m(\cos\theta) e^{im\varphi} \qquad (1)$$

where $\theta \in [0\ \pi]$, $\varphi \in [0\ 2\pi]$, $k_{l,m}$ is a constant, and P_l^m is the associated Legendre polynomial. The spherical harmonics are orthonormal functions such that:

$$\int_0^{2\pi}\int_0^{\pi} Y_l^m(\theta,\varphi)\overline{Y}_{l'}^{m'}(\theta,\varphi)\sin(\theta)d\theta d\varphi = \delta_{l,l'}\delta_{m,m'} \qquad (2)$$

where $\delta_{u,v}$ is the Kronecker delta function defined as the following :

$$\delta_{u,v} = \begin{cases} 1 & if\ u = v; \\ 0 & otherwise. \end{cases} \qquad (3)$$

Therefore, any spherical function $f : \mathbb{S}^2 \longrightarrow \mathbb{R}$ can be expanded as a linear combination of spherical harmonics :

$$f(\theta,\varphi) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} c_{l,m} Y_l^m(\theta,\varphi) \qquad (4)$$

where the coefficients $c_{l,m}$ are uniquely determined by :

$$c_{l,m} = \int_0^{2\pi}\int_0^{\pi} f(\theta,\varphi)\overline{Y}_l^m(\theta,\varphi)\sin(\theta)d\theta d\varphi \qquad (5)$$

Since f is a real valued function, the coefficients $c_{l,m}$ are related to each other by the following relation :

$$c_{l,-m} = (-1)^m \bar{c}_{l,m} \qquad (6)$$

3 SPHERICAL HARMONIC REPRESENTATION OF STAR-SHAPED OBJECTS

Let M denote a triangulated mesh of an object embedded in $\mathbb{R}^3$. M is said to be star-shaped if there exists a point $c \in \mathbb{R}^3$ such that every line segment drawn from c in any direction intersects the surface of M at exactly one point. Considering that c is the center of the spherical coordinate system, the radial function induced by M and c is a well-defined spherical function $f : \mathbb{S}^2 \longrightarrow \mathbb{R}^+$, where $\mathbb{S}^2$ is the unit sphere. More formally, this function is

defined as the following: for each (θ, φ) let p be the intersecting point with M in the direction (θ, φ)

$$f(\theta, \varphi) = d(c, p) \tag{7}$$

where d is the euclidean distance.

In this section, we will show how to represent a star-shaped object using the spherical harmonics. To do this, we propose to represent the object by its radial function $f(\theta, \varphi)$ with respect to a centre c. The spherical harmonics transform (SHT) is applied to this radial function taking into account that the surface is made of triangles.

3.1 Spherical harmonics transform of a star-shaped mesh

Let M denote a star-shaped triangulated mesh with respect to a point c. In the case that f is a binary function, standard algorithms [KFR03, FMK+03] compute a voxelization of M and then use this discrete approximation to find the coefficients of the spherical harmonics transform of f. The discretization introduces uncontrolled errors in the integrations needed to compute the harmonic coefficients.

In [MCA] we have proposed a fast and robust algorithm to calculate the spherical harmonics transform of triangulated meshes indicator function without prior voxelization, extending the decomposition idea proposed in [ZC01]. The calculations are performed independently over the triangles of M and then are summed up to obtain the final transform of M. We can apply the same algorithm to perform the spherical harmonics transform for the radial function defined by equation 7. Assuming that $f_i(\theta, \varphi)$ is the partial radial function defined over the triangle i with respect to a point c, the global radial function $f(\theta, \varphi)$ can be decomposed in terms of $f_i(\theta, \varphi)$ as follows :

$$f(\theta, \varphi) = \sum_{i \in T} f_i(\theta, \varphi) \tag{8}$$

where T is the set of triangles of M. The spherical harmonics transform of f_i is given by :

$$f_i(\theta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_{l,m}^i Y_l^m(\theta, \varphi) \tag{9}$$

Therefore, the expansion of $f(\theta, \varphi)$ can be rewritten as follows :

$$f(\theta, \varphi) = \sum_{i \in T} \left(\sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_{l,m}^i Y_l^m(\theta, \varphi) \right) \tag{10}$$

3.2 Approximating the signal

Theoretically, the expansion of $f(\theta, \varphi)$ is an infinite sum of the spherical harmonics. However, the high order coefficients $c_{l,m}$ obtained by summing up the $c_{l,m}^i$ are corresponding to finer details of the surface, and maybe noise. To filter the surface under a given precision, the summation is limited to a bandwidth bw. We then obtain an approximated surface.

$$\hat{f}(\theta, \varphi) \approx \sum_{i \in T} \left(\sum_{l=0}^{bw} \sum_{m=-l}^{l} c_{l,m}^i Y_l^m(\theta, \varphi) \right) \tag{11}$$

We introduce an error measure ε between the approximated signal $\hat{f}(\theta, \varphi)$ and the original signal $f(\theta, \varphi)$. ε is defined as follows :

$$\varepsilon = \sqrt{\sum_{j \in V} \left[f(\theta_j, \varphi_j) - \hat{f}(\theta_j, \varphi_j) \right]^2} \tag{12}$$

where V denotes the set of points of M. To have a more accurate precision, we can take V as the union of the set of points of M and the set of centroids of the triangles of M. The error ε can be calculated exactly, it depends on the value of bw. The greater the value of bw the smaller the value of ε. Restricting the error measure to the level of a triangle, we can make it as small as desired by increasing bw. The value of bw mainly depends on two factors :

- the distance d between the centroid of the triangle and the point c, Figure 1(a),
- the angle α between the normal of that triangle and the line connecting c and the centroid of the triangle, Figure 1(b).

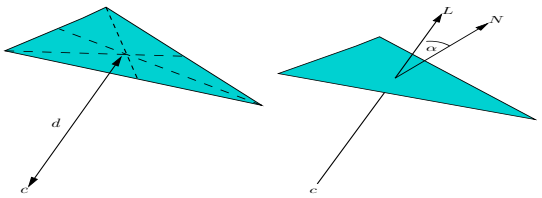

(a) distance (b) angle

Figure 1: The orientation and the distance of a triangle with respect to the point c.

Figure 2 gives a graphical analysis of the value of bw with respect to the two previous factors, for a fixed quality error ϵ. When the distance d increases, the projecting area of the triangle on the unit sphere decreases. Therefore, the bandwidth bw has to increase so as to compensate the distortion of the triangle due to this decreasing of the projecting area. In a similar manner,

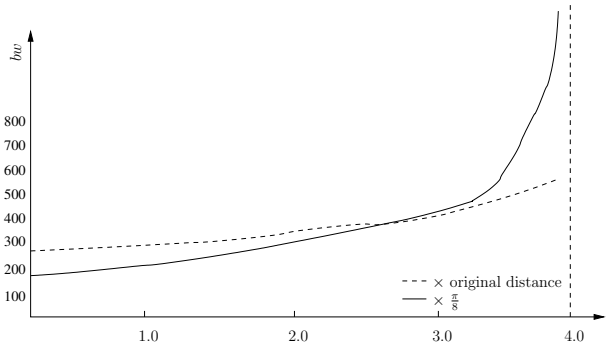

Figure 2: The effect of the distance and the orientation angle of a single triangle with respect to a fixed point c.

the bandwidth bw increases when the angle α increases. The triangle has a maximum distortion when $\alpha = \frac{\pi}{2}$. Figure 3 shows some examples representing some star-shaped models using the spherical harmonics. The bandwidth bw in this case has been fixed to 64 and the center c is the center of mass of the model. On those examples, we observe that $\varepsilon \leq 0.005 * \frac{D}{2}$, where D is the diagonal of the bounding box of the model.

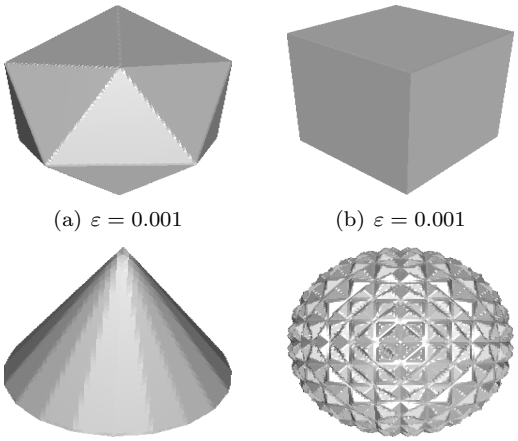

(a) $\varepsilon = 0.001$ (b) $\varepsilon = 0.001$

(c) $\varepsilon = 0.002$ (d) $\varepsilon = 0.005$

Figure 3: Star-shaped objects represented by the spherical harmonics transform of their radial functions, $bw = 64$.

4 FREQUENCY-BASED REPRESENTATION OF GENERAL MODELS

In this section, we extend our framework for representing general 3D models using spherical harmonics. This method consists of three main steps. In the first step, we segment the object into star-shaped parts using a robust segmentation technique [DGG03]. In the second step, we apply the spherical harmonics transform to each part separately as described in section 3. In the last step, we represent each filtered part as an implicit surface, and the whole object is obtained by blending together these implicit representations.

4.1 Segmentation

There are many decomposition techniques used to break complex models into convex or star-shaped sub-models. Lien et al. [LA04] have proposed a concavity measure to partition the models into approximately convex pieces, according to that measure. Dey et al. [DGG03] have decomposed the volume enclosed by a set of points into smaller sub-volumes. Their segmentation is based on topological persistence [ELZ00]. Firstly, they find the critical points of the distance function to the nearest point of the model and they classify them as maxima, minima and saddle points. Then they define what they call stable manifolds based on those maxima. Those stable manifolds are compact regions and decompose the interior volume of the model. Moreover, they are often convex regions.

In this paper, we have used this latter technique. Its advantage is that it also offers a good candidate for the choice of a center (the maximum attached to the region of manifold). However, since further fusions can be performed by their segmentation algorithm between the obtained parts, the choice of one center is less direct. This is the reason why we took the center of mass of the region instead. Nevertheless, this decision could be improved. After the choice of the center for each part, the latter is represented by its radial function $f : \mathbb{S}^2 \rightarrow \mathbb{R}$ with respect to this center.

Note that the initial object may be star-shaped with respect to a point c, but a great number of its triangles may not have a good orientation with respect to c, see Figure 4. So in this case, it is

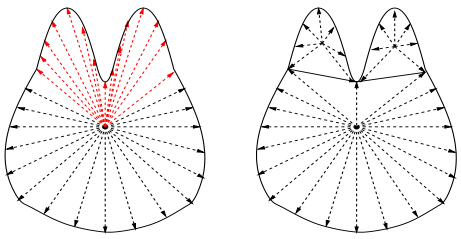

Figure 4: Left: bad orientation of some triangles (red) of a star-shaped object, right: improving the orientation of the triangles by segmenting the object.

recommended to decompose the object into simpler parts. This case can be detected by determining the ratio between the farthest and the nearest points of the object from c. If this ratio is greater than a threshold, then we decompose the object using the technique of Dey et al.[DGG03].

4.2 Conversion into implicit representation

The segmentation of the object results in a finite number of sub-objects, each of which can be represented by a radial function $f : \mathbb{S}^2 \rightarrow \mathbb{R}$ with respect to a center. Consider that $\{M_i : \ i = 1, \ldots, n\}$ is the set of sub-parts and let $\{f_i(\theta, \varphi) : \ i = 1, \ldots, n\}$ denote their radial functions with respect to the set of points $\{c_i : \ i = 1, \ldots, n\}$ respectively. Recall that each $f_i(\theta, \varphi)$ measures the extent of M_i from c_i in the direction (θ, φ).

Consider $\{g_i(r, \theta, \varphi) \ i = 1, \ldots, n\}$ as the set of functions defined as the following : for each point $q \in \mathbb{R}^3$ whose coordinates with respect to c_i are $(r_q^i, \theta_q^i, \varphi_q^i)$

$$\begin{array}{rcl} g_i(r_q^i, \theta_q^i, \varphi_q^i) & = & f_i(\theta_q^i, \varphi_q^i) - d(c_i, q) \\ & = & d(c_i, p) - d(c_i, q) \end{array} \quad (13)$$

where p is the intersection point of M_i with the line $\overrightarrow{c_i q}$, and d is the euclidean distance. Each g_i has the following property :

$$sign(g_i(q)) = \begin{cases} + & if\ q\ is\ inside\ M_i, \\ - & if\ q\ is\ outside\ M_i, \\ 0 & if\ q\ is\ on\ M_i. \end{cases} \quad (14)$$

Therefore, the surface of M_i can be considered as the 0-level of the potential function g_i. The volume of the whole object M is the union of the volumes of these implicit surfaces. The theory of R-functions [PS95, Rva87] provides a useful set of operations on the potential functions. The union operation of two potential functions g_1 and g_2 is defined as follows :

$$g_1 \cup g_2 = \frac{1}{1+a}\left(g_1 + g_2 + \sqrt{g_1^2 + g_2^2 - 2a g_1 g_2}\right) \quad (15)$$

where $a(q) = s(g_1(q), g_2(q))$, s is an arbitrary continuous function satisfying the following condition:

$$-1 < s(t_1, t_2) \leq 1$$

The max function is a special case with $s = 1$. Let M_1 and M_2 be two neighboring parts represented by the radial functions f_1 and f_2 corresponding to the centers c_1 and c_2, respectively. Let g_1 and g_2 denote the corresponding potential functions induced from these radial functions. Therefore, before applying the spherical harmonics transform, we have for any point $q \in \mathbb{R}^3$:

$$g_1(q) = f_1(\theta_q^1, \varphi_q^1) - d(c_1, q) \quad (16)$$

$$g_2(q) = f_2(\theta_q^2, \varphi_q^2) - d(c_2, q) \quad (17)$$

Using equation 15, the potential function g representing the union of the parts represented by g_1 and g_2 is defined as the following:

$$g = g_1 \cup g_2 \quad (18)$$

Let $\hat{f}_1$ and $\hat{f}_2$ denote the spherical harmonics transform of f_1 and f_2, respectively. Therefore after restricting the frequencies to a bandwidth bw, we have for any point $q \in \mathbb{R}^3$:

$$\hat{g}_1(q) = \hat{f}_1(\theta_q^1, \varphi_q^1) - d(c_1, q) \quad (19)$$

$$\hat{g}_2(q) = \hat{f}_2(\theta_q^2, \varphi_q^2) - d(c_2, q) \quad (20)$$

The potential function $\hat{g}$ representing the union of the parts represented by $\hat{g}_1$ and $\hat{g}_2$ may have unsmoothness along the boundaries shared by the parts due to restricting the frequencies to a bandwidth bw, as shown in Figure 5(a). To overcome

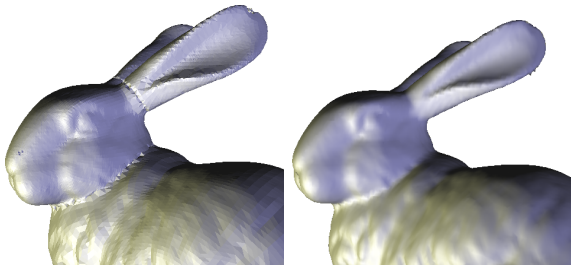

(a) (b)

Figure 5: Left: unsmoothness along the segmentation boundaries, right: using the blending.

this problem, we apply the set-theoretic blending operator based on R-functions [PS94, PS95]. The corresponding blending operator of the two potential functions $\hat{g}_1$ and $\hat{g}_2$ is defined as follows :

$$\hat{g}_1 \oplus \hat{g}_2 = R(\hat{g}_1, \hat{g}_2) + u \quad (21)$$

where R is the corresponding R-function (the union in our case), and $u(q) = w(g_1(q), g_2(q))$, w is a displacement function that has a maximal absolute value $w(0, 0)$; i.e. at the boundary, and asymptotically approximates a zero value with increasing absolute values of the arguments. The general form of w is as follows [PS94, PS95, Rva87]:

$$w(t_1, t_2) = \frac{1.0}{1 + (t_1/a_1)^2 + (t_2/a_2)^2} \quad (22)$$

where a_1 and a_2 are constants which control the blending operator. We need to choose them optimally with respect to the object.

The error between $g = g_1 \cup g_2$ and $\hat{g}_1 \oplus \hat{g}_2$ is defined as follows:

$$\varepsilon = \sqrt{\sum_{v \in V} (g(v) - (\hat{g}_1 \oplus \hat{g}_2)(v))^2} \quad (23)$$

where V is the set of vertices. By varying the values of the constants a_1 and a_2, we can consider ε as a function of these values. So the objective is to minimize the error function $\varepsilon(a_1, a_2)$. In this paper, we use the genetic algorithms [Gol89] as

a minimization tool on $\varepsilon(a_1, a_2)$ to choose heuristically the two parameters a_1 and a_2 to remove the unsmoothness along the boundary of the parts. This method gives very good results as we can see in Figure 5(b).

5 APPLICATION : FILTERING

One of the important points of our frequency-based representation is that it allows to filter the surface of n-genus 3D object easily. Surface filtering is useful in smoothing and noise removal [Tau95, GSS99, KG00]. The underlying assumption is that high order frequencies correspond to noises or finer details of the surface. Therefore, removing those frequencies yields a removing of noises or some finer details of the surface. Here, the required filter is applied separately to each triangle of the object and the results are combined as explained previously. Zhou et al. [ZBS04] have presented some interesting frequency filtering functions $h(l, m)$. For example, the ideal low-pass filtering can be performed by setting :

$$h(l,m) = \begin{cases} 0 & if\ \sqrt{l^2+m^2} > K_l \\ 1 & otherwise \end{cases} \qquad (24)$$

Figure 6 shows smoothing the surface of Armadillo using the previous low-pass filter.

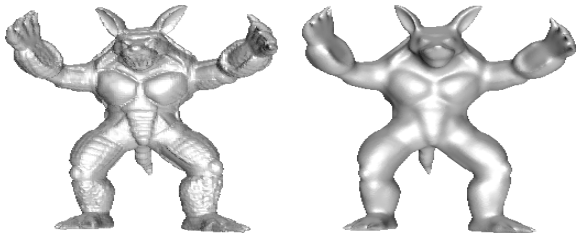

(a) original (b) $K_l = 60$

Figure 6: Smoothing the surface of Armadillo using ideal low-pass filter.

Additionally, surface filtering is also useful for compression and progressive transmission of 3D models. The underlying assumption is that a relatively good approximation may be obtained using only a small number of low-frequency basis functions, see Figure 7. That is, we can send a small number of low-frequency coefficients through the network and progressively send the higher ones to have finer details. Figure 7 shows levels-of-details of Happy Buddha. These levels-of-details can be considered as compressed versions of the model. For example, Figures 7(b) and 7(c) have compression ratios 87% and 99.3% respectively with regards to an initial OFF (Object File format) representation, and without further compression of the obtained sequence of coefficients.

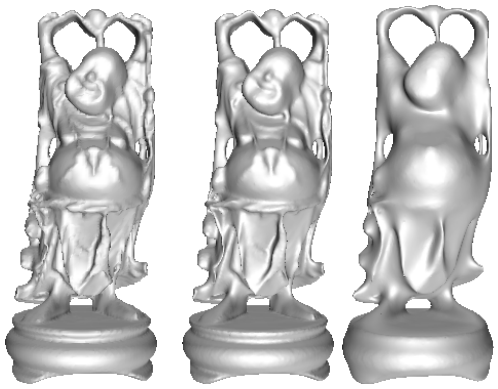

(a) original (b) $K_l = 256$ (c) $K_l = 60$

Figure 7: Several levels-of-details of Happy Buddha using ideal low-pass filter. (a) corresponds to 543,652 points described by doubles and 1,087,716 faces, (b) and (c) correspond to $\frac{(k_l)^2}{2}$ coefficients described by complexes. The compression ratios for (b) and (c) are 87% and 99.3% respectively.

6 EXPERIMENTAL RESULTS

Our method is implemented in C++; we have used a 3GHz Pentium IV PC with 1GB memory for the experiments. The input meshes are considered triangulated. Otherwise, a preprocessing step is required to triangulate the polygons of the mesh. The segmentation step does not take more than 4 minutes for each of the models used in this paper. Generally, the number of parts is in average 50 parts. We have visualized our result by reconstructing the surfaces of the models using the marching cube algorithm proposed by Lewiner et al. [LLVT03]. Table 1 shows a summary of the experimental results.

Model	no. of triangles	SHT time	no. of parts
Bunny	69,451	3min	20
Triple Hecate	180,364	5min	30
Victoire	187,072	6min	50
Buddha	1,087,716	8min	50
Armadillo	345,944	7min	50
Hand	654,666	5min	50

Table 1: Summary of the models used in this paper.

Figure 8 presents some examples of general models, represented using spherical harmonics frequencies. Each model is segmented into subparts, each of which is transformed into a set of frequencies using the spherical harmonics. Taking $\varepsilon \leq 0.01 * \frac{D}{2}$, where D is the diagonal of the bounding box, the corresponding bandwidths are 128, 256, 256, 128, 256 and 256 for the Bunny, Armadillo, Happy Buddha, Hand, Triple Hecate and Victoire respectively.

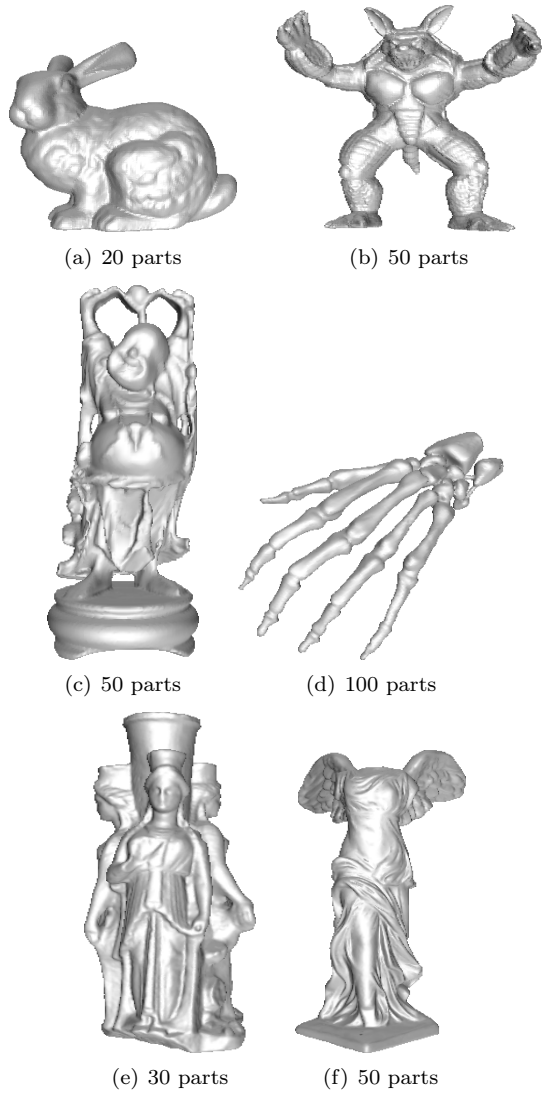

(a) 20 parts (b) 50 parts

(c) 50 parts (d) 100 parts

(e) 30 parts (f) 50 parts

Figure 8: Some models represented using the spherical harmonics. The bandwidth bw has the following values (a) 128 (b) 256 (c) 256 (d) 128 (e) 256 and (f) 256.

7 CONCLUSION AND FUTURE WORK

This paper presents a new technique to represent general 3D models using the spherical harmonics. This representation allows us to filter the surfaces of these models, although they are not topologically equivalent to the sphere, and to describe them compactly. The spherical harmonics transform is computed using a fast, combinatorial and robust algorithm [MCA]. The implicit framework techniques guarantee the avoidance of unsmoothnesses on the surfaces.

Concerning the segmentation technique, the technique proposed in [DGG03] is robust and yields a good segmentation of the objects. However, we need as well to have an optimal segmentation technique that produces approximately convex (or star-shaped) sub-parts while keeping the number of parts as small as possible.

On the other hand, we have chosen the center of each part as the center of the mass of that part. Since the radial function of each part is dependent on the choice of the center, the final representation is affected by this choice. We work to improve this choice of the center.

ACKNOWLEDGEMENT

We acknowledge Dey et al. for providing us with their code of the segmentation algorithm. We also acknowledge the support of ART3D project funded by the french ministry of research. The model of Triple Hecate is one of the statues disponible at the Louvre Museum [Lou]. Thanks to the Stanford graphics laboratory website and the large geometric models archive of Georgia Institute of Technology for putting many models at disposal.

REFERENCES

[Bye59] W. E. Byerly. *Spherical Harmonics*, chapter 6, pages 195–218. New York: Dover, 1959. An Elementary Treatise on Fourier's Series, and Spherical, Cylindrical, and Ellipsoidal Harmonics, with Applications to Problems in Mathematical Physics.

[Bül02] T. Bülow. Spherical diffusion for 3d surface smoothing. In *3DPVT'02: the first International Symposium on 3D Data Processing Visualization and Transmission*, page 449, Padova, Italy, 2002.

[CDR00] U. Clarenz, U. Diewald, and M. Rumpf. Nonlinear anisotropic diffusion in surface processing. In T. Ertl, B. Hamann, , and A. Varshney, editors, *Proceedings of IEEE Visualization 2000*, pages 397–405, 2000.

[DGG03] T. K. Dey, J. Giesen, and S. Goswami. Shape segmentation and matching with flow discretization. In F. Dehne, J.-R. Sack, and M. Smid, editors, *WADS '03: Proceedings of the* 8^{th} *International Workshop on Algorithms and Data Structures*, number 2748 in LNCS, pages 25–36, Carleton Univ., Ottawa, Canada, July-August 2003. Springer Verlag.

[DMSB99] M. Desbrun, M. Meyer, P. Schroder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 317–324, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[EDD+95] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 173–182, New York, NY, USA, 1995. ACM Press.

[ELZ00] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. In *FOCS '00: Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, page 454, Washington, DC, USA, 2000. IEEE Computer Society.

[FMK+03] T. Funkhouser, P. Min, M. Kazhdan, J. Chen, A. Halderman, D. Dobkin, and D. Jacobs. A search engine for 3d models. *ACM Transactions on Graphics*, 22(1):83–105, January 2003.

[Gol89] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Pub. Co., 1989.

[GSS99] I. Guskov, W. Sweldens, and P. Schroder. Multiresolution signal processing for meshes. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 325–334, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[Hob55] E. W. Hobson. *The Theory of Spherical and Ellipsoidal Harmonics*. New York: Chelsea, 1955.

[JDBP04] J. Jin, M. Dai, H. Bao, and Q. Peng. Watermarking on 3d mesh based on spherical wavelet transform. *Journal of Zhejiang University Science*, 5(3):251–258, 2004.

[KFR03] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In *SGP '03: Proceedings of the Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164. Eurographics Association, 2003.

[KFR04] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Symmetry descriptors and 3d shape matching. In *SGP '04: Symposium on Geometry Processing*, pages 116–125, July 2004.

[KG00] Z. Karni and C. Gotsman. Spectral compression of mesh geometry. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 279–286, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[LA04] Jyh-Ming Lien and Nancy M. Amato. Approximate convex decomposition. In *SCG '04: Proceedings of the twentieth annual symposium on Computational geometry*, pages 457–458, New York, NY, USA, 2004. ACM Press.

[LLVT03] T. Lewiner, H. Lopes, A. Wilson Vieira, and G. Tavares. Efficient implementation of marching cubes cases with topological guarantees. *JGT (Journal of graphics tools)*, 8(2):1–15, 2003.

[Lou] The Louvre Museum. http://www.louvre.fr/.

[MCA] M. Mousa, R. Chaine, and S. Akkouche. Direct spherical harmonic transform of a triangulated mesh. *JGT: Journal of Graphics Tools*. to appear.

[PS94] A. Pasko and V. Savchenko. Blending operations for the functionally based constructive geometry. In *CSG 94: Set-theoretic solid modelling–Techniques and Applications*, pages 151–161, Winchester, UK, April 1994.

[PS95] A. Pasko and V. Savchenko. Constructing functionally defined surfaces. In B. Wyvill and M. P. Gascuel, editors, *Implicit Surfaces '95: The first international workshop on implicit surfaces*, pages 97–106, Grenoble, France, April 18–19 1995.

[Rva87] V. L. Rvachev. *Theory of R-functions and some applications*. Naukova Dumka, Kiev, 1987. (in Russian).

[SS01] W. Sweldens and P. Schröder. Digital geometric signal processing, course notes 50. In *SIGGRAPH 2001 Conference Proceedings*, 2001.

[SV01] D. Saupe and D. V. Vranic. 3d model retrieval with spherical harmonics and moments. In *Proceedings of the 23rd DAGM-Symposium on Pattern Recognition*, pages 392–397. Springer-Verlag, September 2001.

[Tau95] Gabriel Taubin. A signal processing approach to fair surface design. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 351–358, New York, NY, USA, 1995. ACM Press.

[ZBS04] K. Zhou, H. Bao, and J. Shi. 3d surface filtering using spherical harmonics. *Computer-Aided Design*, 36(4):363–375, 2004.

[ZC01] C. Zhang and T. Chen. Efficient feature extraction for 2d/3d objects in mesh representation. In *ICIP '01: Proceedings of the International Conference on Image Processing*, pages 935–938, October 2001.